



HAL
open science

Intellectual affectivism : intuition experiences are epistemic feelings

Wjatscheslaw Loev

► **To cite this version:**

Wjatscheslaw Loev. Intellectual affectivism : intuition experiences are epistemic feelings. Philosophy. Université Paris sciences et lettres, 2019. English. NNT : 2019PSLEE065 . tel-03159784

HAL Id: tel-03159784

<https://theses.hal.science/tel-03159784>

Submitted on 4 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE DE DOCTORAT

DE L'UNIVERSITÉ PSL

Préparée à Institut Jean Nicod, UMR8129

Intellectual Affectivism: Intuition Experiences are Epistemic Feelings

Soutenue par

Wjatscheslaw LOEV

Le 18 12 2019

Ecole doctorale n° 540

**Lettres, Arts, Sciences
humaines et sociales**

Spécialité

Philosophie

Composition du jury :

Kathrin, GLÜER-PAGIN Full Professor, Stockholm University	<i>Président</i>
Bénédicte, VEILLET Associate Professor, University of Michigan – Flint	<i>Rapporteur</i>
Fabrice, TERONI Associate Professor, Université de Genève	<i>Rapporteur</i>
John, BENGSON Associate Professor, University of Wisconsin–Madison	<i>Examineur</i>
Elisabeth, PACHERIE Directrice d'études, CNRS	<i>Directrice de thèse</i>
Jérôme, DOKIC Directeur d'études, EHESS	<i>Co-Directeur de thèse</i>

Institut | Nicod

CNRS-EHESS-ENS

Intellectual Affectivism

Intuition Experiences are Epistemic Feelings

Wjatscheslaw (Slawa) Loev

Doctoral Thesis supervised by:

Elisabeth Pacherie

Jérôme Dokic

Institut Jean Nicod
DEC, ENS, PSL, Paris

February 27, 2020

to Suze

Acknowledgements: This PhD thesis, however neatly structured it might appear now, has a rich and messy history of which only a tiny, highly polished fraction has made it into the present version. The very format of this thesis is, in fact, somewhat misleading: it allows only for one place and one name on it. However, the story of this thesis features a substantial number of places, some of which I have learned to call my home. More importantly, the story involves countless characters, with names known and unknown, and with powers ranging from the mundane to borderline supernatural. One of the most significant aspects of this thesis' story is about how my relationships to these people have flourished throughout; about how these people have supported and inspired me in an unfathomable number of ways; and about how they have become (or been) my friends and family. In the following I want to thank those whose names I remember (in the moment of writing this).

I want to thank my supervisors, Elisabeth Pacherie and Jérôme Dokic. Thank you for guiding and supporting me along the way. It was an honour and pleasure to work with you.

Special thanks goes to the other members of my defense committee: Kathrin Glüer-Pagin, Bénédicte Veillet, Fabrice Teroni, John Bengson and François Recanati. Your feedback was invaluable and the real-time interaction with you made my defense a real growth moment, apart from being a lot of fun.

I am very grateful to Sven Rosenkranz, Chiara Panizza and the other Diaphora-members for putting together and efficiently managing the Diaphora Early Stage Network from whose generous PhD fellowship I profited aplenty.

I also would like to thank the team of NGO *Search for Common Ground*, especially Omar Salem and Fien De Baere, for giving me a glimpse into the fascinating world of NGO-work.

Many thanks to Nathalie Evin Abitbol and Vincent Gaudefroy for their unceasing help with the inevitable practicalities of a PhD life (in Paris!).

For helpful advice and discussion I thank Miguel Egler, Joëlle Proust, Uriah Kriegel, Pepa Toribio, Manolo Martinez, Pierre Jacob, Laura Candiotto, Till Vierkant and the participants of numerous conferences, workshops and talks.

For their unwavering support, care, joyful presence, the innumerable beautiful moments and joint adventures and, of course, many many conversations about many many things, including topics concerning my thesis, I extend my infinite gratitude to: Andrés Soria Ruiz, my good, ol' Spanish friend in Paris; Sandra Lasry, my PhD-supervisor-sister; Tricia Magalotti, my Diaphora-sister; Pablo Fernández Velasco, the best collaborator I could wish for; Marco Inchingolo, my favourite Italian person; Antoine Marie, my favourite source of unsolicited fashion advise; Mauricio Martins, my long-lost, sciencey older brother from another mother; Nura Sidarus, who taught me what assertiveness really is; Matheus Valente, my Brazilian role model in Proustian perceptiveness; Yaron Wolf; Anna Giustina; Agostino Pinnapintor; Takuya Niikawa; Martina Cirese; Gabriel Darmon Paes; Clothilde Facon; Céline Boisserie-Lacroix; Angel Hyeyoung Kim; Guido Robin Löhr; Dan Hoek; Hualin Xiao; Boian Nikiforoff-Kharisanoff Vega; Andreas Falck; Zuzana Pundova; Anouk Hajage; François Kammerer; Nils Franzén; Ari Fabrizi De' Biani; Dora Fotopoulou; Nicolas Rogge; Edoardo Genova; Deb Marber; Lisa Vogt; Maddy Hyde; Aïda Elamrani-Raoult; Louis Rouillé; Géraldine Carranante; Tristan Thommen; Santiago Arango-Muñoz; Paula Marchesini; Isi Assalin; Victor Fernández Castro; Manvir Singh; Amir Ahmad Anvari; Andrew Lee; Elena Ayza Aschmann; Agathe Huet; Benjamin Maise; and Mora Maldonaldo.

I also want to extend my thanks to a couple of my good longer-standing friends who, incidentally, all came to share my PhD fate. Leonid Berov, Jan Dirk Capelle, Lisa Steinmetz, Antonia Kaiser and David Richter, thank you for simply continuing to be there for me.

A big thanks for being my bedrock of support (if regrettably sometimes only at a distance) goes to my (extended) family: Olga Loev; Oleg Loev, Natalie Moliker, Marlies Rouweler, André van Scheers, Rutger van Scheers and Marie-Jose Rouweler.

Finally, my greatest thanks goes to my partner (in crime), Suze Lian van Scheers, who I met shortly before starting my PhD and who got to see and support me from beginning to end. Without you this PhD life would have been much more solitary, poor, nasty, brutish, and (perhaps) short.

Abstract: The guiding question of the present thesis is: “What kind of states are intuitions?” The answer developed here is *Intellectual Affectivism* or Affectivism (about intuitions). Affectivism claims that intuitions are affective experiences, or more precisely: they are specific instances of epistemic feelings, feelings of truth and feelings of falsity. The thesis proceeds as follows: First, the target state of which Affectivism is a theory is delineated from other things we call “intuition”. Then the feature profile of “intuitions” in the target sense of intuition experiences is outlined: Intuition experiences are occurrent conscious mental states that are (partially) characterised by their characteristic (but not necessarily *sui generis*) phenomenology. They are furthermore 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced (i.e. there are positive intuitions concerning truth and negative intuitions concerning falsity) and 7) nonvoluntary. It is argued that this feature profile needs to be accommodated by a good theory of intuition experiences, i.e. a good answer to the guiding question should be able to acknowledge and explain these features. Extant intuition theories provide the following answers: *Eliminativism* claims the term “intuition” has no extension—intuitions do not exist. *Doxasticism* claims intuitions are doxastic states. *Perceptualism* claims intuitions are similar to perceptual experiences. It is shown that all the existing answers are unsatisfactory. Either they cannot acknowledge the features of intuitions or they cannot explain them (or both). The rest of the thesis is dedicated to the development of a new intuition theory: Intellectual Affectivism. The answer it gives to the guiding question is the following: intuitions are affective experiences, or more precisely: they are specific instances of epistemic feelings, feelings of truth and feelings of falsity. To motivate this answer, the psychological kind of feelings or affective experiences is introduced and characterised: affective experiences, of which bodily feelings such as bodily pain or pleasure and emotional feelings such as fear or joy are paradigmatic subclasses, are valenced, arousing, motivational and richly intentional by engaging in a division of representational labour with other mental states. Then the class of epistemic feelings is introduced and characterised. The thesis proceeds to make a case for epistemic feelings being affective experiences. Having established that, it goes on to identify and analyse specific epistemic feelings as promising candidates for an identification with intuition experiences: feelings of rightness and feelings of wrongness. It turns out that a propositional variety of these feelings, feelings of truth and feelings of falsity, has the same feature profile as intuition experiences. Thus, the claim goes, positive intuitions are to be identified with feelings

of truth and negative intuitions are to be identified with feelings of falsity. In virtue of these feelings being affective experiences, Affectivism cannot only acknowledge the features of intuitions but also explain them. Intuitions have the features they have for essentially the same reasons as bodily and emotional feelings have them—because they are (specific) affective experiences. Before concluding on the implications of Affectivism, the remainder of this thesis makes a first exploration of the relationship between feelings of truth and falsity and actual truth and falsity.

this page intuitively left blank

Contents

1. Introduction: Intuitions about Intuitions	12
2. The Marks of Intuition Experiences	23
2.1. Introduction: Acknowledging Intuitions	23
2.2. Intuition Experiences and Their Features	27
2.2.1. Intentionality	29
2.2.2. Assertiveness	33
2.2.3. Pushiness	35
2.2.4. Being Noncommittal and Related Features	36
2.2.5. Gradeability: Content and Pushiness	39
2.2.6. Phenomenal Epistemic Valence	40
2.2.7. Nonvoluntariness	43
2.3. Recap and a Phenomenally Contrastive Mixed List	43
3. Intuitions and Where (Not) to Find Them: Extant Theories of Intuitions	46
3.1. Introduction: Intuitions for Philosophy	46
3.2. No Intuitions?	49
3.3. Intuitions and Doxastic States	51
3.3.1. Intuitions as Doxastic Attitudes	52
3.3.2. Intuitions as Doxastic Tendencies	61
3.4. Perceptualist Intuitions	66
3.4.1. Classic Perceptualism: Intuitions as Intellectual Seemings	67
3.4.2. Advanced Perceptualism: Intuitions as Intellectual Presentations and Rich Inclinations	70
3.4.3. Issues with Perceptualism (and Doxasticism)	74
4. Looking for Intuitions Elsewhere: Appealing to Feelings	81
4.1. Introduction: Where to go from here?	81

4.2.	Clarifying Affective Experiences	84
4.3.	Affective Phenomenology	89
4.3.1.	The Hallmark of Affect: Phenomenal Valence	90
4.3.2.	Other Hallmarks of Affect: Felt Arousal and Motivation	94
4.4.	Affective Intentionality	98
4.4.1.	The Particular Objects of Affective Experiences	99
4.4.2.	The Formal Objects of Affective Experiences	102
4.4.3.	Interjecting Remarks on Valence and Intentionality	106
4.4.4.	The Bases of Affective Experiences	108
4.5.	Affective Experiences and Intuitions	113
5.	Epistemic Feelings	121
5.1.	Introducing Epistemic Feelings	121
5.2.	Getting a Feel for Epistemic Feelings	123
5.3.	The Epistemic Feeling Machinery	130
5.3.1.	The FOK (and Co.) Machinery	131
5.3.2.	The FOF Machinery	141
5.3.3.	Epistemic Feelings and Metacognition	145
5.4.	Chapter Recap	147
6.	Epistemic Feelings are Affective Experiences	150
6.1.	Introduction: Reinforcing a Case	150
6.2.	Affective Markers of Epistemic Feelings	151
6.3.	Misplacing Affect in Interesting Ways	155
6.3.1.	Constitutive Affect in Epistemic Feelings	156
6.3.2.	Conscious Affect in Epistemic Feelings	158
6.4.	Understanding Epistemic Feelings better as Affective Experiences	161
6.5.	Epistemic Feelings as Affective Experiences	169
6.5.1.	The Mild Affective Phenomenology of Epistemic Feelings	169
6.5.2.	The Affective Intentionality of Epistemic Feelings	175
6.6.	Chapter Recap	178
7.	Intuition Experiences are Epistemic Feelings	181
7.1.	Introduction: Looking Back and Ahead	181
7.2.	Getting a Feel for Rightness and Wrongness	183
7.3.	The Science of FOR/Ws	187
7.3.1.	The Paucity of FOR/W Research and Ways of Being Right or Wrong	187

7.3.2.	A First Look on FOR/W Research Through a Distinction	189
7.3.3.	Liking Rightness	195
7.3.4.	The Truth Effect	197
7.3.5.	Truth, Fluency and Affect	202
7.4.	FOR/Ws as Affective Experiences and Intuitions as FOT/Fs	209
8.	The Mechanisms of Truth: Of Fluency, Coherence and Comprehension	221
8.1.	The Mechanisms behind Truth and Falsity	221
8.2.	The Status of Fluency	222
8.3.	Getting to the Truth	231
8.3.1.	Coherent References and Truth	231
8.3.2.	Epistemic Monitoring	238
8.3.3.	The Truth in Intuitions	241
9.	Concluding Affectivism: The End and the Beginning	246
9.1.	Recounting the Journey	246
9.2.	Back to the Future of Affectivism	248
A.	A Brief History of Affectivism	257
A.1.	Chapter 2	257
A.2.	Chapter 3	258
A.3.	Chapter 4	260
A.4.	Chapter 5	261
A.5.	Chapter 6	263
A.6.	Chapter 7	265
A.7.	Chapter 8	267

Chapter 1

Introduction: Intuitions about Intuitions

A Curious Sort of Animal

What are intuitions? Few things are sure about the answer to this question. However, one is: intuitions are a curious sort of animal. We all have intuitions and we all talk about intuitions. But how often do we think about what it is that we have and talk about—about what intuitions are? Well, I assume most of us have an intuitive idea. At the same time, this intuitive idea likely differs quite dramatically depending on one’s background. Consequently, things one expects from an answer to the question will vary. If you are a layperson (about intuitions), you may expect to find something about feelings and the infamous “Sixth Sense”. If you are a philosopher, you might expect to find something about intellectual activity, belief and truth. And if you are a psychologist, you probably expect to find something about unconscious processes and mechanisms. Here is an announcement: All these things will be featured in the exhibition to come—yes, even the Sixth Sense.¹

However, the way they feature here will at times not be the way everybody has expected. Insofar, there will be some surprises: The philosophers will perhaps be perplexed by the fact that this thesis is not (primarily) about the epistemology of intuitions.² To get clear on this from the start. This thesis is *really* on the question: What are intuitions? This is, it is about the ontology, metaphysics or psychology of intuitions rather than on their

¹See section 3.4.

²See section 3.1.

epistemology. Don't get the "metaphysics" wrong here. However, this thesis will not be too much on the Sixth Sense either—hopefully to nobody's chagrin. What it will rather be about are unconscious processes and mechanism. First and foremost, however, it will be about feelings. About feelings and how they relate to the other mentioned things such as unconscious processes, intellectual inquiry, truth and belief. One of the upshots will be that the elements that laypersons, philosophers and psychologists bring to the table are not as far apart as would appear to some at first sight. The various elements concerning intuitions are related—through feelings.

Thus, the most obvious purpose of the present thesis is simply to state a theory of what intuitions are. This theory essentially consists in developing the — perhaps for some unsurprising or even trivial — idea that *intuitions are feelings*. I call this theory *Intellectual Affectivism* or, in short, *Affectivism* (about intuitions). And it is supposed to be a philosophical theory of intuitions that is empirically informed. Insofar the present thesis is also a first stab on a *cognitive science* of intuitions, trying to go beyond "just" philosophy or "just" psychology.

As I said, this theory is also supposed to bring the elements mentioned by the layperson, philosopher and psychologist *together*. This is significant insofar as (at least) two of these elements are often seen in opposition to each other: feelings and intellectual activity. So, on a bigger picture, this thesis is also an attempt to rehabilitate feelings into the intellectual realm, a realm from which they have been often expelled. This exile is motivated by a suspicion: the suspicion that feelings obstruct intellectual activity. But now, if intuitions are feelings and if intuitions play central roles in our intellectual activity, then, it turns out that feelings play central roles in our intellectual activity. And seemingly they have done so all along—sneakily. As a participant in intellectual activity, philosophy, say, it might thus be time to reconsider one's stance on feelings. They might turn out to not only *somehow contribute* to intellectual activity but to *be* some of its very building blocks: intuitions.

On the other hand, one of the reasons why feelings have acquired their dubious reputation is their occasional opacity. What's going on when we have feelings and why do we have them? Why can't they be more like perception, every epistemologist's darling and alleged paragon of transparency? Interestingly, while this occasional opacity of feelings appears to have scared philosophers away, it has inspired the interest of psychologists, eager to explore the mentioned unconscious processes and mechanisms that bring them about. This has led the disciplines to diverge: psychological research increasingly demonstrates

how little we are actually *aware* of what’s going on *in general*, not only when we have intuitions, feelings or are engaged in intellectual activity. More and more we discover that, actually, there’s not *so much* going on consciously and that the mind’s heavy lifting is—for better or worse—going on below the radar of consciousness. For the most part, consciousness is more the request, relay, output and quality control station of complex operations “in the basement” rather than the data crunching machinery itself. That does not mean that we cannot (to some degree) reconstruct what’s going on in the basement—but it is nevertheless a reconstruction.

Now while psychologists tend to outsource a lot of the operational business from consciousness, philosophers do rather the opposite and serve consciousness a hefty workload. Having said that, we should not confuse the explicit contents and operations found in a well-argued piece of intellectual prose with the contents and operations found in the consciousness of the author. Surely, most of the contents that ended up in the piece have *passed* through consciousness—but they have not been conceived in it, not in any transparent way. We might be very good in (swiftly) reconstructing or naming these operations or steps but we should not deceive ourselves into thinking that they took place consciously. Even if many things seemed immediate and obvious, clear and distinct to the author, for a good part, she has been led by her intuitions—her feelings, not by operations in her conscious mind. Conceiving of intuitions as feelings emphasizes that they are the conscious outputs of complex and opaque operations, “sent up” to guide conscious (intellectual) activity.

Perhaps I am belabouring this point. Perhaps, in fact, I am mistaken. Perhaps nobody thinks that intellectual activities such as reasoning mostly *take place in* consciousness and I am railing at a Pre-Freudian Strawman or an overly naïve version of myself in the past. Or perhaps nobody, not even me, has ever thought that. But if not, then this thesis is also a reminder not to slip into this way of thinking.

Enough of the “big picture”—let’s get down to the theory. The central claim of Intellectual Affectivism is that intuitions are feelings.³ The “Intellectual” in Intellectual Affectivism refers to intuitions. This is because intuitions are supposed to play important roles in our *intellectual* (and epistemic) activity (see e.g. section 3.1). The “Affectivism” in Intellectual Affectivism stands for the claim that intuitions are feelings, i.e. that they are *affective* in nature. More specifically Affectivism states that: intuitions are (specific) epistemic feelings (which themselves are feelings). However, in this introduction I will

³Or: intuitions are *affective experiences*. I use “feelings” synonymously with “affective experiences” and elaborate on this use in chapter 4, specifically in section 4.2.

use “intuitions are feelings” as a shorthand for that. This should suffice for now in terms of precision.

The bulk of this thesis will be dedicated to *developing* Intellectual Affectivism. That is, it will mostly consist of paving the way to and arriving at Affectivism about intuitions. However, here and there we will also explore some implications of the theory and see how it can be applied to some problems. The most straightforward implication of Affectivism is also one of its biggest advantages: it can not only acknowledge the features of intuitions but also go some way in explaining why intuitions have these features: because intuitions are feelings. The features that feelings have can, in turn, be explained by what we know about feelings, using the conceptual and empirical resources that we have for feelings (e.g. concerning their unconscious determinants). Arguably, we have much more of these resources for feelings than we have for intuitions. So by identifying intuitions with feelings, we can use the resources that we have for feelings to explain intuitions and their features. I think this is progress.

Julien Deonna has pointed out a similar benefit when it comes to Perceptualist Theories of Emotions:

It has been argued that emotions are very much like perceptions or even that they are kinds of perceptions. [...] The appeal of such a view, if true, is clear enough. It is the promise of illuminating a more elusive phenomenon by means of one that is slightly more familiar. (Deonna 2006, p. 29)

Note, however, that the explanatory power gained from saying that intuitions “are very much like” something, i.e. (very) similar to something, is radically different from the explanatory power gained from saying that intuitions *are* something (or a subclass of it). So to be clear: I am *not* saying here that intuitions are similar to feelings. I am saying that intuitions *are* feelings.

Let me elaborate on the advantage of Affectivism about intuitions in comparison to taking intuitions to be *similar* to feelings: The appeal of Intellectual Affectivism rests on the natural idea that many observations (e.g. its characteristics) about X can be explained in virtue of classifying it as an instance of A and in virtue of identifying it with Y, if one knows more about A and Y than about X. To illustrate: many observations and characteristics of, say, Samuel can be explained in virtue of Samuel being a representative of homo sapiens and in virtue of homo sapiens being a subclass of homo, of homo being a subclass of primates, of primates being a subclass of mammals etc. Due to this fact

about Samuel we can explain that Samuel is warm-blooded, has a remarkably large and resource-hungry brain and opposable thumbs, uses all kinds of tools, (usually) locomotes bipedally and has command of an astonishingly sophisticated system of communication. We can explain even more about Samuel by identifying him — apart from being an American etc. — by his pseudonym Mark Twain (it turns out that Samuel’s last name is Clemens). Knowing this about Samuel, we can also go some way in explaining that he sometimes likes to seclude himself in his study, that he is a well-known writer and that he lost a lot of money in poor investments.

Consider also how much we can explain about, say, Donald in virtue of him being an instance of *homo sapiens*. Now, for illustrative purposes, imagine one day Donald takes a walk. This walk leads him into the marshes. Suddenly, a thunderstorm comes up. And — maybe you can see it coming — poor Donald is hit by lightning (it turns out that Donald’s last name is Davidson). However, Donald is not pulverised — he survives and, surprisingly, emerges virtually unscathed. Something that puts a significantly heavier load on Donald’s psyche is what happens next: A second lightning hits the tree next to him. This results in that—by sheer chance, of course—the molecules of the tree are rearranged so as to create a being that Donald recognizes as looking quite similar to himself (how similar exactly, I leave up to your imagination). In fact, this “Swampman”, as baptised by Donald, seems to have opposable thumbs, locomote bipedally and to make utterances that sound suspiciously similar to American-accented English—these are, at least, the observations of flabbergasted Donald (Davidson 1987).

However, explaining these observations about Swampman does not seem to admit of the same resources as in the case of the observations about Donald. Swampman certainly *looks* similar to Donald and thus *appears* to be an instance of *homo sapiens* etc. But...is he? Or it? The best we can do is to say that, well, Swampman seems (very) *similar* to an instance of *homo sapiens* etc. However, this does significantly less work in explaining what Swampman is and what Swampman does. This point is reinforced by considering a case in which the lightning would have rearranged the tree molecules, not into a human-like Swampman but, say, into “Orfolo”, an exemplar of “Orfolei”, a wholly alien, carbon based intelligence (Bisson 1991).⁴ Presumably, in explaining the observations about Orfolo, such as its (or his? or hers?) mind-numbing ability to produce complex arrays of sounds, flashes of colours (is it communicating?) and movements (is it locomoting?), we would have to start all over with our classification of what we’re dealing with. We could, perhaps, say that Orfolo is in some circumscribed ways similar to this or similar to that.

⁴I assume that “Orfolo” is the singular of “Orfolei”.

After all, we're to a large degree carbon-based as well. Still, so are bicycle frames. In any case, identifying such similarities wouldn't by itself explain much about Orfalo—we would simply be describing our observations and associations with it.

Now, identifying intuitions as *similar to* feelings would bring us into a Swampman- or Orfalo-like explanatory situation. Identifying intuitions as feelings, on the other hand, is more akin to the Samuel- or Donald-scenario. It seems plain to me that the latter situation is the more favourable one. I suggest that this is where Intellectual Affectivism takes us.

A Roadmap to Come Back to

Here is the roadmap to Affectivism: In chapter 2 I will delineate the target state of which Affectivism is a theory from other things we call “intuition”. After making some qualifying remarks, I will first say what I do *not* mean by “intuition” and then I will paint a positive picture of what I *do* mean. What I mean by “intuitions” are intuition experiences. Intuition experiences are occurrent conscious mental states that are (partially) characterised by their characteristic (but not necessarily *sui generis*) phenomenology—in other words, there is something it is like to have an intuition in my sense. This phenomenology is part of the specific “feature profile” of intuition experiences.

In the remainder of chapter 2, I will use examples to zero in on this feature profile of intuitions. It will come to the fore that intuitions are not only phenomenally conscious but also intentional states, often taking a proposition or propositional content as their intentional object. This content they represent assertively, i.e. as true or false. And they motivate or push the subject to assent or dissent to what they represent as true or false. In doing so, however, they fall short of fully committing the subject to their contents, akin to perceptual experiences and unlike beliefs and judgments. Furthermore, intuition experiences are gradable in two ways: On the one hand, similar to the content of perceptual experiences, they can be more or less determinate in the way they represent their contents. On the other hand, they can push you to assent or dissent more or less strongly. A feature of intuition experiences is that they exhibit phenomenal epistemic valence, i.e. they can directly represent their contents either as true (positive intuitions) or as false (negative intuitions). Positive intuitions feel genuinely different from negative

intuitions even if their contents are the same. They constitute phenomenal polar opposites. Finally, intuition experiences are nonvoluntary, that is, intuitions are not under voluntary control but happen to one.

To sum up on the feature profile: intuition experiences are 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced and 7) nonvoluntary. How good a theory of intuition experiences is will depend on its ability to accommodate this feature profile, i.e. on its ability to *acknowledge and explain* these features

So, after outlining this feature profile, I will put it on the wish list for a good theory of intuitions and go shop through the aisles of extant philosophical intuition theories in chapter 3. What do we find there on the subject matter of intuitions? Besides a strong epistemological focus, we find Eliminativism lurking in a corner, Doxasticism — being somewhat of the majority view — occupying most of the shopping shelves and then, finally, we arrive at an assortment of fairly popular niche products: Perceptualism.

Eliminativism about intuitions states that whatever it is that we call “intuitions” and whatever we think they are, they — like Swampmans and Orofolei — do not exist. This will appear to be a position difficult to sell since at least some things that we call “intuitions” certainly do appear to exist. For instance, sometimes we call beliefs or judgments “intuitions” and most people — except, perhaps, for some *intuition-unspecific* Eliminativists — seem to agree that beliefs and judgments exist. Or aren’t you *judging* right now that I am on the wrong or the right track with Eliminativism?

In order to demystify intuitions, Doxasticists use exactly this datum to their advantage and claim that intuitions *are* doxastic states, i.e. beliefs, judgments or something in the ballpark. If this turns out to be a viable theory of intuition experiences, then that would be helpful in not only granting intuition experiences their features but — since we allegedly know things about doxastic states — also in explaining these features. In other words, we would find ourselves in a scenario akin to where we could explain observations about Samuel and Donald by classifying them as *homo sapiens*. In fact, Doxasticism turns out to be a fair description of *some* mental states that we call “intuitions”. However, it also turns out to be a poor characterization of *intuition experiences*. This is simply because doxastic states such as beliefs or judgments do not fit our wish list of features — particularly the *phenomenal* features — outlined in chapter 2. So instead of getting us into the coveted Samuel/Donald-scenario, Doxasticism gets us into a situation where we try to explain observations about Donald by classifying him as, say, a bat.

Calling Donald an instance of a bat will perhaps attract Thomas Nagel's attention but it won't help us make sense of many observations about him—while it might help us to make sense of *some*, such as him being warm- (or after being called a bat) hot-blooded.

Perceptualism takes the phenomenal impotence of Doxasticism as a point of departure and makes its *start* from the phenomenal features of intuitions. They proceed by describing and elaborating the features of intuitions, using perceptual experiences as an analogy. Arguably, this approach is more helpful than Doxasticism for acknowledging the phenomenal features of intuitions. The deficit of Perceptualism lies in its inability to *explain* these features. Perceptualism claims that intuitions emerge as in some ways *similar to* perceptions but they do not claim that intuitions *are* perceptions: intuitions are similar but not perceptions. So what *are* intuitions and what explains that intuitions are the way they are? Here Perceptualism does not provide a satisfactory answer. Ultimately, Perceptualism seems to bring us into a situation that can be located somewhere in between a Swampman- and an Orfalo-scenario. We might be able to state that Swampman and Orfalo have such and such features and try to make sense of them by associating them with knowledge about other things, say, Donald, *homo sapiens* or bats. Stating these *similarities*, however, appear of relatively little explanatory power to explain the observations made.

To sum up: extant intuition theories are either unable to acknowledge the features of intuitions (Doxasticism) or they cannot explain them (Perceptualism)—or both (Eliminativism).

Chapter 4 lays the groundwork for a fresh look at intuition experiences as we find it in Intellectual Affectivism. One goal of the chapter is to establish a first motivation to examine the hypothesis that intuitions are what Affectivism claims them to be: feelings. One part of the chapter is dedicated to clarifying what I mean by feelings: not qualia, not bodily sensations—but affective experiences. After that, I will dive into an exploration of what we know about affective experiences, specifically their phenomenal and intentional features. The central feature of affective experiences is their phenomenal valence, i.e. the felt positivity or negativity of certain experiences. Affective experiences are essentially *valenced* experiences. This is what sets them apart from mere qualia and bodily sensations. Another characteristic phenomenal aspect of affective experiences is felt arousal: during an affective experience, the subject feels a more or less localised increase or decrease (i.e. *change*) in the level of activation, energy, or excitement. Both valence and arousal are closely associated with another crucial feature of affective experiences: they

are directly motivating and so move us to behave or act in more or less stereotypical ways. So affective experiences are valenced, arousing and motivational. Furthermore, they are gradable along these dimensions: they can be more or less arousing as well as more or less positive or negative and motivating.

When it comes to intentionality, feelings have two parts: a *particular object* and a *formal object* where the feeling represents the former as bearing the latter. In this context, the formal object refers to feeling-specific properties such as the painfulness of pain, the fearsomeness of fear, the funniness of amusement, etc. (section 4.4.2). Representing something to be painful or amusing constitutes an evaluation which is *phenomenally grounded* in the *valence* of affective experiences (section 4.4.3). Concerning their particular objects, affective experiences can be in principle about all kinds of things, ranging from bodily events and parts over objects in the physical and imaginary realm (including fictions and propositions themselves), states of affairs in past, present and future over to propositional contents (section 4.4.1). To accommodate this broad range of intentional objects, feelings rely on a division of representational labour with other mental states such as bodily sensations, perceptions, judgments, memories, imaginings etc. This *base* of a feeling is a mental state (or a set of states) that supplies a feeling with its particular object: the feeling does not access the particular object directly but through other mental states.

Putting these explanatory resources on the table goes some way demonstrating that we know quite a bit about affective experiences. Thus, classifying something as affective experiences would amount to something akin to the Samuel/Donald-scenario. As a consequence, Affectivism about a mental state, if successful on other counts, emerges as a good thing. The last section of chapter 4 probes the idea that intuition experiences might be affective experiences by juxtaposing their respective features. The result of this comparison is favourable, establishing initial plausibility to explore Affectivism about intuitions further. The next chapters will be the steps to intellectualise Affectivism.

Chapter 5 takes the first step into this direction and introduces a class of mental states that appears well suited to house intuition experiences: epistemic feelings. Epistemic feelings have been broadly described as “feelings that enter into the epistemic processes of inquiry, knowledge and metacognition” (de Sousa 2008, p. 189). I will first give the reader a pre-theoretical grasp for some prominent instances of epistemic feelings such as the feeling of knowing (FOK) and the feeling of familiarity (FOF). Then we will dive into the research of what brings about these feelings. Apart from looking behind the

scenes of FOKs and FOFs, we will touch upon an influential *general* framework which understands (some) epistemic feelings as a form of metacognition.

Chapter 6 will make the case for epistemic feelings as a subclass of affective experiences. Based on empirical findings reviewed in section 6.2 and 6.3 as well as a mix of theoretical and phenomenological considerations interspersed throughout section 6.4 and 6.5.1 I will try to establish that epistemic feelings are affective experiences and analyse them along these lines (section 6.5). Furthermore, section 6.4 will consider how Affectivism about *epistemic feelings* can contribute to a better understanding of several aspects of epistemic feelings. Eventually, epistemic feelings emerge as experiences with a valence, arousal and motivational dimension. In other words: epistemic feelings are affective experiences. And that they are despite the fact that their inherent valence and arousal are mild and the behaviours they motivate often of a covert mental or cognitive nature. As all affective experiences, epistemic feelings have particular objects, formal objects and bases that provide their particular objects that are epistemically evaluated by epistemic feelings. The feeling-specific formal objects of epistemic feelings are epistemic properties broadly construed: familiarity, uncertainty, understanding, knowing, rightness etc.

Chapter 7 identifies intuition experiences as specific epistemic feelings: feelings of rightness and feelings of wrongness. To be more precise: intuitions are specific varieties of FOR/Ws, namely feelings of truth (FOTs) and feeling of falsity (FOFs, from now on). While identifying Donald to be a mammal is informative and roughly accurate, identifying him to be a primate is more informative and accurate—and identifying him to be an exemplar of *homo sapiens* is yet still more informative and accurate. So: intuitions are feelings. More precisely: epistemic feelings. Still more precisely: FOR/Ws. Yet still more precisely: FOT/Fs. In section 7.2 I will give you a first feel for FOR/Ws. Then, as I have done previously for other epistemic feelings, I will take you to a FOR/Ws science exhibition (section 7.3).

After this, I will provide an analysis of FOR/Ws based on the resources established for affective experiences and epistemic feelings (section 7.4). In the course of this analysis, I will delineate FOT/Fs among FOR/Ws as those FOR/Ws that take propositions as their particular objects and represent them as right or wrong, amounting to an evaluation of the propositions as true or false. This analysis will conclude in the realization that the feature profiles of FOT/Fs and positive and negative intuition experiences are identical. Thus, by inference to the best explanation, intuitions are identical to FOT/Fs and Intellectual Affectivism is successful. Affectivism emerges as a good theory of intuitions:

it can not only acknowledge the features of intuitions but also explain them. Intuitions have the features they have for essentially the same reasons as bodily and emotional feelings have them—because they are (specific) affective experiences.

One of the major advantages of Affectivism, as developed here, is that intuition experiences become ripe for a deeper mechanistic exploration. This is why the final main chapter 8 will consist in an attempt to elucidate the psychological mechanisms behind intuitions/FOT/Fs and their relation to actual truth and falsity. In this context, we will consider the role of processing fluency and, importantly, coherence with currently activated information in memory. We will see that the mechanisms that regularly produce FOT/Fs do so based on what appears to be an excellent psychological implementation of truth. Against this background, we will also be able to establish a connection between intuitions, on the one hand, and understanding and beliefs, on the other.

Chapter 2

The Marks of Intuition Experiences

2.1. Introduction: Acknowledging Intuitions

It is important to acknowledge that we use the word “intuition” (and its cognates) in philosophical, scientific and everyday contexts in manifold ways (cf. Andow 2015). The present thesis does not aim at providing a theory of how we use the word “intuition”. As Ole Koksvik notes:

[I]t is not an aim [...] to investigate the use of the word ‘intuition’ and its cognates in ordinary English. Uses of the word are highly varied, and of little value to the investigation of our target mental state. [...] The target of this inquiry is not the usage patterns of words, either in everyday situations or in philosophy. (Koksvik 2011, pp. 16 sq.)

I acknowledge that sometimes we mean this and sometimes we mean that when speaking of “intuition”; there is no one *single* item that we call “intuition”, but many.

Relatedly, the aim here is not to provide a theory of *all* the various things we call “intuition”. What I offer here is an account of a specific phenomenon among the many things we call “intuition”. You might have noticed that I previously used the plural “intuitions” rather than “intuition”. This is to delineate the topic of this thesis — as something that we refer to when using “intuition” as a *countable* noun — from things that we call “intuition” using the word as an *uncountable* noun. As an uncountable noun “intuition” usually refers to traits of people (Jung 1971; Myers 1962) or mental

capacities that we possess (e.g. Glöckner and Witteman 2010). Psychologists often understand “intuition” in this sense: as referring to a capacity or system — often going by names such as the intuitive “System 1” that contrasts with the deliberate “System 2” (Kahneman 2011) — that engenders “intuitive”, i.e. largely automatic and unconscious, processes (e.g. Epstein 2008; Evans 2008; Evans and Stanovich 2013; De Neys 2017). This system and the processes it brings about, in turn, are often taken to ground the trait of intuition in people (Betsch 2008a). This capacity- or trait-reading of “intuition” will not be at the forefront of my interest here.

Instead, I will be primarily concerned with “intuition” in its countable sense—and more specifically: when understood as mental states: intuitions. Note that by doing so I grant that *sometimes* when speaking of “intuition” or “intuitions”, we do not mean mental states at all (cf. Cappelen 2012). Perhaps when we say dismissively about a claim that somebody makes “Yeah...that’s just an intuition...” we don’t refer to a mental state but to a possibly non-mental content (which might sometimes be the content of a mental state) (see e.g. Bealer 1992, footnote 7; Molyneux 2014, p. 457). I grant that such cases might well exist. However, in this thesis such cases are not the ones I am mostly interested in. I am interested in “intuitions” understood as mental states.

At the same time, the present thesis will not be concerned with *all* the various things we call “intuition” in its countable sense—and not with all the mental states that we call intuitions. It will be concerned with mental states that we call “intuitions” and that exhibit certain features to be specified in this chapter—this will be my target state here. So I think the only datum required by my theory is that *sometimes* we use the word “intuition” to refer to mental states and to the target state in particular.⁵

Ideally, the theory to be elaborated here should be a *good* theory of the target state. What makes a good theory of the target state? For starters: capturing and explaining the features of the target state. I will present a selection of these features. Mostly, I haven’t come up with these features myself. One finds these properties scattered in the intuition literature. The resulting list, call it the “feature profile of intuitions”, will not be exhaustive. It will be, however, fairly extensive. This is to grant that by and large my fellow intuition theorists are collectively onto something right with their descriptions of intuitions. Attention and resources are limited, however, and individual authors attend to some features and not to others.

⁵Or perhaps not even that is required (cf. Bengson 2014).

This neglect is sometimes coupled with a denial that intuitions have the features other authors have identified and used in their diagnosis of what kind of state intuitions are. Such denials seem to be often motivated by metaphysical (or phenomenal) parsimony which seems to mandate being austere in the features one ascribes to intuitions so as to avoid commitment to the existence of a mysterious state with high complexity.

This principle of parsimony, be it metaphysical or phenomenal (or both), is widely endorsed in the whole gamut of philosophy of mind. Levy, for instance, voices that “we ought to avoid multiplying mental states unnecessarily, we need to ensure that we are postulating exotic states and processes only when they are truly needed” (Levy 2016, p. 9; see also Egan 2011, p. 67; Quilty-Dunn 2015, p. 277). An analogous parsimony principle concerning *phenomenal* properties is outlined by de Vignemont: “one should not posit additional phenomenal properties in one’s mental ontology when one can account for them by appealing to other properties” (de Vignemont 2019, p. 268). Specifically, with intuitions in mind, Lynch reminds us of “a good commandment to live by, philosophically speaking. Namely, *thou shalt not posit mysterious faculties without necessity*” (Lynch 2006, p. 231).

It is, thus, unsurprising that Paul Boghossian’s repeated worry targets the obscurity of intuitions resulting from their somewhat mysterious ontological status:

To be sure, the idea that we possess a quasi-perceptual faculty—going by the name of ‘rational intuition’—the exercise of which gives us direct insight into the necessary properties of the world, has been historically influential. It would be fair to say, however, that no one has succeeded in saying what this faculty really is nor how it manages to yield the relevant knowledge. ‘Intuition’ seems like a name for the mystery we are addressing, rather than a solution to it. (Boghossian 2000, p. 230)

The single most influential consideration against rational insight theories can be stated quite simply: no one has been able to explain—clearly enough—in what an act of rational insight could intelligibly consist. That is, no one has been able to say how some cognitive act, of a sort that we might plausibly enjoy, is able to yield immediate knowledge of the modal properties of properties

If the theory of rational insight is to serve as a genuine explanation for how we are able to have such a priori knowledge, rather than simply acting as a

placeholder for such an explanation, it must consist in more than a suggestive label; it must somehow lay bare, in appropriate detail, how some capacity that we have gets to work on the properties we are able to think about so as to disclose their natures. (Boghossian and Bonjour 2001, p. 635)

One of the upshots of the present thesis will be that the complexity of the allegedly mysterious state is real but that the mystery around it is not. Restraint on behalf of metaphysical parsimony is thus unwarranted. We can embrace the existence of the state with all its complexity—with all its features.

But since the list will not be exhaustive: Why *these* features? First, as already mentioned, this is what others say. I am, by and large, persuaded by what they say and find the mentioned features in my own phenomenology (and hope that you will too). Second, I think that together they constitute a cluster of properties that will capture *only* instances of the target state, although probably not all of them.⁶

I think this is a good start. And starting from there will also help us see how one can accommodate the fact that some instances of the target state might lack some of the features on the list (or have additional ones).⁷ The third point is methodological: the bulk of the features are *phenomenal* properties. Now it seems that phenomenal properties constitute a comfortable point of departure for philosophical inquiry for the simple reason that we are able to access and to say something informative about such features without leaving the armchair. Phenomenology, in contrast to subjects better reserved for experimental work such as descriptive claims about unconscious causes (i.e. aetiology), appears as a proper province for a philosophy of mind that is concerned with describing the psychological nature of a specific mental state. This is not to say that such analytic armchair phenomenology is not without limitations but to say that it is a viable approach to *phenomenally* circumscribed targets. Furthermore, it does not imply that we have to stick to phenomenology and refrain from theoretical and empirical investigation beyond appearances. The idea is rather that appearances are a promising point of departure to engage in such further inquiry.

⁶I think that there are features among them that are found in all and only in instances of the target state. That is, features that are had by all and only by instances of the target state. These features are *most* diagnostic for intuitions while other features intuitions might share with other states. To anticipate: I think these features are assertiveness and epistemic phenomenal valence that I will introduce in 2.2.2 and 2.2.6 respectively (you find reasons for this assessment in chapter 7, especially section 7.4). Nevertheless, sticking only to these features would conceal much of the exciting complexity of intuitions.

⁷Note that some of these features might not be independent and might (together with other features) imply or entail other features on the list and not on the list.

So let me put some flesh on the bones of these abstract considerations and without further ado give you a quick outlook on what is to come in this chapter.

2.2. Intuition Experiences and Their Features

Let me start by presenting some paradigmatic propositions which intuition theorists take to prompt intuitions. Taking this as a point of departure I will then proceed to characterise what is prompted. Consider propositions such as:

- “If p, then not-not-p.” (Non-Contradiction)
- “Causes necessitate their effects.” (Causality)
- “Physical objects continue to exist when we do not perceive them.” (Permanence)
- “Torturing kittens for fun is wrong.” (Torture)
- “Non-conscious physical duplicates of conscious beings are possible.” (P-Zombies)
- “A cylinder with a certain base and height encloses a greater volume than a cone with the same base and height.”
- “Gettier’s Smith does not know that he will get the job.” (Gettier)⁸
- “There are more than fifty stars in the universe.” (Fifty Stars)

It is assumed — and so will I assume here — that reading these propositions — at least sometimes, perhaps after some iterations and reflection — prompts distinctive conscious episodes: intuitions. Having an intuition has a phenomenal quality: there is something *it is like* to have an intuition.

That means that alongside the myriad of other phenomenal states you are undergoing now (such as your momentary perceptual experiences, emotions and mood), an intuition makes a contribution to your overall phenomenal state—to how it is like to be you in this very moment.

⁸Assuming one is sufficiently familiar with Gettier cases (Gettier 1963).

Each of us has a multitude of experiences each day: gustatory, olfactory and visual experiences; moods, emotions, and bodily sensations. Each has a phenomenal character; there is something it is like to have it. There is also a global phenomenal character, something it is like overall to be a particular conscious being at a particular time. To say that there is an intuitional experience, that intuition has a phenomenal character, is to say that a person's having an intuition contributes to the phenomenal character of her overall experience (Koksvik 2011, p. 103).

This overall phenomenal state would have been different without the intuition. Perhaps when reading one of these propositions, you, at first, did not have an intuition concerning it while essentially having the same other phenomenal states. Or — instead of having an intuition — you tried to assess the standing of the proposition by (perhaps quite automatically) recalling someone's testimony concerning it or recalling the fact that you had an intuition concerning it in the past. This supposedly felt different from having an intuition. Consider for instance:

- “The negation of a disjunction is the conjunction of the negations.” (De Morgan's Law)

If you have never considered *De Morgan's Law* it is likely that you won't instantly have the intuition that it is true. Though, after some reflection, something might happen: suddenly, it just *seems* true to you (Bealer 1992, p. 101). This is a modification in your overall phenomenal state brought about by the occurrence of an intuition. Something similar is the case for the phenomenal contrast between considering $2 + 3 = 5$ and $6253 + 4773 = 11026$ ⁹ or the above-mentioned propositions and propositions such as:

- “A cylinder with a certain base and height encloses three times the volume as a cone with the same base and height.”
- “1729 is the smallest number expressible as the sum of two positive cubes in two different ways.”
- “Dogs are better pets than cats.”

⁹These arithmetic operations are not strictly speaking propositions. If you take issue with this, consider them as a shorthand for propositions such as “Two plus three equals five”.

- “There are no differences but differences of degree between different degrees of difference and no difference.”¹⁰

This is not to say that one cannot have intuitions concerning these propositions. Yet, without idiosyncratic preconditions such as exceptional giftedness, expertise, enculturation, emotional dispositions, or intoxication, these propositions are just less likely than e.g. *Non-Contradiction* or *Permanence* to instantly elicit intuitions.

Against this background, we begin to realize that there is something characteristic happening in our consciousness upon the occurrence of an intuition. This is the sense in which intuitions are occurrent, conscious, phenomenal states (cf. Talbot 2013, p. 318; Pust 2014; Koksvik 2017). What its (phenomenal) nature consists in specifically I discuss shortly. As is common in the philosophical literature, I will use “seem” or “appear” to generically refer to this specific phenomenology without committing to views that flesh out these locutions when characterizing intuitions (for that see 3.4).

2.2.1. Intentionality

During the conscious episodes that constitute intuitions, it seems or appears to one that *Non-Contradiction* or that *Torture*. Intuitions are thus not like oranges, trees, kangaroos, fires or smokes that – assuming they’re not part of a symbolist artwork – stand for nothing apart from themselves. By “just being there” oranges, trees, kangaroos, fires or smokes don’t “say” anything.¹¹ This is different for intuitions. They are more like pictures, street-signs, lists of items, loudspeaker announcements, or fortune cookie slips that — besides having their physical dimension — also have an intrinsic *informational* dimension: They “say” something and what they say is about or is directed at things that are not the intuition itself, i.e. are external to the intuition (Crane 2009). If you have the intuition that *Permanence*, then the intuition is about the *Permanence* proposition, not the intuition itself.

A shorthand for saying that, is to say that intuitions are *intentional* states; that the things they are about or directed upon, say, propositions, are the intentional objects

¹⁰This is William James telling us something he felt to be profoundly true—while on laughing gas (James 1882, p. 207, see also section 7.2).

¹¹This, of course, does not mean that these things do not carry information for a suitably equipped interpreter (Dretske 1981, 1986).

of the intuitions; and that the intuitions can be said to be about or to represent¹² their intentional objects, say, the specific propositions (cf. Koksvik 2011, p. 3, 2017, p. 2). This is not to exclude that intuitions can be about non-propositional objects (see e.g. Bonjour 2001, pp. 676 sqq.; Chudnoff 2011b, 2013a; Molyneux 2014, pp. 450 sqq.).¹³ However, if not indicated otherwise I will focus on propositions as the intentional objects of intuitions, i.e. on *propositional* intuitions for now (cf. Sosa 2007a, p. 52). I will refer to the intentional object of an intuition in the form of a proposition as the content of the intuition.

Saying that a state is intentional leaves a free parameter as to what kind(s) of intentionality, “aboutness” or “directedness” the state exhibits and what kind(s) it does not. Although here I will remain neutral on this issue, let me make some qualifying remarks: There are various (and variously demanding) notions of intentionality on the market, for instance: tracking or co-variational theories (e.g. Dretske 1981; Millikan 1984; Millikan 1993); inferential or conceptual role semantics (e.g. Block 1986; Harman 1987); and phenomenal intentionality theories theories (e.g. Horgan and Tienson 2002; Kriegel 2013). Intuitions might be intentional in one or more of these ways and not intentional in one or more others. Thus far, all that is needed is that intuitions are intentional in at least *one* way.

Different kinds of intentionality are not mutually exclusive.¹⁴ For instance, part of the information or content that a state carries can be “broadcasted” and become phenomenally conscious. In doing so, *phenomenal intentionality* might become instantiated in relation to part of a state’s content. This might, in turn, engender certain (e.g. global and person-level) functions for which e.g. a transition over phenomenal to access con-

¹²For now I use “to represent X” synonymously with “to be about X” or “to be directed onto X”. I thus take being representational as synonymous with being intentional and not, as some do, as being intentional *and* assertive (see section 2.2.2). This is similar to Bengson’s distinction between “merely contentful” (here: intentional) and representational states (here: intentional and assertive) (Bengson 2015, footnote 11).

¹³*A fortiori* there is no commitment to the controversial and more encompassing thesis of Propositionalism according to which *all* intentional states are relations to propositions or something proposition-like (for a critical review see Montague 2007). See McGahhey and Van Leeuwen 2018 for a refreshing critical discussion of Propositionalism specifically in relation to intuitions.

¹⁴On the other hand, different *views* or *theories* of intentionality are often pitted against each other in a mutually exclusive fashion. That is, each view describes a property that it argues to be what intentionality *really* is. In effect, a view of intentionality does not usually deny the existence of the property that another view of intentionality describes. What it denies instead is that this property *is* (a form of) intentionality. So different theories of intentionality usually come together with the assumption of monism about intentionality. For my purposes it is for now acceptable to be pluralist (or, in fact, agnostic) about intentionality, taking the distinct properties that different views describe as varieties of *intentionality*. This is how my talk of kinds of intentionality is to be understood.

consciousness is necessary (Bayne and Chalmers 2003). At the same time, plausibly, not all intentional features need to be “phenomenally encoded”, i.e. be part of *phenomenal* intentionality. Consequently, other parts of the information or content that the state carries are not “broadcasted” but are nevertheless fed to under-the-phenomenal-radar (e.g. more specialised or local) processes, assuming thereby *part* of their functional (e.g. inferential) role that does not require broadcasting or consciousness.¹⁵ Consequently, the information or content that a state carries can be more than is apparent via its phenomenal – and phenomenally intentional – character. That is to say that the state in question has both, phenomenal intentionality *and* functional role semantics. At the same time, it seems that both these mechanisms enjoy a degree of independence and might also come apart as the ventral and dorsal visual pathways, blindsight or hemispatial neglect demonstrate.¹⁶

Within the intuition literature, intuitions are generally assumed to be intentional. However, it is more often than not left open what kind(s) of intentionality applies to intuitions. Nevertheless, those that concern themselves explicitly with intuition experiences can usually be positioned at least in relation to phenomenal intentionality, i.e. the thesis that the intentional content of a state is fixed by the phenomenology of the state.¹⁷ Chudnoff seems to attribute phenomenal intentionality to intuitions (Chudnoff 2013b). Bengson is less clear on that matter but, with his emphasis on the phenomenal directness of the content of intuitions, seems plausibly understood to lean towards phenomenal intentionality (Bengson 2015). Koksvik explicitly attributes phenomenal intentionality to perceptual experiences but not to intuition experiences. On his view, phenomenal intentionality can only be realised via content-specific phenomenal properties which are, according to him, lacking in intuitions (Koksvik 2011, pp. 200 sq., 248).¹⁸ To see this

¹⁵To give it a gloss of a metaphorical hierarchy: part of the informational value of a state might travel “upwards” and be broadcasted while part might remain “on the same level” or travel downwards in the cognitive hierarchy.

¹⁶The ventral and the dorsal visual pathways both process and represent visual information but do not seem to equally give rise to *phenomenally intentional* content (cf. Clark 2009). In blindsight, patients report to be not conscious of any visual objects (the “blind” part) but when forced to interact with objects in ways that typically depend on vision, perform significantly better than one would expect against the background of their report (the “sight” part) (Stoerig and Cowey 2007). This suggests that some information processing is going on without phenomenology. Similar lessons can be drawn from research on hemispatial neglect where patients loose awareness of *part* of the visual field (Rafal et al. 2006).

¹⁷The phenomenal intentionality thesis is easily confused with its mirror image: the intentionalist or representationalist thesis about phenomenal consciousness. Intentionalism is roughly the thesis that the phenomenology of a state is fixed by the intentionality of the state (e.g. Tye 1995; Chalmers 2004; Crane 2009).

¹⁸That is to say that according to Koksvik intuitions do not have content-specific phenomenal properties when stripped down to their *essence*. This leaves open that intuitions can sometimes, i.e. incidentally,

we need to acquaint ourselves with Koksvik's very useful distinction between content-specific and attitude-specific phenomenology. In his own words:

What sort of contribution does intuitional experience make? We can distinguish between two different ways such contributions can be determined. On the one hand, the contribution a mental state makes can depend on the *content* of that mental state. Suppose that if I perceive something red, this makes a different contribution to the character of my overall phenomenal experience than perceiving something green does. In that case, perception has *content-specific phenomenology* (and if not, it does not). On the other hand, it may be that *whatever* I visually perceive, the fact that I am having *a* visual perceptual experience makes a contribution to the character of my overall phenomenal experience: perhaps it contributes a certain 'visualness' [...] In that case I shall say that perception has *attitude-specific phenomenology*. Of course, perception may have, and indeed actually does have, both attitude-specific and content-specific phenomenology. (Koksvik 2011, p. 104)¹⁹

So according to this distinction, a state can have content- and attitude-specific phenomenology, while Koksvik seems to grant *phenomenal* intentionality only to content-specific phenomenology (see also footnotes 24, 69 and section 4.4).

have content-specific phenomenal properties. For instance, when they happen to have a content with phenomenal properties (e.g. iconic content). In order to exclude this, it must be additionally essential to intuitions that they can *only* have content without phenomenal properties.

¹⁹Koksvik goes on to helpfully clarify the distinction:

The terms 'content-specific phenomenology' and 'attitude-specific phenomenology' are not perfect. For one, 'attitude-specific phenomenology' might make it sound like the phenomenology suffices to distinguish one attitude from another. As I will be using the terms, this is a substantive question, and in fact I shall argue that intuition and perception *share* aspects of their attitude-specific phenomenology. The terms may also suggest that the ultimate *origin* of the phenomenology is in the content, or in the attitude, respectively. However, the issues here are subtle, and the terms are intended to leave questions of origin open. It is possible, for example, that a certain attitude only admits content of a particular kind. Some think that perception is like this; it admits only non-conceptual content. Suppose that it does, and further that nonconceptual content always makes a different contribution to the character of a person's overall phenomenal experience than conceptual content does. Though suppose that, contrary to fact, no-matter what the content is, the contribution is always the same. In that case, the way the terms are used here, perception would have attitude-specific phenomenology but not content-specific phenomenology, even though the origin of the phenomenology is in the content. Thus the terms are intended to indicate *variation with*, rather than *ultimate origin in*, content and attitude (although, again, two different attitudes can share aspects of their attitude-specific phenomenology). (Koksvik 2011, p. 104)

In effect, depending on what kind of intentionality one ascribes to intuitions, intentionality itself might or might not count as a phenomenal feature. In the case of intuitions, phenomenology and intentionality might or might not come apart.

2.2.2. Assertiveness

Apart from being intentional in one way or another, i.e. having a certain *kind of intentionality* (be it phenomenal or not), a feature that is often explicitly assumed to be phenomenal is the specific *mode of intentionality* that intuitions have.²⁰ That is, an intuition is not only about or directed onto its content, this general intentionality being of a specific kind, but an intuition is about or directed onto its content *in a specific way*. In the case of intuitions propositions are not represented as funny or doubtful, neither are they suggested for consideration or inquisition. Rather, the intuition “asserts” or “claims” the proposition. In other words, an intuition represents its content *as true*.²¹ This is how intuitions are *assertive* or represent assertively.²² According to intuition theorists, intuitions appear to have this assertiveness in common with beliefs and perceptual experiences. And it stands in contrast to e.g. imaginings which represent their content without making a claim about its truth or desires that insist on the fulfilment rather than the truth of their content.

Let me add some caveats. There is a notorious problem with spelling out what assertiveness amounts to beyond conversational metaphors such as assertive states “*purport[s] to be saying how the world is*” (Pryor 2013, p. 209). Even if we assume that we can spell out what assertiveness means in the case of, say, beliefs, it does not seem to follow that the same notion of assertiveness straightforwardly applies to other states usually taken to be assertive such as perceptual experiences (Briesen 2015). That is, it might well

²⁰The distinction between kind and mode of intentionality amounts to the distinction between a general account of intentionality and a specific account of intentionality of intuitions under a general account of intentionality. Two distinct kinds of states, say beliefs and desires, can have the same kind of intentionality, say, one specified by inferential role, but differ in the way or the mode in which they realise this intentionality, say, the content of beliefs play other inferential roles than that of desires or – in accordance with their inferential roles – beliefs represent their content as actual while desires represent it as to be actualized. In other words, once you pick a kind of intentionality, there are sometimes different modes of this kind of intentionality and sometimes the kind of intentionality might be exhausted by a single mode.

²¹Or false, see section 2.2.6 further below.

²²The term “assertive” is ambiguous in that it can mean either “claiming” or “convincing”. In philosophy it is usually used to denote the former while in everyday life it is rather used to denote the latter. Here it is meant in the former way, i.e. by an intuition being assertive is meant that the specific way in which it represents is “claiming”, not that it is *per se* “convincing” (see also sections 2.2.3 and 2.2.4 further below).

be that the assertiveness of some states, say, perceptual experiences is distinct from the assertiveness of other states, say, beliefs. Alternatively, it might be that the assertiveness is on some level of abstraction the same while the sources or grounds of this assertiveness might differ.²³ This echoes the distinction between kinds and modes of intentionality. One kind of intentionality might have multiple distinct assertive modes. Or distinct kinds of intentionality might have (multiple) assertive modes. Note that assertiveness does not have to be a feature of the intentional content (as endorsed by Chudnoff and perhaps by Bengson) but can be a feature of the intentional mode or attitude of intuition. Understanding assertiveness as a phenomenal feature of intuitions does thus not entail it being a *phenomenal intentional* feature of intuitions. At least not in the sense of phenomenal intentionality which has so far only been developed as a thesis about intentional content, not about intentional mode or attitude.²⁴

This is why Koksvik infers from his premise that intuitions lack content-specific phenomenology the conclusion that intuitions cannot have phenomenal intentionality. At any rate, for now, it is enough to note that intuitions certainly appear assertive whether it is the same kind of (grounds of) assertiveness as for beliefs and/or perceptual experiences or not.

²³Assume for a moment that we cash out assertiveness in terms of a mind-to-world direction of fit (which is itself a vague metaphorical notion) (Briesen 2015, pp. 2245 sqq.). Then in order to be assertive a state must have a mind-to-world direction of fit. It is, I think, uncontroversial that beliefs as well as perceptual experiences satisfy this criterion, and are thus both assertive in this way. However, the way in *how* they cater to fit the world seems to be quite distinct. Phenomenally, they appear to have markedly different qualities. Chudnoff (in a similar vein but more clearly than Bengson), for instance, notes that in addition to representing a proposition as true (like belief), an intuition additionally (like perceptual experience) presents us with truth-makers for these propositions by making it “as if objects and their features are directly before the mind” (Chudnoff 2011b, p. 636).

²⁴That is, the phenomenal intentionality thesis has so far largely been silent on the intentional properties of intentional modes or attitudes, if such properties can be coherently conceptualized (for one notable exception see Horgan and Tienson 2002, p. 522). On the face of it, it does not seem to be straightforward that intentional contents are the only bearers of intentional features. Intentionalism or representationalism, a theory of phenomenal consciousness via intentionality and a mirror image of phenomenal intentionality which is a theory of intentionality via phenomenal consciousness, sometimes draws on the phenomenal features of attitudes in addition to those of contents to explain the phenomenal qualities of mental states (e.g. Crane 2009). I might be confused, but it seems to me that something similar might be possible for intentional features, and here I am not wholly alone (see Weiss 2016, footnote 45)

2.2.3. Pushiness

Another aspect of intuitions which is sometimes confounded with assertiveness is their *motivational* or *pushy character*.²⁵ It is not only that they describe or “say” something (i.e. represent) and do it in a certain way (i.e. assertively). At the same time, they also direct one to behave or act in a specific way towards their content, i.e. towards what they assertively represent or say: they motivate, push or incline you to assent²⁶ to their content (Van Inwagen 1997; Sosa 2007c; Koksvik 2011; Chudnoff 2014b). Consequently, I will use pushiness or being pushy to denote the character of being motivational with respect to a specific mental behaviour or action: assent (or dissent) towards a content. That is, being pushy is a special case of being motivational.²⁷

A metaphor might help: a politician may say many things. She may say many things assertively. Now, if one happens to be one of her supporters then one will probably be motivated to assent to what she assertively states. On the other hand, if one does not happen to support her one will probably lack the motivation to assent to what she says assertively. Perhaps one will be motivationally neutral concerning assent to what is assertively stated or one will be motivated to dissent instead. Nevertheless, the fact that she will still be saying things and assertively so is unaltered. What changes is one’s motivation to assent (or dissent). In this sense, assertiveness and pushiness can come apart (cf. Siegel 2014, pp. 55 sqq.). Some kind of assertiveness is present in beliefs, perceptual experiences and intuition experiences alike. They all “say” things assertively. However, intuition experiences additionally single-handedly and directly motivate one’s assent (or dissent) to what they assertively state. In contrast, beliefs do not seem to be pushy but to consist of past assent and perceptual experiences do not seem to be pushy without additional desires.²⁸

²⁵I borrow the term “pushy” from Koksvik 2011.

²⁶Or dissent, see section 2.2.6 further below.

²⁷In the literature, there is an ambiguity on whether pushiness is to be understood as a *felt* push or as just somehow establishing a *disposition* to do something without implying pushy phenomenology. I take this to be the difference between talk of intuitions as inclinations (e.g. Sosa 2007c) and dispositions (e.g. Sosa 1996). The former is more demanding and includes the latter. The account to be developed here will focus on pushiness understood phenomenally. Insofar it can also be taken as an elaboration of the nature of intuitions as inclinations.

²⁸This leaves open, however, the possibility that beliefs and perceptual experiences *non-phenomenally* dispose the subject to assent to their content (see footnote 27).

2.2.4. Being Noncommittal and Related Features

In contrast to a belief or judgment, however, an intuition does not fully *commit* the subject to its content, insofar it is *noncommittal*.²⁹ Unlike an intuition experience “belief, when it is conscious, [...] is like a reminder of what I have already agreed is the case. Conscious belief has phenomenology of already present *commitment*, not of pushiness” (Koksvik 2011, p. 185). Instead, intuitions are “the tendencies that make certain beliefs attractive to us, that “move” us in the direction of accepting certain propositions without taking us all the way to acceptance” (Van Inwagen 1997, p. 309). To illustrate, consider:

- “For all conditions, there is a set containing all and only the things meeting this condition.” (Naïve Comp)

If, after some reflection, *Naïve Comp* seems true to you then you are in good company. Many philosophers report that it seems true to them even though they have been convinced by Russell’s paradox that *Naïve Comp* is false:³⁰

I have an intuition—it still *seems* to me—that the naïve comprehension axiom of set theory is true; this is so despite the fact that I do not believe that it is true (because I know of the set-theoretical paradoxes). (Bealer 1992, p. 102; see also Williamson 2007, p. 217)

Similar observations can be made for paradoxical clusters of propositions that all appear true but cannot all be true without leading to paradoxes. As in the case of *Naïve Comp* resolving some paradox by rejecting some of its assumptions often does not dispose of one’s intuitions about the dismissed assumption.

Intuitive seemings remain distinctive conscious states in their own right, without collapsing into beliefs, as is shown by paradoxes like the liar, or the sorites. Each proposition in a paradoxical cluster exerts a powerful intuitive attraction, despite how compelling it also is that they cannot all be true together. Even when one eventually settles on a solution, moreover, the

²⁹This is only to say that having an intuition does not by itself amount to a commitment to its content. For that, we must *additionally* assent to it by judging accordingly. And that we *usually* do. In fact, that might be our *default mode* towards the deliverances of intuitions (cf. Döring 2009, pp. 293 sq.). In some cases, however, we might resist (see below).

³⁰The paradox is due to the set of all sets that do not contain themselves. If it does not contain itself then it must contain itself and contradict its condition.

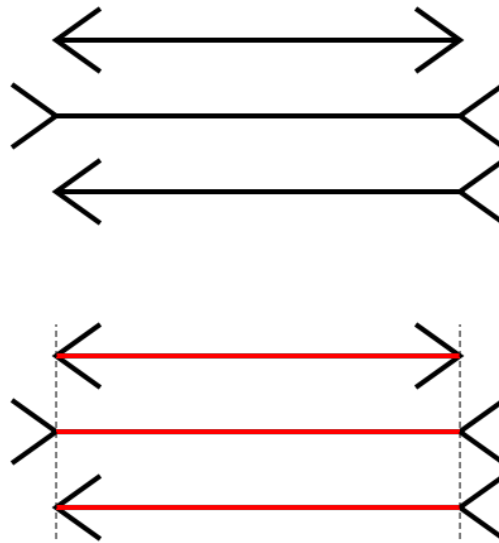
pull of the rejected proposition is not removed but overcome. (Sosa 2007a, p. 51; see also e.g. Sosa 2007c, pp. 47 sqq.)

Here is another proposition that most likely will seem true to you (it certainly does to me, much more so than *Naïve Comp*):

- “There are more natural numbers than there are even numbers.” (Numbers)

I get the intuition that *Numbers* is true. Surely, even numbers are merely numbers such as 0, 2, 4 and so on while natural numbers include all even numbers as well as all odd numbers. Though in fact, this intuition exploits our somewhat shaky grasp on infinities: the set of natural numbers, as well as the set of even numbers, contain (countably) infinite numbers. As a consequence, all elements of the former set can be exhaustively matched up with all elements of the latter set. Therefore, *Numbers* is false. However, I still get the intuition that it is true. Perhaps you do, too.

Some philosophers make sense of these kinds of “cognitive illusions” in analogy with optical illusions such as the Müller-Lyer illusion:



Depiction of the Müller-Lyer illusion

I can't and I suppose you can't either help seeing the lines as having different lengths even though they are of the same length and we both know them to be of the same length (Bealer 1996, p. 6). However, this analogy has its limits and should not be taken

to suggest that one cannot rid oneself of one's intuitions after learning about them being mistaken. As Nimtz points out:

[O]ptical illusions are peculiar in that changes in our beliefs do not affect our perceptual experiences. Everyone will continue to have the impression of the lines differing in length in the Müller-Lyer illusion, whatever her knowledge of their actual length. But this does not hold true in the cases of intuitions [...]. Russell's proof [...] might well make you *lose* the intuitions you used to have. (Nimtz 2010, p. 366)

Just as there is a gradient in the ease with which one can have (or obtain) intuitions about certain propositions, i.e. some intuitions are easier had than others, there seems to be a gradient in the rigidity of different intuitions, i.e. some intuitions are easier lost than others. Presumably, however, many of them can be lost (with sufficient efforts such as plentiful exposure to and recapitulation of counter-evidence). This is a dissimilarity in comparison to optical illusions that seem to prevail no matter what.³¹³² Thus, although intuitions seem to be less plastic than e.g. beliefs they are more plastic than perceptual experience.

Furthermore, there appear to be differences in plasticity within the class of intuitions. Consider for instance:

- “The sentence ‘the boy the man the girl saw chased fled’ is ungrammatical.”
(Grammaticality)

As *Naïve Comp* and *Numbers, Grammaticality* will likely seem true to you. However, the sentence *is* grammatically (although not stylistically) sound.³³ English grammar

³¹At least if one has them in the first place (Segall et al. 1966; Henrich et al. 2010).

³²Perhaps the “no matter what”-formulation is too strong. However, to my knowledge there are no reported cases where subjects (without neurophysiological damage) could be made to lose the illusory Müller-Lyer experience.

³³The example is taken from Koksvik 2011, pp. 2 sq., 135. Koksvik claims that one has the intuition that *the sentence ‘the boy the man the girl saw chased fled’ is ungrammatical* is true (the text in italics is the content of the intuition). However, it is not clear whether such a heavy-handed description (reminding of the intuitions with nested contents) is the most plausible (or the only) one. To me, it seems more natural to describe what occurs as an intuition that represents the sentence “*the boy the man the girl saw chased fled*” straight as incorrect or ungrammatical. Here, an assessment of correctness or grammaticality, and not of truth, is part of the attitude, rather than part of the content of the intuition. Furthermore, I — at least initially — rather seem to have a negative intuition concerning the correctness or grammaticality of the sentence rather than a positive intuition regarding the truth of the whole *Grammaticality* proposition (although the proposition seems to be an OK re-description in terms of truth). It might be that one can have both intuitions in quick succession. However, that opens up the possibility that intuitions are more versatile regarding the properties they are receptive to. It seems plausible that we do not only have intuitions about the truth of propositions but also

allows for such *centre embedding* where a phrase is embedded in the middle of another phrase (e.g. De Roeck et al. 1982).

Now, read it again. The sentence might still seem ungrammatical to you. Let me provide you with what the sentence means: “The boy fled; the boy, that is, who was chased by the man, who, in turn, the girl saw” (Koksvik 2011, p. 3). Now read it a couple of times more. Presumably, with the information provided and accumulating iterations the intuition might gradually loosen its grip on you. After a while, the proposition might cease to appear true to you (though it does not have to appear false to you either).³⁴ However, it took a while until your intuition could be persuaded even in the face of strong evidence (testimony, explanation) against it. In other words, you probably ceased to *believe* that *Grammaticality* is true (perhaps after I told you) but it took a good deal more time and effort to cease to have the *Grammaticality* intuition.

What this shows is that, first, intuitions are quite clearly fallible³⁵ and, second, they are to some (varying) degree *belief-independent*.³⁶ The fact that one believes that *Naïve Comp* or *Grammaticality* are false does not (straight away) adjust one’s illusory positive intuitions towards the propositions. This reminds of the (more pronounced) belief-independence in cases of visual illusions.

2.2.5. Gradeability: Content and Pushiness

Another feature of intuitions is their *gradeability*. In fact, there are two aspects in which intuitions are said to be gradable. Call the first *content-gradeability*. Content-

about their normative standing, be it relative to moral, logical or linguistic norms. Considering e.g. *Torture* against this background might also lead to a reconstruction of the involved intuition as representing *torturing kittens for fun* as morally wrong rather than as representing *torturing kittens for fun is wrong* as true. This also introduces the potential to confuse the different properties intuitions might (come to) be receptive to: an intuition that tracks one’s — perhaps idiosyncratic — moral norms or subjective tastes might be mistaken to track truth. For the sake of simplicity, I will bracket this issue for now.

³⁴Alternatively, the sentence might cease to appear ungrammatical to you.

³⁵This can be taken to be implied by the fact that intuitions are intentional and assertive. If they represent something as true or false, they can naturally misrepresent it in so being.

³⁶What might explain this on a functional level is that the mechanisms generating intuitions are “informationally encapsulated”, that is, they are to some degree (initially) unaffected by (relevant) information that is present elsewhere in one’s cognitive system but currently lies outside the system generating the intuition. In contrast, beliefs and judgements can seem to be “informationally parasitic” in that they initially rely for their formation on sources of information external to the belief-system.

gradeability concerns the way an intuition experience represents its content.³⁷ It builds on a perceptual analogy: some perceptual experiences are more *determinate* or *vivid*, i.e. clear and distinct in representing their contents than others. Glasses, for instance, make the content delivered by visual experience more or less determinate.³⁸ Now, some philosophers claim that there is a parallel quality to the content of intuitions, i.e. things that we see with the “mind’s eye” or hear with the “mind’s ear” can be more or less clear and distinct (Chudnoff 2011b; Bengson 2015). In some cases the content of an intuition might be inexplicit: The subject might not be able to fully articulate the content of her intuition (Chudnoff 2011b, pp. 642 sq.; Bengson 2015, pp. 730 sq.). On the other hand, there is *pushiness-gradeability* (Weinberg 2007, p. 335; Koksvik 2011). Naturally, we can be more or less strongly pushed towards assent (or dissent).

One may wonder whether and — if so — how content-gradeability and pushiness-gradeability co-vary. On the face of it, they can be (largely) dependent or independent. We can experience a strong push to assent *that p* because *p* is so clearly represented: I feel a strong push towards assenting that there is a bottle on the table before me *because* I see it clearly and distinctly. On the other hand, something might be clearly and distinctly represented without me being strongly, or even at, all pushed towards assenting to it. Koksvik points out that “it might *clearly but weakly* seem to someone that torturing one innocent person to save 20 others is permissible (while it clearly and strongly seems to that person that torturing an innocent person to save a thousand is). Put differently, “there may be no doubt about what the content of the state is—no “haziness,” [...]—but the intuition may still be weak.” (Koksvik 2017, p. 6)

2.2.6. Phenomenal Epistemic Valence

A peculiar feature of intuitions that sets them apart from states such as perceptual experiences, beliefs or imaginings is their *phenomenal epistemic valence* (Koksvik 2011). So far we have considered *positive* intuitions, i.e. intuitions that represent some content as true. Now suppose it is 1963 and you believe that knowledge is justified true belief. You’re reading Gettier 1963 for the first time. Following Gettier’s exposition, you finally

³⁷Note content-gradeability is to be understood as a feature of the experience (Bengson 2015, footnote 15, 26; cf. Harman 1990) and, further, as a feature of its content-specific phenomenology (Koksvik 2017, p. 8).

³⁸One way of accounting for this is to say that a visual experience with and without glasses represent the same content, but with a varying degree of determinacy or vividness. Another way would be to say that they actually don’t. What is in fact going on is that the visual experience with glasses simply has *another* content than the visual experience without glasses (Bourget 2017).

arrive at Smith’s belief that *the man who will get the job has ten coins in his pocket*. As it turns out, this belief is justified and true. Thus: Smith *knows* that the man who will get the job has ten coins in his pocket (Inverse Gettier). Yet *Inverse Gettier* just seems false to you. Something similar might happen if you consider propositions such as:

- “ $2 + 3 = 4$ ” (Calc Error)
- “People prefer pain over pleasure.” (Pain)
- “An undermined house will rise up.” (Rise)
- “The earth is a perfect square.” (Square Earth)
- “Water flows up by itself.” (Flow-Up)
- “There are as many numbers between 0 and 1 as there are between 0 and 1000000.” (Numbers 2)

In such cases, what one experiences is a *negative* intuition where a proposition seems false instead of true.³⁹ And this it does by *directly* representing a *specific* proposition as false. This stands in contrast to representing *Smith does not know that the man who will get the job has ten coins in his pocket* or $2 + 3 = 5$ as true in response to *Inverse Gettier* or *Calc Error* or to having positive intuitions that represent a *nested* content such as *it is false that (Inverse Gettier)* as true.⁴⁰ Crucially, positive intuitions differ *phenomenally* or in *how they feel* from negative intuitions (Koksvik 2011). This phenomenal contrast is only there when you compare a positive to a negative intuition and not there if you compare two positive intuitions or two negative intuitions. Interestingly, one can have a positive *or* a negative intuition that, say, *non-conscious physical duplicates of conscious beings are possible*.

One way to capture such contrasts is to say that the truth or falsity, i.e. the epistemic valence, of the intuition is not expressed in the content but in the *attitude* and that their valence is part of their attitude-specific *phenomenology* (ibid., pp. 188 sq.). Consequently, positive intuitions differ phenomenally from negative intuitions not due to a

³⁹Being positive or negative is different from being correct or illusory. For instance, the *Naïve Comp* intuition is positive because it represents *Naïve Comp* as true but illusory because *Naïve Comp* is false. Similarly, the intuition that represents *Numbers 2* as false is negative and incorrect because *Numbers 2* is actually true.

⁴⁰That is not to say that if one has an intuition in response to propositions such as *Pain* or *Flow-Up* one *always* has directly corresponding negative intuitions. It is possible that one can just as well have positive contradicting intuitions (with nested contents).

difference in content — this can be, in fact, identical — but due to a difference in attitude. In other words, positive intuitions are not the same kind of attitude as negative intuitions in the same way that happiness is not the same kind of attitude as sadness.

Positive and negative intuitions seem to constitute a pair of what Kevin Mulligan has helpfully called ‘polar opposites’ (Mulligan 2007, pp. 217 sqq.). I like to add that they are not only polar opposites but *phenomenal* polar opposites.

Note that variation along the feature dimension of phenomenal epistemic valence covaries with characteristic differences along other (phenomenal) feature dimensions. A difference in valence leads to a valence-specific difference in assertiveness and pushiness: Positive intuitions assert a proposition to be true while negative intuitions assert a proposition to be false; positive intuitions push one to assent to its content while negative intuitions push one to dissent to its content. This opens up the possibility that the features of assertiveness and pushiness might be not independent of the feature of phenomenal epistemic valence. It might, for instance, turn out that epistemic phenomenal valence is (part of) the phenomenal grounds of the assertiveness (or pushiness) of intuitions (see section 2.2.2 and 4.5).

Phenomenal epistemic valence is a peculiar feature of intuitions that sets them apart from perceptual experiences, beliefs or imaginings. Perceptual experience, in general, does not directly represent some (propositional) content as being false. Although one can *infer* that some proposition is false from the content of one’s perceptual experience. At most, it only directly represents some (propositional) content as being true. The contrast with belief is different: belief seems to be able to directly represent that some proposition is false. Nevertheless, the most plausible way it would do so is on the level of content via the above-mentioned nested contents. Although we can believe that p (is true) and we can believe that p is false (or *not* p), the belief that p and the belief that p is false seem 1) only to differ in terms of their content and 2) lack a phenomenal contrast. A (conscious occurrent) belief (or judgment) that *P-Zombies is true* will not phenomenally contrast with a belief that *P-Zombies is false*. Imaginings, on the other hand, represent contents but they do not by themselves orient their contents in relation to truth or falsity. They have therefore no epistemic valence, let alone *phenomenal* epistemic valence.

2.2.7. Nonvoluntariness

Another feature of intuitions is that they are *nonvoluntary*. The formation of an intuition is not an intentional act and is not subject to direct voluntary control. One does not choose or decide to form or have an intuition. While conscious choices, decisions, judgments, guesses or imaginings are sometimes under a certain degree of conscious control, intuitions are passively received—they happen or fail to happen to one (Bengson 2015, see also footnote 67).⁴¹ Consequently, one is not responsible or rationally criticisable for one's intuitions in the same way as one is for one's choices, decisions, judgments, guesses or imaginings (Koksvik 2011; Chudnoff 2014b).

I don't mean to say here that e.g. judgments or beliefs are under full voluntary control. It certainly doesn't seem that we are free to believe or judge whatever we want (Williams 1970). Nevertheless, there seems to be more leeway for agency in the case of beliefs and judgments than in the case of intuitions and perceptual experiences (perhaps there is still more leeway in the case of the former than in the latter). That is why we are often held responsible or are rationally criticisable for our beliefs and judgments and less so for our intuitions and perceptual experiences (Koksvik 2011; Chudnoff 2014b). In the former case, we seem to be norm-bound epistemic agents while in the latter we are at least to a much greater extent epistemic patients.

Of course, being to some degree acquainted with the dynamics of such nonvoluntary states one may be able to have an indirect effect on their occurrence by performing conducive actions over which we *do* have control. Just as one does not choose to form the visual experience that there is an orange lying on the table, one does not choose to have (or have not) the *Gettier*, *P-Zombie* or *Naïve Comp* intuition. One may, however, choose to look elsewhere or to attend to specific features of a case. As discussed above, one is not free to simply and directly rid oneself of the intuitions one believes to be mistaken. In the case of intuitions (but not perceptual experience), however, it is possible to influence one's dispositions to experience certain intuitions in the long-run (cf. Holroyd 2012).

2.3. Recap and a Phenomenally Contrastive Mixed List

In this chapter I hoped to illustrate that intuitions are not only phenomenally conscious but also intentional states, often taking a proposition or propositional content as their

⁴¹ Although one might have actively worked towards getting or losing them.

intentional object. This content they represent assertively, i.e. as true or false. And they motivate or push the subject to assent or dissent to what they represent as true or false. In doing so, however, they fall short of fully committing the subject to their contents, akin to perceptual experiences and unlike beliefs and judgments. Furthermore, intuition experiences are gradable in two ways: on the one hand, similar to the content of perceptual experiences they can be more or less determinate in the way they represent their contents. On the other hand, they can push you to assent or dissent more or less strongly. A peculiar feature of intuition experiences is that they exhibit phenomenal epistemic valence, i.e. they can directly represent their contents either as true (positive intuitions) or as false (negative intuitions). Positive intuitions feel genuinely different from negative intuitions even if their contents are the same. They constitute phenomenal polar opposites. Finally, intuition experiences are nonvoluntary, that is, intuitions are not under voluntary control but happen to one. So intuition experiences are 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced and 7) nonvoluntary.

In concluding this chapter, I want to address a caveat: it is possible that the summarised phenomenal features were not always salient to you in your own phenomenology. This is in part due to the relative homogeneity of the grouped propositions. The various clusters of propositions were supposed to jointly illustrate specific features characterizing instances of (correct or incorrect) positive, (correct or incorrect) negative or no intuitions. Sometimes, however, such a homogeneous grouping fails to draw a phenomenal contrast between instances where we have intuitions and instances where we don't (cf. Dechêne et al. 2009). So here is a more heterogeneous list of propositions that might let the mentioned phenomenal features come to the fore in a starker contrast:

1. "1729 + 4773 = 6052"
2. "2 + 3 = 5"
3. "3 + 2 = 4"
4. "If P then not not P."
5. "If P then not not not not not not P."
6. "1729 is the smallest number expressible as the sum of two positive cubes in two different ways."

7. “for all conditions, there is a set containing all and only the things meeting this condition.”
8. “There are as many numbers between 0 and 1 as there are between 0 and 1000000.”
9. “There are no differences but differences of degree between different degrees of difference and no difference.”
10. “Physical objects continue to exist when we do not perceive them.”

I hope that upon reading this mixed list of propositions, some of the described features have figured more prominently into your total phenomenal state. In case you wonder, here is how somebody whose phenomenology is perfectly attuned to the illustrative purposes of this exercise (*you*, I hope) would have experienced the propositions: 1) 1, 5, 6, 9 should have prompted no intuition experiences (although you might have engaged in inferential judgments upon e.g. counting the “nots” in 6 etc.). 2) 2, 3, 4, 7, 8, 10 should have prompted intuition experiences. 3) 2, 4, 7, 10 should have prompted positive intuition experiences whereas in the case of 7 the positive intuition experience is incorrect. 4) 3 and 8 should have prompted negative intuition experiences whereas in the case of 8 the negative intuition experience is incorrect.

You should have experienced various phenomenal contrasts between instances of 1), 3) and 4). 1) should have contrasted with 3) and 4) in that the phenomenal features characteristic of intuition experiences are absent in 1) and present in 3) and 4). Furthermore, although propositions in 3) and 4) prompt intuition experiences with most characteristic phenomenal features being roughly the same, they still contrast quite starkly. This is because the phenomenal epistemic valence is positive in 3) and negative in 4). In fact, this difference in valence leads also to a valence-specific difference in assertiveness (asserting a proposition to be true vs. asserting a proposition to be false) and pushiness (push to assent vs. push to dissent). Finally, correct and incorrect intuitions prompted by 3) and 4) should not significantly differ when it comes to their phenomenology (although there might, of course, be differences in degrees of content-determinacy and pushiness).

Chapter 3

Intuitions and Where (Not) to Find Them: Extant Theories of Intuitions

3.1. Introduction: Intuitions for Philosophy

Now that we have a firmer grasp on the target state, namely intuitions, it is time to go into the theories that have been proposed about it. In the present chapter, I will review extant theories of intuition. These theories, however, should be savoured with some caution. The reason why is that they tend to be developed with a specific goal in mind: intuition theorists engage with the general question “What are intuitions?” in order to provide an answer to the more specific question “Are intuitions justifiers?”. As Elijah Chudnoff notes:

Some philosophers think that intuitions are a source of justification [...] for beliefs [...]. Other philosophers are more skeptical. They doubt that intuitions are a source of justification for beliefs [...]. The motivation to defend one or the other of these orientations, non-skeptical or skeptical, drives most discussions about intuition. As a consequence most of these discussions focus on issues that seem to bear immediately on the epistemological status of intuitions [...] Prior to all these questions, however, is the question: What are intuitions? Depending on what intuitions are, they might or might not be reliable, they might or might not possibly justify beliefs [...], they might

or might not be embarrassed by recent experimental studies, and they might or might not be coherently foresworn. (Chudnoff 2011b, p. 625)

This narrow epistemological take is a consequence of the largely metaphilosophically motivated intuition debate within which intuition theorizing is embedded. This is because many philosophers accept Herman Cappelen’s “Centrality thesis”:

Centrality (of Intuitions in Contemporary Philosophy): Contemporary analytic philosophers rely on intuitions as evidence (or as a source of evidence) for philosophical theories. (Cappelen 2012, p. 3)⁴²

The implications of this thesis, in turn, are well captured by Michael DePaul and William Ramsey:

In contemporary analytic discussions [...] “intuition” has become the name for whatever it is that might provide philosophy with a distinctive method and hence preserve it as a separate (in principle) intellectual domain. Our disagreement about the nature and epistemic authority of intuitions is at root a battle for the preservation of philosophy as an autonomous field of inquiry (DePaul and Ramsey 1998, p. 7).

I won’t belabour the fact that contemporary intuition theorizing seems to be overly metaphilosophically motivated and thus concerned with the epistemology of intuitions but simply note it for now. One of the consequences, anyway, is the occasional lack of clarity on whether claims about e.g. the similarity between intuition and other mental states are to be understood epistemologically or metaphysically (or both). On the face of it, it seems it is one thing to say that some state is like another state when it comes to its epistemology and quite another to say the same when it comes to its metaphysics. Baring, of course, the possibility that the epistemology of a state exhausts its nature, i.e. metaphysics. In fact, one does find such approaches:

I propose to identify intuitions [...] by their putative justificatory role in philosophical practice: intuitions are responses that are putatively either foundational [...] or quasi-foundational justifiers for philosophically relevant non-empirical belief (Kauppinen 2013, p. 361).

⁴²There are some who deny this descriptive claim about philosophy (Earlenbaugh and Molyneux 2009; Cappelen 2012; Molyneux 2014; Deutsch 2015). Nado calls members of this comparatively small faction in the debate “intuition deniers” (Nado 2016). Some have explicitly attempted to refute intuition denial (Bengson 2014; Andow 2017; Nado 2017) while others have directly argued that philosophers *do* use intuitions as evidence (Climenhaga 2018).

Be it as it may, after noting that the theories to come could be understood as making claims about the epistemology of intuitions *or* about the metaphysics of intuition (or both), I will assume that they are to be understood as proposals about the nature of intuitions and not (or not only) as proposals about the epistemology of intuitions. After all, what I am after here is a psychological kind, not an epistemological kind (cf. Koksvik 2011).

Note further that *individual* claims that a theory makes about intuitions can be true without making thereby the whole theory true. Conversely, a theory as a whole can be false without negating the truth of some of its individual component claims. Why this obvious point? Because, on the one hand, I think that the claims that different theories make regarding the individual features of intuitions are by and large on the right track. This is what Chapter I was about. On the other hand, I think that the overarching theoretical frameworks in which these claims are embedded are mistaken.

Here's a bird's eye view on the frameworks proposed so far. A few say intuitions probably do not exist (e.g. Smith 2000, pp. 23 sq.). Call this *Eliminativism* about intuitions. Others claim that intuitions *are* judgments or beliefs (see e.g. Lewis 1983, p. x; Plantinga 1993; Van Inwagen 1997, p. 309; Williamson 2004, 2007; Lynch 2006; Ludwig 2007). Call this *Simple Doxasticism*. Yet others claim that intuitions *are* dispositions (e.g. Sosa 1996, 1998; Earlenbaugh and Molyneux 2009) or inclinations (or attractions or temptations) to judge or believe (e.g. Van Inwagen 1997, p. 309; Sosa 2006, 2007c,a; Boghossian 2009). Call this *Dispositional Doxasticism* and *Inclinal Doxasticism* respectively. Taken together, these views constitute the family of intuition theories called *Doxasticism*.⁴³

Then there are those opposed to Doxasticism. They maintain that intuitions are pre- and non-doxastic. This faction typically uses perceptual experiences as an analogy to analyse intuitions. Such *Perceptualism* is not to be understood literally: No one is claiming that intuitions *are* perceptual experiences. Instead, they either claim that intuitions and perceptual experiences belong to the same kind of superordinate state such as “seemings” or “presentations” (e.g. Bealer 1992; Bengson 2015) or that there are instructive similarities between intuitions and perceptual experiences without committing to further ontological linkage (e.g. Koksvik 2011; Chudnoff 2013c). Call them *Quasi-Perceptualism* or *Comparative Perceptualism* respectively. In what follows I'll zoom in on the various views and their issues.

⁴³Chudnoff was, as far as I can tell, the first one to talk of “Doxasticism” and “Perceptualism” in the context of intuitions (Chudnoff 2011b).

3.2. No Intuitions?

There are those who deny the existence of mental states that are appropriately called “intuitions”. Here is a rather clear statement of such “Eliminativism” about intuitions:

[What] exactly is an intuition? One rarely encounters clear statements of their nature. If an intuition is a thought, why employ a term suggesting it is anything less than that? If intuition is a particular type of thought, what type? If an intuition is an emotion or feeling, what distinguishes intuition from ill-founded feelings? [...] Are intuitions desires? Hunches? Stubborn convictions that a person refuses to surrender? The point is, we cannot be sure whether we have such things, let alone what role they play in providing moral guidance, until we know precisely what intuitions are. One suspects that the absence of definition, keeping intuition afloat as a hazy “something” between a thought and a feeling, may hide the fact that there are no such things. (Smith 2000, pp. 23 sq.)⁴⁴

Taking such a position is often motivated by the air of mystery surrounding the topic of intuition (see section 2.1). The most straightforward way to understand Eliminativism is as the thesis that the term “intuition” has no extension. It might have an intension but there are no things that fit it and to which the term could extend. Intuitions are thus like Swampmans, Orfolei, Flying Spaghetti Monsters or Phlogiston: they do not exist.

Eliminativism: The term “intuition” has no extension—intuitions do not exist.

Another motivation to adopt Eliminativism is the highly varied use of the word “intuition” and its cognates, sometimes not referring to mental states at all, but e.g. to commonsensical ideas or linguistic devices used to hedge epistemic risk.⁴⁵ Embracing Eliminativism, however, seems to be a premature response to such polysemy. “Intuition” sometimes means this and sometimes means that, depending on the context etc. That is, intuitions are not like Swampmans or Flying Spaghetti Monsters but more like bats.

⁴⁴For other philosophers who seem to embrace something resembling Eliminativism see Ayer 1956, p. 31; Fumerton 1990, p. 6; Cappelen 2012.

⁴⁵Looking outside of philosophy one finds psychologists noting something similar about “intuition”: “There are as many meanings for the term intuition as there are people using it.” (Betsch 2008b, p. 3) In fact, intuition “has been given so many different meanings, some opposite to others, that it makes one wonder whether the term has any meaning at all.” (Epstein 2008, p. 23)

The polysemy of the word “bat” does not lead us to deny the existence of a winged mammal that navigates space through sonar; nor the existence of an elongated tool used to hit objects. It seems perfectly fine to continue to engage in debates about how exactly bats navigate space, what it is like to be a bat or whether wooden bats are preferable to metal ones. Just be clear what you mean by “bat”.

Perhaps intuitions are not (only) like bats but they’re also like chefs. Colloquially, “chef” is often synonymously used with “cook”. So perhaps “intuitions” are synonymously to some other, less confusing term. In fact, perhaps everything that can be said about intuitions can be said in less ambiguous terms, maybe using a few more words than just one. If (some of) the things that we call intuitions exist and if *all* these things can be called by less ambiguous names, then perhaps the most reasonable position to adopt is not Eliminativism but Reductionism about intuitions, i.e. the thesis that intuitions are best understood as other familiar mental states. Perhaps one could then adopt Eliminativism regarding the term “intuition”, petitioning to purge it from our language practice on behalf of clarity. This is what Smith in part alludes to in the quote above. In any case, Eliminativism as *a theory of intuition experiences* appears clearly insufficient:

Insufficiency of Eliminativism (IE):

- (IE-P1) Eliminativism is true if it can accommodate (i.e. acknowledge and explain) the feature profile of intuition experiences (“intuitions”, from now on).
- (IE-P2) Eliminativism cannot (and does not want to) accommodate the feature profile of intuitions.
- (IE-C1) Therefore, Eliminativism is false.

Coming back to the things we call intuitions: another possibility is that (some of) the things that we call intuitions exist and that *some but not all* of these things can be called by less ambiguous names. That is: there are states among the various things to which the term “intuition” extends for which there is no better pre-theoretical term than “intuition”. Perhaps there is something characteristic and substantial that we call “intuition” and that has not yet been sufficiently theoretically elucidated to count as well-understood. This is what I, together with non-reductive Perceptualists, think is the case.⁴⁶ I’ll introduce positions from both sides in turn, starting with one that is often stylised as a form of Reductionism about intuitions: Doxasticism.

⁴⁶However, I do not think that this has to lead to Non-Reductionism about intuitions.

3.3. Intuitions and Doxastic States

“Doxasticists” (Chudnoff 2011b) are considered “reductionists” (Koksvik forthcoming) or “minimalists” (Bengson 2015) since they take intuitions to be reducible to a relatively “familiar” class of mental states such as (some subset of) beliefs or judgments (Simple Doxasticism) or (some subset of) dispositions or inclinations to believe or judge (Dispositional Doxasticism, Inclinal Doxasticism).⁴⁷ The subsets of the states in question are often taken to be individuated by a specific kind of content (e.g. modal, abstract, commonsensical or counterfactual) and/or aetiology (e.g. positively such as based on sheer understanding or conceptual competence or negatively such as *not* based on memory, inference, introspection etc.).⁴⁸ Here are Ludwig, a Simple Doxasticist, constraining the aetiology of intuition to understanding or conceptual competence and Sosa, a Doxastic Inclinationist, constraining the content and aetiology of intuitions to modal content and to our capacity for understanding respectively (for a discussion of these constraints see next section):

I will use “intuition” to mean an occurrent judgment formed solely on the basis of competence in the concepts involved in response to a question about a scenario, or simply an occurrent judgment formed solely on the basis of competence in the concepts involved in it (in response, we might say, to the null scenario). (Ludwig 2007, p. 135)

S rationally intuits that p if and only if S’s attraction to assent to <p> is explained by a competence (an epistemic ability or virtue) on the part of S to discriminate the true from the false reliably (enough) in some subfield of modally strong propositional contents that S understands well enough, with no reliance on introspection, perception, memory, testimony, or inference (no further reliance, anyhow, than any required for so much as understanding the given propositional content). (Sosa 2007a, p. 58)

⁴⁷I borrow the overarching terms “doxastic attitudes” for beliefs and judgments and “doxastic tendencies” for doxastic dispositions and inclinations from Bengson 2015. I use the term “doxastic states” to refer to and capture both.

⁴⁸We will see that introducing such constraints is not unique to the doxastic camp but can also be found among Perceptualists (Bealer 1999; Brogaard 2013, p. 279).

3.3.1. Intuitions as Doxastic Attitudes

Assume, with Simple Doxasticism, that intuitions are doxastic states such as beliefs or judgments. Unfortunately, what “doxastic” exactly means is not clear. The way the distinction between doxastic and non-doxastic accounts is drawn, often suggests that being “doxastic” just *is* being a belief or judgment, i.e. it denotes belief- or judgment-hood.⁴⁹ Thus:

Simple Doxasticism: Intuitions are beliefs or judgments.

At first glance, it seems plausible that in certain contexts we do in fact call some kinds of judgments or beliefs “intuitions”. Van Inwagen, for example, states that philosophers “call their philosophical beliefs intuitions because ‘intuition’ sounds more authoritative than ‘belief’” (Van Inwagen 1997, p. 309). It thus appears uncontroversial that *sometimes* when philosophers speak of “intuitions” they mean some kind of judgment or belief.⁵⁰

What kinds? Let me clarify a bit more: beliefs can be dispositional or they can be occurrent; judgments, while essentially occurrent, are sometimes unconscious and sometimes conscious. In fact, perhaps the only way beliefs can be conscious and occurrent is in the form of conscious judgments (Crane 2013). A state or event that is conscious can be phenomenally and/or access conscious (Block 1995; Bayne and Chalmers 2003). It might also be reflexively conscious. By that I mean that a subject is conscious or aware of being in a certain mental state, sometimes but not always additionally to this mental state being phenomenally and access conscious as well. Consider Crane’s illustration for the case of worry:

[C]onsciously worrying is not the same as being conscious *that* one is worrying. One could become conscious that one is worried about one’s finances by discovering something about one’s behaviour, say; and one could do this without undergoing the kinds of inner events which constitute conscious worrying. (Crane 2001, p. 106)

⁴⁹In other philosophical contexts “doxastic” is sometimes used to mean that the content of a mental state has or can be described as having a propositional structure. This, however, cannot be meant here since both, doxastic and non-doxastic theorists of intuitions consider intuitions to be states that can be described as having propositional contents.

⁵⁰And, as Koksvik points out, *sometimes* when we talk of conscious belief we actually mean intuition experiences (Koksvik 2011, p. 192).

Consciously worrying is here to be understood as a phenomenally conscious state. On the other hand, whatever state one is in when one becomes conscious *that* one worries, say thinking about one's behaviour, one is at this point reflexively conscious of one's worry while not necessarily being consciously worrying, i.e. one's worry being phenomenally conscious.⁵¹

I take it that if something has a phenomenology it is also conscious and if something is conscious it is also occurrent. On the other hand, not everything that is occurrent is also conscious and not everything that is conscious must be phenomenally conscious or have a phenomenology; it might, for instance, be access but not phenomenally conscious (cf. Strawson 2009, pp. 158 sqq.). Now it still seems fine to grant that by calling a certain mental state an "intuition" we mean *sometimes* to refer to beliefs of the dispositional, unconscious occurrent, and conscious occurrent variety; and to unconscious and conscious judgments. Furthermore, there seems to be no need to deny that some of the conscious beliefs and judgments will have a certain phenomenology, i.e. be phenomenally conscious. If these are the "intuitions" that Simple Doxasticists try to account for, there is no problem.

The point at which things get moot is when Simple Doxasticism is considered as an account for *intuition experiences* as characterised in chapter 2. When I say "I have an intuition" or when I simply have an intuition and in so doing I have an intuition experience, the state that I am in doesn't seem in any straightforward way identifiable as some sort of a less ambiguously specifiable state such as a belief or judgment. Of course, this leaves open that, at the end of the day, one might in a *non-straightforward way* succeed in identifying intuition experiences as, say, some belief or judgment. As Nitz notes:

Phenomenal data only goes so far. Mental states do not come with labels revealing what kind of states they are to the glance of the mind's eye. Dissecting our mental life into separate phenomena and sorting those into kinds is a theoretical task that is decided by the descriptive and explanatory powers of the competing accounts. (Nitz 2010, p. 365)

However, the phenomenal data *is* data that needs explaining by an account. And this is where Simple Doxasticism struggles.

⁵¹This is something to bear in mind when trying to understand what is meant by occurrent conscious belief. As Crane argues it cannot mean consciously believing, strictly speaking (Crane 2001, pp. 106 sqq.).

Let me work up to that slowly. For a start, our target state is occurrent and phenomenally conscious. So only phenomenally conscious occurrent beliefs and phenomenally conscious judgments can be considered as candidates for identifying our target state. The gap that is difficult to bridge lies, however, in the way phenomenology figures in making the state the kind of state it is. As should be apparent now, a belief that is not occurrent and not phenomenally conscious is still a belief; a judgment that is not phenomenally conscious is still a judgment. In contrast, if a state is not occurrent and not phenomenally conscious then it is not an intuition experience. In other words, phenomenology appears to be only incidental to beliefs and judgments. As Klausen points out: “Beliefs are not essentially or typically phenomenal states” (Klausen 2013, p. 188; see also Lyons 2018, p. 187). At the same time, phenomenology is essential to intuition experiences: It is (in part) what individuates the psychological kind with which I am concerned here. The resulting argument against Simple Doxasticism can be simply summarised as follows (cf. Mitchell 2018, p. 2):

- (P1) Intuition experiences essentially have a phenomenology.
- (P2) The state to which intuition experiences are reduced do not.
- (C) Therefore, the reduction fails and Simple Doxasticism is false.

Now, I have granted that beliefs and/or judgments *can* be phenomenally conscious. It seems then that Simple Doxasticists would be well advised to appeal to this datum in order to make sense of intuition experiences. However, actually, they don’t. Quite the opposite: they deny that there is a phenomenology in need of explanation. This seems consequential: If all intuitions are essentially beliefs or judgments and phenomenology is not essential to beliefs or judgments then phenomenology cannot be essential to intuitions. Assume for a moment that a Doxasticist would nevertheless advance a claim that is out of character:

I insist, intuitions are essentially beliefs or judgments. And yes, intuitions also do comprise intuition experiences. Yet, intuitions, as judgments and beliefs, are sometimes phenomenally conscious and sometimes not. In case intuitions are not phenomenally conscious then they are not intuition experiences but (phenomenally) unconscious intuitions. But when intuitions are phenomenally conscious then they are intuition experiences. Intuition

experiences are essentially beliefs or judgments that are phenomenally conscious. In other words, intuition experiences are beliefs or judgments that are essentially phenomenally conscious.

Such a move seems *ad hoc* and borders on circularity. In trying to account for intuition experiences, it was stipulated that there are beliefs or judgments that are essentially phenomenally conscious. This way, intuition experiences have been “explained” but now one has to explain that there are beliefs or judgments that are essentially phenomenally conscious. This datum seems quite curious since beliefs and/or judgments are not considered to be essentially phenomenally conscious. It appears that the explanatory baton has just been passed on without really explaining anything along the way. Furthermore, one may justifiably be tempted to ask: Which judgments or beliefs among all judgments or beliefs are those that are essentially phenomenally conscious? Those that constitute intuition experiences, presumably.

But then one may plausibly object that at least many instances of beliefs or judgments that are phenomenally conscious just don’t seem to have a phenomenology resembling intuition experiences. Perhaps that’s because there are also *incidentally* phenomenally conscious judgments or beliefs? But then what distinguishes these incidentally phenomenally conscious judgments or beliefs from essentially phenomenally conscious judgments or beliefs? The phenomenology specific to intuition experiences, presumably. I think by now the circular nature of such an argument becomes obvious.

This illustrates another point: identifying intuitions straight with beliefs or judgments does not only struggle with the datum that intuition experiences have a phenomenology at all but also does little to explain the *specific* phenomenology of intuition experiences. So even if one stipulates that intuition experiences are beliefs or judgments of the (essentially) phenomenally conscious type in order to account for the sole possession of phenomenology, one would have to make further qualifications on the *specific* phenomenology. Such qualifications appear necessary since we do not learn much if anything about the phenomenology of a mental state in virtue of learning that it is a belief or judgment.⁵² After all, a distinctive phenomenology is ordinarily not thought part of beliefs and judgments.⁵³ Thus such a phenomenology is something that either has to be

⁵²Perhaps I should add that we do not learn much *positive* about the phenomenology. We might learn something negative about it, i.e. about what the phenomenology is *not* like.

⁵³By saying that I do not want to take a stance on the cognitive phenomenology debate which states that besides the phenomenology of perceptual, somatic and affective states there is also a (distinctive) phenomenology to purely cognitive states such as beliefs and judgments (see e.g. Bayne and Montague 2011). However, Doxasticists usually do not allude to cognitive phenomenology in order

denied (see above) or will have to be explained not in terms of the kind of state intuitions are, namely judgments and beliefs, but in terms of something that is specific to the subset of beliefs and judgments that constitute intuition experiences (cf. Glüer 2009). That is, one has to explain some features of (some) intuitions *on their own* terms and *in contrast* to the kind of states intuitions are. Thus, the assimilation of all intuitions to beliefs and judgments is uninformative when it comes to explaining important features of at least some intuitions, i.e. intuition experiences. Such further phenomenological qualifications, however, would make the position all the more ad hoc.

For the sake of the argument, let's grant to the Simple Doxasticist that intuition experiences are essentially beliefs or judgments that are characterised by a specific phenomenology. A Simple Doxasticist would now be quite free in painting a picture of the specific intuitive phenomenology. However, in virtue of having introduced his Simple Doxastic bit, one constraint needs to be satisfied: the phenomenology he paints as essential to intuition experiences cannot be in outright contradiction to it being a belief or judgment. Nevertheless, on the face of it, some of the phenomenal features in chapter 2 are. That is to say: Specific phenomenal features of intuition experiences appear incompatible with the idea that they are beliefs or judgments.

For instance, beliefs and judgments are surely committal (rather than noncommittal). In fact, they seem to be essentially characterised by making or being reminded of a commitment. "What is distinctive of [...] belief (the attitude) is a certain kind of commitment to the truth" of its proposition (Crane 2001, p. 103). As a consequence, if intuitions were beliefs or judgments, we would have contradictory beliefs or judgments for propositions such as *Naive Comp* or *Grammaticality* where we have an intuition (that is occurrent and conscious) that p and where we at the same time (occurrently

to explain the phenomenology of intuitions, i.e. they remain silent on the issue of cognitive phenomenology. In fact, Doxasticists would probably deny that there is any phenomenology in need of explanation, cognitive or otherwise. It is thus a yet unexplored question whether the alleged cognitive phenomenology of judgments and beliefs can account for the phenomenology of intuitions. One can grant that intuitions have a phenomenology and even that beliefs and judgments have a cognitive phenomenology. However, it is not clear whether the phenomenology of intuitions is the same as the cognitive phenomenology of beliefs or judgments, whether intuitions have a unique cognitive phenomenology of their own or whether the phenomenology of intuitions is a *cognitive* phenomenology at all. On the face of it, the aspects of cognitive phenomenology discussed in the cognitive phenomenology literature are rather general in kind and seem not straightforwardly related to the kind of concrete phenomenological features that intuitions are claimed to possess, regardless of whether intuitions possess the discussed aspects of cognitive phenomenology as well. For the moment I want to note that the topic of cognitive phenomenology itself is controversial and for now best bracketed out. When discussing Perceptualism in section 3.4, we will see that Perceptualists might be thought to turn perceptual phenomenology into a cognitive phenomenology.

and consciously) believe or judge that *not p*. An agent will then have two contradicting beliefs or judgments at the same time, and would appear rationally criticisable or even irrational.⁵⁴ This would not be a very good result because in such cases we would not want to call the ones that have intuitions that contradict their beliefs irrational. It seems to be a feature of intuitions that they are noncommittal and at the same time relatively belief-independent. One appears *not* rationally criticisable when one can't help having positive intuitions in response to propositions which one knows, believes and judges to be false (Koksvik 2011).

Similarly, beliefs and judgments do not seem to be (more or less) pushy, at least not without additional desires.⁵⁵ That is, they do not motivate assent; rather they consist *in* assenting or having assented to a specific content. *If* beliefs or judgments motivate anything at all, then it is some kind of assenting behaviour towards propositions that are somehow (logically) related to the proposition that is their propositional content. However, their content is not in need of assent since it is either about to receive it or

⁵⁴See Glüer 2009 for a strategy to resolve similar conflicts that recommend against a doxastic theory of perceptual experience. Taking as a point of departure the observation that the contents of perceptual experiences seem to be quite different from the contents of ordinary beliefs, Glüer argues that perceptual experiences are beliefs with *phenomenal* contents. Instead of ascribing *sensible* properties such as redness or roundness to concrete objects, phenomenal contents ascribe phenomenal properties such as looking red or looking round (cf. Chalmers 2006). This diffuses the apparent irrationality-generating conflict between an illusory perception and one's better knowledge when both are gauged in terms of beliefs. The two beliefs have different contents and thus do not contradict each other: In the case of the Müller-Lyer illusion, the illusory perception is the belief that one line *looks* longer than the other while one's informed ordinary belief is that both lines are of the same length. No contradiction. Can something similar be said to resolve the conflict between the intuition experience that *Naive Comp* is true and the belief that *Naive Comp* is false? On the face of it: No. What motivated Glüer's argument is the observation that the content of perceptual experiences and ordinary beliefs is quite different. In contrast, the content of intuition experiences and ordinary beliefs seem to be very similar. As Koksvik points out:

If there is a distinction between the kind of content which thought and belief have, on the one hand, and the kind of content which perception has, on the other (conceptual and non-conceptual content, respectively), it seems clear that intuition has the kind of content which thought and belief have, and not the kind which perception has. (Koksvik 2011, p. 105)

Thus even if Glüer's strategy might be successful for perceptual experiences, it is unavailable for intuition experiences. Except, perhaps, if one would *still* insist on specific phenomenal contents for intuition experiences, maybe instead of *looks* contents some kind of *seeming* contents. One *could* do that, I think, but it would lack the plausible motivation that motivated Glüer's position and appear *ad hoc*. The contents of perceptual experiences seem not only quite different from those of beliefs but also from those of intuitions (which are similar to those of beliefs). Why would beliefs have ordinary contents and intuitions phenomenal contents?.

⁵⁵One might try to construct a belief and desire account of intuitions as has been tried for intentions (e.g. Audi 1973; Sinhababu 2013) and emotions (e.g. Marks 1982; Reisenzein 2012) elsewhere. To my knowledge this theoretical possibility has not been worked out. and I won't pursue it here.

has already received it. It is also not clear how one can cast content-gradeability as an integral part of a belief or judgment. Could content-gradeability perhaps be cast in terms of credence or confidence with which a certain belief or judgment is held? Here is a point against this idea:

[V]isual and logical illusions [...] indicate that this is not so: it might be clearly or vividly presented to one that the lines are of different lengths, or that every predicate defines a set, even though one has little or no credence or confidence that this is so (because one knows better). As this suggests, clarity and vividness look to be qualitative features or modes of the presentational state itself. (Bengson 2015, footnote 15)

What about then the mental imagery in form of inner speech or pictorial and grapheme visualizations that often accompany beliefs and judgments? Doesn't this mental imagery at least sometimes appear to have content-gradeability? In fact, doxastic states can often be said to be phenomenally conscious in virtue of mental imagery which might be regarded as an (occurrent) expression or trace of beliefs or judgments. However, such mental imagery appears to be only derivatively linked to the content of beliefs and judgments. In the first place, it is the content of other mental states: *imaginings*. It seems more plausible, then, to attribute content-gradeability to these imaginings, not to the beliefs or judgments that may give rise to them.

Also when it comes to the peculiar feature of phenomenal epistemic valence we encounter tension thinking of intuition experiences as beliefs or judgments. As mentioned in Chapter 2, it is true that beliefs or judgments are able to represent that some proposition is false. Although we can believe that p (is true) and we can believe that p is false (or not p), the belief that p and the belief that p is false seem 1) only to differ in terms of their content and 2) lack a phenomenal contrast. A (conscious occurrent) belief (or judgment) that *P-Zombies* (is true) does not phenomenally contrast with a belief that *P-Zombies is false* or *P-Zombies, not*.

The resulting argument against Simple Doxasticism can be simply summarised as follows:

- (P1) Intuition experiences essentially have a specific phenomenology.
- (P2) The specific phenomenology of intuition experiences is incompatible with them being the state to which they are reduced, even if it is granted that these states essentially have a phenomenology.

(C) Therefore, the reduction fails and Simple Doxasticism is false.

Of course, these remarks do not have to be taken as a knock-out argument against parsimonious Simple Doxasticism. It *may* be that *while* being beliefs and judgments, intuitions do have specific features that distinguish them from other beliefs and judgments. Intuitions might just be fairly different from what we consider *paradigmatic* cases of beliefs or judgments that are normally taken to be evidence-sensitive and to lack special phenomenology (Reimer 2010). One will have to do so avoiding a kind of trivializing “Elastic Strategy” that consists in stretching the notion of belief or judgment in order to accommodate the counterexamples as which intuition experiences appear (Scarantino 2010). As an auxiliary strategy, one might try to argue that some of the described features of intuition experiences that stand in conflict with them being beliefs or judgments are *not* (essential) features of intuition experiences. If this can be done, there might be some hope for Simple Doxasticism.

However, it seems little promising to appeal to qualifications such as the mentioned content or aetiology restrictions, since, if the aim is to individuate a general psychological kind (partly) characterised by its phenomenology, such amendments will appear unduly restrictive, unmotivated and *ad hoc* (Lynch 2006; Koksvik 2011, pp. 23-34, 2013; Chudnoff 2014a; Bengson 2015, footnote 32). It is hard to see what kinds of *content* or *aetiology* restrictions might give us the specific *phenomenal* properties of intuition experiences needed. Discussing content-constraints, Lynch and Koksvik remark:

If intuiting is a distinct kind of attitude, why can't we, given the right circumstances, take up that attitude towards almost any proposition, in the way that, given the right circumstances, we can find ourselves hoping or fearing, or believing almost any proposition? Without argument, it is difficult to see how intuition would be restricted in a more comprehensive way than other attitudes. (Lynch 2006, p. 230)

Bealer regards rational intuition as a subclass of intuitions, a species of the genus. There are many other subclasses likewise singled out by content [...] The question is why we should think that each such subclass, or even just one of them, corresponds to a psychological kind. The simple fact that there is a difference in content is not sufficient on its own; since nothing otherwise stops us from thinking that there are psychological kinds restricted in content in arbitrary ways. And clearly there are not: there is no kind corresponding to fear of bicycles, for example. (Koksvik 2011, p. 28)

In a laudable display of frankness, Sosa instructively admits that there is no deep reason for making such restrictions:

One might quite properly wonder why we should restrict ourselves to modal propositions. And there is no very deep reason. It's just that this seems the proper domain for philosophical uses of intuition. (Sosa 2007b, p. 101)

It seems that such regimentations are best understood as consequences of the metaphilosophical intuition debate and the specific role intuitions play in philosophy. Content and aetiology restriction, thus, are rather reflections of the idiosyncracies of philosophy as a discipline, not of intuitions as a psychological kind.

Finally, talking of aetiology, let me briefly consider a possible way of understanding Simple Doxasticism, that has been so far overlooked among intuition theorists and to which I am somewhat sympathetic. This idea, other than the accounts considered so far, takes dispositional beliefs as a starting point. Understanding beliefs as dispositions enjoys large common ground in contemporary analytic philosophy (Ryle 1949; Geach 1957; Price 1969). It is largely a functional characterization: to have a belief amounts to “being disposed to do and experience certain things” (Schwitzgebel 2002, p. 250).

Now, intuition experiences are things we experience and things we can be disposed to experience. So, among the dispositions characterizing a dispositional belief, we will plausibly find “the disposition to feel assent to an internal utterance of P (or to think silently to oneself, “P”)” (ibid., p. 252). This “felt assent to an internal utterance of P” can be plausibly understood as an intuition experience and the disposition in question as a disposition to experience intuition experiences towards P. Thus, against this background, intuition experiences, at least sometimes, can be plausibly understood as (partial) phenomenal expressions or manifestations of dispositional beliefs without *being* these beliefs (cf. Prinz 2007, p. 350).⁵⁶ Cohen, although not talking of intuitions, hits the mark of this idea:

Belief that p is a disposition, when one is attending to issues raised, or items referred to, by the proposition that p, normally to feel it true that p and false that not-p, whether or not one is willing to act, speak, or reason accordingly. [...] The standard way to discover whether you yourself believe that p is by introspecting whether you are normally disposed to feel that p when you consider the issue. [...] Belief is a disposition, not an occurrent feeling.

⁵⁶Perhaps on other occasions, they might be considered phenomenal expressions of *dispositions to believe* rather than beliefs (Audi 1994).

Although, when you listen, you can hear the relentless downpour through the curtains, you may from time to time stop thinking about the rain. But you do not then stop believing that it is raining—as presumably you would do if belief were an occurrent feeling. (Cohen 1992, pp. 4-5)

In fact, the present thesis is an extended defence of the idea that intuitions are these feelings to feel true or false to which Cohen alludes in making sense of belief. But I get ahead of myself. So, if all Simple Doxasticism wants to say is that there is such a functionally-causal relationship between dispositional beliefs and intuitions, then, I think, there is hope. Nevertheless, this still leaves open the question of what intuition experiences are apart from phenomenal expressions of beliefs (see section 8).

Be it as it may, I think I have established a case for looking beyond the class of doxastic attitudes to find a comfortable home in our ontology of mind for at least some states that are plausibly called “intuitions”. This section argued that Simple Doxasticism as usually proposed does not fit the feature profile of intuition experiences, not even if we grant that some beliefs or judgments (essentially) have a phenomenology. This case gets even stronger against the background that the identification of intuitions with doxastic attitudes is rarely argued for, but rather stipulated. Lewis for example casually and without much argument remarks: “Our “intuitions” are simply opinions [...]. Some are commonsensical, some are sophisticated; some are particular, some general; some are more firmly held, some less. But they are all opinions” (Lewis 1983, p. x). It seems to be a kind of default strategy to assume that any state that is somehow related to truth is a belief. This is not to suggest that the word “intuition” does not extend to some sort of judgments or beliefs but to say that there is something else, perhaps something more “substantial”, it extends to as well and that this something is not properly captured when intuitions are assimilated to doxastic attitudes.

3.3.2. Intuitions as Doxastic Tendencies

This becomes clearer when one considers the accounts that claim that intuitions are not doxastic *attitudes* but doxastic *tendencies* which in contrast to doxastic attitudes can be noncommittal: intuitions are dispositions or inclinations to believe or judge. What it means to be a “doxastic” tendency can perhaps be straightforwardly understood: a tendency *to believe or judge*. This relaxes the notion of being “doxastic” compared to understanding “doxastic” as belief- or judgment-hood. Being “doxastic” is then only insofar informative about the nature of the tendency as it tells us what it is a tendency

for, i.e. what output it can generate or make more likely: namely beliefs or judgments. It tells us something about the *functional* (rather than e.g. phenomenal) profile of the state so classified (cf. Nimtz 2010, pp. 363 sq.; Werner 2014, pp. 1768 sq.). However, many different states can have beliefs or judgments as their outputs. To name a few: perceptual experiences can lead to corresponding judgments and beliefs (or make them more likely), beliefs lead to corresponding judgments, judgments lead to corresponding beliefs.

3.3.2.1. Dispositional Doxasticism

Let's have a look at Dispositional Doxasticism first:

Dispositional Doxasticism: Intuitions are dispositions to believe or judge.

On the face of it, Doxastic Dispositionalist accounts on their own appear unequipped to account for even very basic features described in chapter 2, such as being occurrent and conscious. Being a disposition does not by itself explain phenomenological features. Seemingly, it is rather in tension with a state having a phenomenology at all.⁵⁷ It seems generally accepted that experiential “mental phenomena are necessarily occurrent. There are no dispositional experiential phenomena” (Strawson 2009, p. 159). It is thus unsurprising that Koksvik points out:

[T]he phenomenology of having an intuition is one of the facts about it. It is one of the things that needs to be explained. An account need not preserve appearances in the sense of vindicating our initial view of the phenomenology. But there is pressure to *either* do that, or to explain the appearances *away*: to explain why appearances are misleading. The problem with the dispositional account is that it does not seem capable of doing either. (Koksvik 2011, p. 95; see also Pust 2014)

What should be clear is that by classifying intuitions as doxastic dispositions virtually nothing concerning the phenomenological features of intuition is elucidated. Rather, it remains a fact about intuitions that doesn't square well with them being dispositions and has to be *additionally* explained or stipulated. Note that this is even so if the functional profile of dispositions fits well with the functional profile of intuitions. Thus, classifying

⁵⁷This is not to suggest that there is no sense in which one can be said to be disposed to experience or re-experience an intuition. This is presumably what we mean when we say “Anna has the intuition that torturing kittens for fun is wrong” while Anna is sleeping.

intuitions as doxastic dispositions can elucidate *functional* features of intuitions. All the same: it doesn't fit well with the *phenomenal* profile of intuitions.⁵⁸

3.3.2.2. Inclinal Doxasticism

We are left with the candidate idea that intuitions are inclinations to believe or judge.

Inclinal Doxasticism: Intuitions are inclinations to believe or judge.

However, what are these inclinations? Inclinations do not seem to be some established canonical kind of mental state such as beliefs, desires, perceptual experiences, emotions or intentions. The notion of canonicity is supposed to capture that the existence and place of some mental states such as beliefs and perceptions are less controversial than the ones of others—such as the mentioned inclinations.⁵⁹ On some reading, often invoked by opponents of Doxasticism, they just *are* dispositions. This reading runs into the aforementioned problems. However, this is not the only possible interpretation. For instance, Sosa 2007a defends an account of intuitions as some kind of inclinations to believe or attractions to assent which he calls “intellectual seemings”.⁶⁰ Although he does not make it the centre of his attention he describes these inclinations at some point as being conscious and having a phenomenology:

[I]ntuitive seemings remain *distinctive conscious* states in their own right, without collapsing into beliefs as is shown by paradoxes [...] On this account,

⁵⁸There is also an important caveat concerning the functional role of intuitions. If intuitions are dispositions, then we might be able to explain that intuitions *cause* beliefs or judgments by e.g. making them more likely. However, when it comes to the functional role of intuitions, the central project does not revolve around its causal role (which is, in fact, rather neglected) but around its rationalising or justifying role for judgments or beliefs (Bengson 2015, p. 723). While we can see how dispositions *cause* judgments or beliefs, it is unclear how dispositions could *rationalise or justify* judgments or beliefs: A disposition to judge *that p* might explain *that* I judge *that p* but it does not explain *why* it might have seemed reasonable or justified to me to judge *that p* (Bengson 2015, p. 727; cf. Quinn 1993, pp. 235 sqq.; for this point being made against a dispositionalist account of pain see Bain 2013, S75 sq.). In other words, dispositions are explanatory or motivating but not normative or justifying reasons for judgments (cf. Alvarez 2017).

⁵⁹Neither are inclinations canonical kinds of *phenomenal* states such as perceptual experiences and feelings (prominently comprising bodily feelings, emotional feelings and moods), i.e. mental states that relatively uncontroversially have a distinctive phenomenology.

⁶⁰Note that “Doxasticism” is a term used by opponents of Doxasticism to subsume a rather heterogeneous class of theories. It is questionable whether this charitably captures what is essential about some theories so subsumed. It appears well possible that Sosa's account of intuitions as conscious inclinations is one of the accounts which is not straightforwardly characterised as “doxastic”.

intuition is a *conscious* state of felt attraction (Sosa 2007c, pp. 51 sq., my emphasis).

What are these [intellectual] seemings? It is helpful to compare deliberation on a choice or the pondering of a question, where we “weigh” reasons pro or con. Switching metaphors, we feel the “pull” of conflicting considerations. No matter the metaphor, the phenomenon itself is familiar to us all. There is something it is *like* to feel the pull of contrary attractions as we deliberate or ponder. (Sosa 2007a, pp. 47 sq., his emphasis)

This seems to distinguish inclinations from dispositions.⁶¹ However, we are not told what kinds of states inclinations are. Although Sosa claims that we are familiar with the phenomenon, he does not tell us whether inclinations just belong to some canonical kind of mental state or whether they are a mental kind of their own.

John Bengson is the only one who appears to make some steps in the former direction, drawing on Tamar Schapiro’s treatment of inclinations (Bengson 2015; cf. Schapiro 2009). Schapiro analyses inclinations as cases of “unmotivated” desires characterised as “motivational states that exert an influence *on* the will [...] independently of deliberation”, being “necessarily action-oriented” (Schapiro 2009, p. 230). Against this background, Bengson at some point explicitly diverges from a dispositional reading of inclinations. He subsumes inclinations under “familiar” kinds of states (as apparently Sosa) and labels inclination theories as “doxastic” and “minimalist” (e.g. Bengson 2015, pp. 712 sqq.).

At the same time, Bengson, inspired by Schapiro, suggests that inclinations are akin to conscious desires (e.g. *ibid.*, p. 727, fn 19). What does that mean? Are inclinations (some kind of) desires? If inclinations are some kind of desires, then they are indeed familiar and minimalist but it is hard to see in what sense they are “doxastic” apart from being in some curious way a desire *to believe or to judge*, perhaps similar to wishful thinking. This would perhaps equip Inclination Doxasticism with resources to account

⁶¹Or, if one – seemingly contradictory – takes being a disposition as being *compatible* with having a phenomenology (thus being conscious and *occurrent*), it at least spells out further features of the specific dispositions that intuitions are which are not implicated solely in virtue of being a disposition (cf. Werner 2014, pp. 1768 sq.). This is, being a disposition is a functional description and is wholly uninformative when it comes to phenomenology while it might be *compatible* with having a phenomenology (see main text above). Compare that for example with what we learn about the phenomenology of a mental state *merely* by learning that it is a perceptual experience or an emotion.

for the pushiness of intuition experiences. However, it would sit ill with the assertiveness of intuition experiences, and in fact, with their justificatory epistemic role.

Now, if inclinations are only in some way *similar* to desires but not desires, then it is again open what kind of states they are. Are inclinations perhaps just a kind of their own? Then they might somehow be doxastic but a theory that introduces the new kind of inclinations to our mental ontology would hardly be called “minimalist”. It rather appears “inflationist”. In any case, inclinations turn out to have as unclear a place in our mental ontology as intuitions. That makes inclinations unfit to be a point of departure for explaining the features intuitions have in terms of what kind of states they are, for inclinations themselves would require such an explanation.

To emphasize: If we want to classify intuition experiences by identifying them with other states, it would be progress if this identification would be with states that we know more about than about intuition experiences themselves. This would allow us to tap into a richer body of knowledge about intuitions. On any count, such a classification would have to do more explanatory work than the mere deployment of a suggestive new label such as “inclinations” or an “ontologically free-floating” re-description of the features intuition experiences have (or a combination of both).⁶²

To sum up on Doxasticism: Simple Doxasticism and Dispositional Doxasticism have their merits in being ontologically parsimonious and, in virtue of identifying intuitions with beliefs or judgments or dispositions to believe or judge, they seem to straightforwardly satisfy epistemic desiderata. However, these approaches fail quite thoroughly to account for the features of intuition experiences. On the other hand, Inclinal Doxasticism seems to do slightly better on phenomenological grounds but it leaves us with big question marks regarding ontology (as well as epistemology).

Here is a rough and ready comparison between intuitions and doxastic states that illustrates what is wrong with Doxasticism:

⁶²This is not to downplay the value of careful (phenomenological) analysis of intuitions.

intuition experiences	= beliefs	= judgements	= doxastic dispositions	= doxastic inclinations
intentional	✓	✓	✓	✓
assertive	✓	✓	?	?
pushy	✗	✗	?	✓
non-committal	✗	✗	✓	✓
(c-/p-) gradeable	? ✗	? ✗	✗ ?	✗ ✓
ep valence	✗	✗	✗	✗
non-voluntary	✗	✗	?	?

A comparison of features between intuitions and doxastic states

It appears that Doxasticism as a theory of intuition experiences is insufficient:

Insufficiency of Doxasticism (ID):

- (ID-P1) Doxasticism is true if it can accommodate the feature profile of intuitions.
- (ID-P2) Doxasticism can either not acknowledge the feature profile of intuitions (Simple and Dispositional Doxasticism) or it cannot explain it (Inclinalational Doxasticism).
- (ID-C1) Therefore, Doxasticism cannot accommodate the feature profile of intuitions.
- (ID-C2) Thus, Doxasticism is false (or at least incomplete).

On any count, if intuition experiences are some kinds of doxastic states, they will be rather *atypical* doxastic states or “doxastic” in a sense that rather bends the term. This suggests that in the absence of an independent argument, the stipulation that intuition experiences are doxastic states becomes rather unmotivated and *ad hoc*. And even *if* a case can be made *for* the possibility to identify intuitions with doxastic states this does not preclude that a better case can be made for identifying intuitions with some other kind of states. Either way one does have a case for looking further for a more suitable and informative identification.

3.4. Perceptualist Intuitions

I already noted for intuitions and inclinations that it is possible to just say that there is a kind of its own in our ontology of mind which cannot be assimilated to other kinds.

That is, in the case of intuitions we are talking about a *sui generis* kind of mental state. This brings us to the second camp among intuition theorists to which I now turn. “Anti-reductionists” (Koksvik forthcoming) or “non-minimalists” (Bengson 2015) reject the idea that doxastic states are all there is to intuitions. Such theorists are without exception of a distinctive brand: “Perceptualists” (Chudnoff 2011b), contrary to what the label may suggest, do not propose to *identify* intuitions with perceptual states. Rather, they try to make headway in explicating what intuitions as non-doxastic states consist in by stressing and analysing the *similarities* between intuitions and perceptual experiences.⁶³

With a phenomenological focus, perceptualist theories identify significant similarities between perceptual experiences and intuitions. Perceptualists claim that it is due to the specific phenomenal features of intuitions that doxastic theories fail to capture substantial aspects about intuitions, just as doxastic theories of perception fail. That is, Perceptualists engage in a kind of “phenomenal profiling” when they reject doxastic accounts of intuitions. They purport to show that doxastic states do not fit the phenomenal profile of intuitions. I am sympathetic with this claim and I take it to motivate the rejection of Doxasticism.⁶⁴

Now, what does Perceptualism has on offer? I will first sketch classic Perceptualism, whose original idea is that intuitions are seemings. Then I will briefly review more sophisticated perceptualist accounts that flesh out the simple seeming idea in more phenomenal detail. Finally, I will try to canvas what is right about Perceptualism and what is wrong about it, together with extant intuition theorising in general.

3.4.1. Classic Perceptualism: Intuitions as Intellectual Seemings

The classic statement of Perceptualism is that intuitions are to be located among the class of mental states called *seemings* (e.g. Bealer 1992; Pust 2000; Huemer 2007; Brogaard

⁶³Such proposals conceive of intuitions in a way that doxastic theorists find rather mysterious (Boghossian 2000; Boghossian and Bonjour 2001) and missing in their mental lives (Williamson 2007, p. 217; Cappelen 2012, p. 117). Perceptualists provide error theories for Doxasticists who claim to be unable to find perception-like *sui generis* intuitions among their mental states (Chudnoff 2011b; Koksvik 2011; Bengson 2015). So far as I can tell there are no reports of Doxasticists who, guided by such error theories, were eventually able to find perception-like intuitions among their mental states.

⁶⁴As will become clear, I am not sympathetic to the moral Perceptualists draw from it: they take this to motivate carving a new space for intuitions within our mental ontology. This seems to imply that a phenomenal profiling among other than doxastic states in our mental ontology would not be successful. In this, I think, Perceptualists are mistaken. More on that in section 3.4.3 below and the rest of the thesis.

2014). Consider George Bealer’s statement

By intuition, we do not mean a supernatural power or a magical inner voice or anything of the sort. When you have an intuition that *A*, it *seems* to you that *A*. Here ‘seems’ is understood, not in its use as a cautionary or ‘hedging’ term, but in its use as a term for a genuine kind of conscious episode. For example, when you first consider one of de Morgan’s laws, often you draw a blank; after a moment’s reflection, however, something happens: it now really *seems* obvious. You suddenly ‘just see’ it. It presents itself as how things must be. Of course, this kind of seeming is intellectual, not sensory or introspective. (Bealer 1992, p. 101).

Those adhering to the idea of seemings usually hold that all states where “things seem to one a certain way” are seemings, perceptual experiences most prominently among them.⁶⁵ While perceptual experiences are conceived as sensory seemings where something seems a certain way perceptually to one, intuition experiences constitute the seeming-subclass of non-sensory *intellectual* seemings where something seems a certain way intellectually to one. Thus:

Seeming Perceptualism: Intuitions are intellectual seemings.

At first glance, calling something a seeming does not seem very informative. For what are seemings? And what does calling a state a seeming add to our understanding of this state? We might have somewhat of a grasp of our uses of the word “seem” but seemings are not exactly a canonical mental kind of states. Now, there are roughly three views on what seemings are. The *belief view*, the *inclination view* and the *experience view* (Tucker 2013; McAllister 2018). Looks familiar? Indeed. The belief and the inclination view bring us just back to where we’ve already been. I have already expressed my reservations with the idea that intuitions are beliefs or inclinations. As to the experience view: I am currently dealing with the idea that intuitions are phenomenal states—intuition *experiences*. That intuition experiences have a phenomenal character we already knew. In fact, in chapter 2 we analysed their specific phenomenal features in some detail. Labelling intuitions now intellectual seemings does little to explain these

⁶⁵Huemer, for instance, writes:

I take statements of the form “it seems to *S* that *p*” or “it appears to *S* that *p*” to describe a kind of propositional attitude, different from belief, of which sensory experience, apparent memory, intuition, and apparent introspective awareness are species. (Huemer 2007, p. 30)

features. Having outlined these features to begin with is, in fact, more than most seeming theories provide. The most daring accounts of seemings just state that they are *sui generis* phenomenal states, pointing to perception for examples of seemings. And from there on, it gets metaphorical.

As Chris Tucker illustrates:

Characterizing seemings as experiences may help insofar as people generally have a grip of what experiences are. But the progress is limited. On this view, seemings aren't just any experience: they are a special kind of experience. Okay, but what kind? Well, the kind that has this really neat and distinctive phenomenal character. Okay, but what phenomenal character? It's at this point that proponents of the Experience View resort to ostension/examples (e.g., when you are looking at a tree, it seems to you that a tree is in front of you) and metaphor (a seeming that P "recommends" P as true or "assures" the subject of P's truth). If that's the best the proponent of the Experience View can do, one may well doubt that there is any such distinctive phenomenal character and may wonder whether there is any such thing as a seeming at all (Tucker 2013, pp. 5 sq.).

Those who take intuitions to be intellectual seemings rarely engage in a closer examination of the phenomenology specific to *intellectual* seemings. Instead, they sometimes offer something else to delineate intellectual seemings from sensory seemings: aetiology and content constraints.

I believe that there are two kinds of intellectual seemings: intuitions about how people in my community apply concepts and a priori intuitions. Both kinds of intellectual seemings are ideally grounded in semantic memory. If I judge that a subject in a Gettier case lacks knowledge, I make this judgment on the basis of semantic memory of how people in my community apply the notion of knowledge. Likewise, if I judge that p and q entails p, I make the judgment on the basis of semantic memory of the meaning of the logical connectors. So intellectual seemings are a special case of memory-related seemings. (Brogaard 2013, p. 279)

Semantic memory can here be related to a process of understanding or concept application, as suggested by Bealer (Bealer 1999).

The analysis of concept possession is the final step in our account [...] of the evidential force of intuitions. [...] [T]he explanation is that [...] [intuitions] have the right sort of modal tie to the truth. [...] In the case of intuition, determinate possession of our concepts entails that there must be such a tie. [...] Our intuitions are what seem intellectually to be so concerning the applicability of concepts to cases presented to pure thought. (Bealer 1999, p. 47)

This normative stipulation might be successful in ensuring the reliability of the subset of intuitions so delineated. As we have seen in section 3.3.1, however, it is less helpful in shedding light on what kind of psychological state intuitions are.

3.4.2. Advanced Perceptualism: Intuitions as Intellectual Presentations and Rich Inclinations

The features that I have introduced in chapter 2 were largely sourced from Perceptualists that go a step beyond just stating that intuitions are intellectual seemings (with a constraint supplement). By calling intuitions “presentations” rather than seemings, Bengson and Chudnoff try to characterise what the ascribed *phenomenology* consists in: Intuitions (as perceptual experiences) have *presentational* phenomenology (Chudnoff 2011b, 2013c; Bengson 2015).

Presentational Perceptualism: Intuitions are intellectual presentations.

Here is how Bengson describes presentations:

Intuition is neither a doxastic attitude, such as a belief or judgement, nor a mere tendency to form such an attitude, but rather a *presentation*: a conscious state or event that, like perceptual experience, directly and immediately presents the world as being a certain way. (Bengson 2015, p. 708)

What does the phenomenology of a presentation consist in? Drawing on Bengson and Chudnoff, part of the features in chapter 2 describes it. In describing intellectual presentations, Bengson gives a list of features that is largely covered in chapter 1. He specifically describes presentations as (ibid., pp. 720-723, 730, 749):

1. “*baseless*, in the sense that they are not consciously formed, by a subject, on the basis of any other mental state(s).”

2. “*gradable*: their overall quality may vary in different situations, depending upon the manner in which they present in those situations (e.g. more or less clearly, vividly, etc.)” (see also Chudnoff 2011b, pp. 642 sq.).
3. “*fundamentally nonvoluntary* (i.e. passive or receptive)”
4. “*compelling*, in the sense that they tend to dispose or incline assent to their contents”.
5. “they also seem to *rationalize* such assent, in the (psychological) sense that they tend to make formation of corresponding beliefs seem rational or fitting from the first-person perspective”.
6. “sometimes *inexplicit*: one need not be able at the time to formulate explicitly — out loud or in one’s head — the way things are presented as being.”
7. *translucent*, that is, one sees or intuits that p in virtue of being presented with p , not in virtue of being presented with some other content q (see also *ibid.*, pp. 640, 643).⁶⁶

Some aspects of Chudnoff’s highly complex view of intuitions are worth mentioning. Here is Chudnoff’s attempt to capture presentational phenomenology:

Sturgeon’s claim about visual perception—“what it’s like to enjoy visual experience is for it to be as if objects and their features are directly before the mind”—has a true analogue for intuition: what it’s like to enjoy intuition experience is for it to be as if objects and their features are directly before the mind. (*ibid.*, p. 636)

For Chudnoff, when having an intuition *that p* it seems to one *that p*, and at the same time one is aware of a *truth-maker* for p . Bengson, for instance, rests content with postulating that intuitions are presentational in a sense roughly analogous to assertiveness (complemented by the features he outlines). He does not embrace the commitment that in order to be presentational or assertive, intuitions need to present us with truth-makers

⁶⁶Note that it is not clear what relations these features have among each other. Sometimes there appears to be some redundancy and conflation. For instance, the second sentence of his description of *baselessness* appears to crosscut nonvoluntariness: “In fact, presentational states are not states that one forms at all, whether consciously or non-consciously; rather, one simply has — or fails to have — them” (Bengson 2015, p. 720); inexplicitness seems to be a special form of gradeability; and translucence appears like the reason for presentational states being rationalizing.

for the propositions that are their contents (Bengson 2015, pp. 709, 724, 742 sq.; cf. Teng 2018).

Now, for Chudnoff, such a rich perception-like truth-maker phenomenology is possible because he endorses a “molecular view of intuition constitution”. He claims that, phenomenally, intuitions do not occupy a place of their own as do e.g. imaginings or perceptual experiences. Rather, they arise together with other experiences which, in turn, constitute (but are not identical to) the intuition experience.⁶⁷

perceptual experiences fit into your stream of consciousness like experiential atoms; they are not constituted by your other experiences, such as your imaginings and conscious thoughts. [...] Intuition experiences fit into your stream of consciousness like experiential molecules; they are constituted by your experiences, such as your imaginings and conscious thoughts. (Chudnoff 2014b, pp. 10 sq.)

The phenomenal constituents of intuitions, on the other hand, often stem from what Chudnoff calls “imaginative or cognitive endeavours” (Chudnoff 2011a): foregoing reflections on some subject matter. When these reflections “click together” in the right way, they transform into an intuition experience (Chudnoff 2011b, pp. 645 sqq.). This intuition imbues some part of the content of its phenomenal constituents with a “felt presence to mind” (ibid., p. 637) as described above, marking it as a truth-maker for its proposition. Thereby intuitions also guide or motivate the subject in her mental actions to judge in accordance with what the intuitions present. Thus, he takes intuitions to be descriptive and *directive* at the same time, emerging as “pushmi-pullyu representations” (Chudnoff 2014b; cf. Millikan 1995). That intuitions do not only represent things but guide and motivate our mental behaviour in relation to them is noted independently by diverse authors (Vaidya 2010, p. 399; Bengson 2014, p. 573; Molyneux 2014, pp. 450 sqq., see also section 4.5).

In opposition to Bengson and Chudnoff, Koksvik takes intuitions—*in contrast* to perceptual experiences—to *lack* content-specific and therefore presentational phenomenology

⁶⁷This appears to introduce some tension with Bengson’s requirement that intuitions be baseless. This is one of the reasons why I omitted this feature in chapter 1. Baselessness and the related notions of immediacy or non-inferentiality are tricky and hard to spell out in general. Perhaps the best thing to say is that you cannot infer an intuition into being (because it is nonvoluntary), but clearly, you can have intuitions in the process of inferential reasoning (Koksvik 2013). And you can also *infer* that you have the intuitions you currently have due to your inferential reasoning. However, nothing much will hinge on it here.

(Koksvik 2011, 2017).⁶⁸ Rather, intuitions have attitude-specific phenomenology in common with perceptual experiences.⁶⁹

This phenomenology comprises: 1) *Objectivity*, i.e. “the fact that intuitional experience purports to be about objective, mind-independent, facts” (Koksvik 2017, p. 7).⁷⁰ 2) *Pushiness* describes the feature that an intuition experience “is not neutral with respect to its contents” but (to various degrees) pushes “the subject of the experience to accept its content” (Koksvik 2011, p. 175, 2017, p. 7). Finally, *valence* “reflects whether, when a subject enjoys an intuitional experience, the content of the experience seems true or false” (Koksvik 2017, p. 7). To sum up his view:

When it seems to you that if something is red it is coloured, the phenomenology of objectivity ensures that the state is about an aspect of mind-independent reality, and the content is up for consideration as true, not as false (valence). However, the content is not presented neutrally, merely as a proposition to consider. You are *pushed* to believe it. (ibid., p. 7)

Thus, Koksvik gives us something like:

Inclinal Perceptualism: Intuitions are rich inclinations.

I call Koksvik’s conception of intuitions “rich inclinations” because other than Inclinal Doxasticism it positively identifies more features than just pushiness.⁷¹

⁶⁸In the case of Bengson and Chudnoff it is not clear what parts of their characterizations are supposed to pick out content-specific and what parts are supposed to pick out attitude-specific phenomenology.

⁶⁹It is easy to get confused here. In section 2.2.1 I noted that Koksvik takes a lack of content-specific phenomenology to preclude that intuitions can have phenomenal intentionality (taking presentational phenomenology as a kind of it). He does grant attitude-specific phenomenology to intuitions, however. This reverberates my remark in footnote 24: Koksvik, as almost everyone else, conceives of phenomenal intentionality as a thesis exclusive to content-specific phenomenology, i.e. only content-specific phenomenology can have phenomenal intentionality. By implication, there is no phenomenal intentionality to attitude-specific phenomenology. I noted in footnote 24 that it does not seem obvious that phenomenal intentionality has to be restricted in this way.

⁷⁰This is a feature that I have not put on the list in chapter 2. I make some remarks on why in footnote 169 and 204. The general point is: I think that Koksvik significantly underestimates the degree to which our mental states *appear* as subject- or mind-independent to us. Thus, objectivity does not seem to be as diagnostic of intuitions as Koksvik suggests. However, it can be accommodated by the account I am about to provide (see section 6.4). Alternatively, objectivity can be approximated by a state being assertive *and* nonvoluntary. This, at least, seems to capture the contrast Koksvik tries to illustrate by comparing intuitions, which he takes to be pushy (he sometimes confounds pushiness with assertiveness) and objective, with wishful thinking, which he takes to be pushy but not objective (Koksvik 2011, p. 200). I think that wishful thinking might appear as objective as intuitions. This is *why* it is a problem. However, in contrast to intuitions, wishful thinking seems to be a result of the subject’s will and to some degree under voluntary control (see also footnote 204). This is why we are often to blame for wishful thinking.

⁷¹In fact, Koksvik’s conception might perhaps be understood as making sense of what the doxastic

3.4.3. Issues with Perceptualism (and Doxasticism)

I could now set out to criticise the various perceptualist views in their details.⁷² However, I don't think that this is necessary. By and large, I am sympathetic with the central tenets of (especially advanced) Perceptualism. I want to emphasise that I agree with most things Perceptualism has to say: that intuitions have a characteristic phenomenology and, more specifically, that they have the phenomenal features outlined. Here is a rough and ready comparison between the features of intuition experiences as outlined in chapter 2 and the states Perceptualists put forward to account for them:

intuition experiences	~ perceptual experiences	= Bengson's presentations	= Chudnoff's presentations	= Koksvik's rich inclinations
intentional	✓	✓	✓	✓
assertive	✓	✓	✓	✓
pushy	?	?	?	✓
non-committal	✓	✓	✓	✓
(c-/p-) gradeable	✓ ?	✓ ?	✓ ?	? ✓
ep valence	✗	✗	✗	✓
non-voluntary	✓	✓	✓	✓

A comparison of features between intuitions and perception-like states

As you can see, Perceptualists do a fair job in acknowledging the features of intuitions. This is why I agree with them on what they have to say. However, what I (together with Doxasticists, I presume) am less content with, is where this leaves intuitions in our ontology of mind. This is, I am worried about what Perceptualists do *not* say. Let me explain.

Perceptualists do not make the implausible claim that intuitions just *are* another perceptual modality next to vision, audition, proprioception etc. To my knowledge, there is no one who endorses such a kind of *Literal Perceptualism*. On the contrary, intuitions are said to be essentially non-sensory. Put differently: While Doxasticists tend to say that intuitions just *are* (a certain kind of) doxastic states Perceptualists claim that they are *similar to but not quite* perceptions. This is precisely what marks the perceptualist approach as *non-reductive* and comparative, using perception as an *analogy*. This is, on

inclinations in section 3.3.2.2 consist in—if they aren't desires.

⁷²For instance, what about epistemic phenomenal valence in other accounts than Koksvik's (for a problem with Chudnoff's view see footnote 73)?

any count, what *Comparative Perceptualism* rests content with. Chudnoff and Koksvik are in this camp because they look at perceptual experiences in order to outline similarities but also differences between perceptual and intuition experiences. It is true that due to intuitions' seemingly epistemically virtuous similarities with perceptual experiences and apparent lack of "defeater features", Chudnoff and Koksvik consider intuitions and perceptual experiences *epistemically* on a par (for the same point concerning the analogy between perception and emotion see Oliver-Skuse 2016, p. 28). At the same time, Chudnoff considers perceptual experiences as phenomenally atomic but intuitions as phenomenally molecular. For Koksvik, on the other hand, perceptual experiences have and intuitions lack content-specific phenomenology. Thus, they do not seem to be plausibly understood as putting intuitions and perceptual experiences in the same *metaphysical* box.

Quasi-Perceptualists such as Bealer, Huemer and Bengson take the perceptual analogy a step further. They postulate a new superordinate kind of mental state to which *both* perceptions and intuitions belong. Intuitions and perceptions are then assimilated to the non-canonical class of seemings or presentations. This can be thought of as a hybrid strategy combining a suggestive new label with a re-description of intuition features guided by the perception-analogy. Such a move, however, seems exclusively motivated by a need to carve an ontological (and epistemic) niche for intuitions on a par with perception, resulting in a contentious form of *reductionism* of intuitions *and* perceptual experiences to an ontologically dubious class of states.

There is an asymmetry in explanatory value for the assimilated states: perception is used as a model for e.g. presentations to function as a projection surface and to embody the properties it has in common with intuitions. Our notion of perceptual experiences is not informed (but rather obscured) by assimilating them with presentations while our notion of intuitions is mostly constituted by it. The superordinate kind is used to ontologically "launder" and transpose the properties of perceptual experiences to intuitions, creating an air of explanatory value of such a linking common category. However, effectively we are faced with a mystery in order to explain another mystery and this new mystery appears to shake the ontological status of something that seemed relatively firm on the ground: perceptual experiences.

Be that as it may, the problem with Perceptualism is not that it is unable to provide an informative characterization of the features of intuitions but that it leaves them essentially unexplained. While there are good explanations for why perceptual experiences

have the features they do, these explanatory resources are not available for intuitions. This is the downside of a mere analogy. And this problem is not solved by dressing the analogy in a suggestive label (e.g. “seemings” or “presentations”) and postulating a superordinate kind that is no less mysterious than intuition experiences themselves. The Rutherford-model of atoms can informatively characterise the interaction between electrons and the atomic nucleus in analogy with the interaction between low- and high-mass celestial bodies — but this *illustrating* analogy does not *explain* the interaction. Would it be taken to do so, it would be *radically* mistaken. Analogies can become *explanatory* when they tap into a connection in kind such as an identity relation, that is, when they turn out to be *more* than analogies. This is where Quasi-Perceptualism tries to get without success.

Where does this leave intuitions in our ontology of mind? Effectively it leaves intuitions on a place of their own, “ontologically free-floating”, characterised by having certain features which themselves are not (yet) explained. Perception here is only used as an analogy that proves helpful for a *comparative* approach. In contrast to Doxasticists, however, it is not that Perceptualists *identify* intuitions with other kinds of states. So intuitions emerge as *sui generis* mental states with some similarities to perceptual experience (and doxastic states). This leaves us with little resources to explain the features ascribed to intuitions on the basis of what kind of state intuitions are.⁷³ In other words, we find ourselves in a Swampman-/Orfalo-like explanatory situation.

A significant consequence of an ontologically free-floating status of intuition experiences is that Perceptualism appears to be an ontologically unstable position. Put differently: Perceptualism lays open to reductive attacks. Now, as we have seen, Perceptualism has little to fear from Doxasticism. This is because Doxasticism seems to be ill-equipped to lead a successful offensive along this flank due to its inability to overcome the main perceptualist defence line: the phenomenal features of intuitions. However, forms of Non-Doxasticism might fare better on this count. It is one of the professed goals of this thesis to show that Affectivism is such a promising form of Non-Doxasticism.

⁷³To my knowledge only Chudnoff among those who describe intuitions as *sui generis* provides something of a story about what kind of state intuitions are in contrast to doxastic states and perceptual experiences, namely experiential states constituted by other, more familiar (and atomic) states such as thoughts or imaginings. This idea faces many challenges. One of them is: molecular intuitions appear very different from atomic perceptual experiences. In fact, it appears also very different from any doxastic states. Is there *anything* in our ontology of mind that seems to have such a molecular phenomenal structure? Since if not, then intuitions appear truly very special, not to say mysterious. However, there seems to be good news for Chudnoff. I, at least, can think of one class of states that sometimes appear to be composed in something like a molecular way. This kind of states are affective experiences (see chapter 4, especially section 4.4.4).

To sum up on Perceptualism: It is tailor-made to fit the phenomenal profile of intuition experiences. In virtue of choosing perceptual experience as an analogy one might dare to hope that we might obtain similarly favourable epistemic recommendations. Eventually, however, it confronts us with a mystery when it comes to finding a place in our ontology of mind for and explaining the features of intuition experiences. This makes the casual epistemological lessons from the perceptual analogy somewhat dubious and leaves Perceptualism vulnerable along its ontological flank. Perceptualism, therefore, appears to be an insufficient theory of intuition experiences.

Insufficiency of (Comparative and Quasi-) Perceptualism (IP):

- (IP-P1) Perceptualism is true if it can accommodate the feature profile of intuitions.
- (IP-P2) Perceptualism can acknowledge the feature profile of intuitions.
- (IP-P3) Perceptualism cannot explain the feature profile of intuitions.
- (IP-C1) Therefore, Perceptualism cannot accommodate the feature profile of intuitions.
- (IP-C2) Thus, Perceptualism is false (or at least incomplete).

Of course, this is not to say that intuitions must necessarily reduce to some other kind of state in order to be (to some degree) understood. It might well be that intuitions are just *sui generis* as beliefs or desires. Then they will have to be understood on their own terms. So of course, one might just settle with this somewhat pre-theoretical (re-descriptive) state and hope that we will be able to explain the features of intuitions in the advent of further exploration (cf. Maynes and Gross 2013, p. 717). Surely, there was a time where we could only describe the superficial properties of perceptual experiences without deeper insight. However, before one comes to terms with that, one might question whether the conclusion that intuitions are *sui generis*, i.e. distinct from all other kinds populating our mind, must be accepted. In fact, it clearly must not: The conclusion is based on two points whose validity I will grant. The first point consists in the perceptualist verdict that doxastic accounts⁷⁴ fail to account for important phenomenological features of intuitions (which Perceptualists claim to identify). The second point follows from the need to abstain from upgrading the perceptual analogy to something like an identity claim to remain on the plausible side of things: intuitions are distinct from perceptions. The argument then looks as follows:

⁷⁴At least if one ignores the possibility to claim that intuitions are somewhat atypical doxastic states, different from paradigmatic doxastic states such as beliefs or judgments.

- (P1) Intuitions are (similar but) ontologically distinct from doxastic states.
(P2) Intuitions are (similar but) ontologically distinct from perceptual experiences.
(C) Therefore, intuitions are a *sui generis* kind

It is easy to see that the conclusion does *not* follow if *sui generis* does not merely mean distinct in kind from doxastic and perceptual states.⁷⁵

But surely there is more to our ontology of mind than doxastic and perceptual states. In fact, seen in this light one might wonder whether the approach taken so far and inherent in the two premises, that is, comparing (or reducing) intuitions either to doxastic states or perceptual experiences, is even suitable *if* one grants the conclusion that intuitions are a *sui generis* kind. As Chudnoff points out: “Most likely, however, any comparison to other states will be more or less misleading, since intuitions are *sui generis* mental states” (Chudnoff 2014b, p. 25). So why are intuitions taken to allow only for an informative comparison with doxastic and perceptual states?

No doubt, doxastic and perceptual states appear *prima facie* as natural candidates *if* one thinks the nature of intuitions back from their central epistemic role (within philosophy). Both, doxastic and perceptual states, are widely accepted as intimately connected with truth (or what we take to be true) and consequently with our evidential and justificatory practices. They are thus by default *hopeful* candidates since they are epistemically well established. However, apart from the epistemic job some philosophers want intuitions to do, there is little to back restraining one’s analysis in such a way.

For why are the employed comparisons the right ones? Why are the similarities between intuitions and perceptual experiences more characteristic for intuitions than the dissimilarities or why are the similarities to e.g. beliefs less characteristic than the similarities to perception? In fact, it is easy to see that even if there are similarities between intuition experiences and perceptual experiences, there are still quite some dissimilarities. That is, the comparative approach is not wholly successful. For example, perceptual experiences

⁷⁵ *If* this is the only sense in which intuitions are *sui generis* then the conclusion does follow. At the same time, those who favour this conclusion on such a basis must remain (sympathetically) open to approaches that look beyond doxastic and perceptual states to identify intuitions among the states populating our mind. (Or they must reject that our ontology of mind has more on offer than doxastic and perceptual states.) For example, identifying intuitions with some kind of affective experiences would establish that intuitions are ultimately reducible to (a subset of) other (well-known) states in our mental ontology (in fact, this will be exactly my approach here). This would not preserve intuitions as being *sui generis* in a broad sense of being distinct in kind from all other kinds of mental states but it would preserve the narrow sense of intuitions being *sui generis* by being distinct in kind from doxastic and perceptual states.

do not seem to have phenomenal epistemic valence. Also, perceptual experiences, unlike intuition experiences, do not seem to admit of abstract contents. Furthermore, our memorial and imaginative capacities associated with perceptual modalities seem different from the ones associated with intuition: Memories and imaginings with a perceptual content exhibit a remarkable isomorphism to perceptual experiences, i.e. sensory memories and imaginings are quite similar to perceptual experiences (Nanay 2015). It is hard to think about something similar for the case of intuition.

Furthermore, employing the perception-analogy introduces the risk to overstretch the analogy by cross-importing features from perceptions to intuitions and vice versa that the other state, in fact, does not have (see also 4.3.1). For example, content-gradeability is obvious in perceptual experiences but much less so in intuitions. In what sense can the content of an intuition be clear or hazy if it is similar (or identical) to the contents of conscious thoughts (Koksvik 2011, p. 105, see also footnote 54)? To take an example that has the reverse import direction: It appears obvious that intuitions have pushiness. It is less obvious for perceptual experiences. So even perceptualist accounts that model intuitions on perceptual experiences do not fit the phenomenal profile of intuitions that well (while nevertheless fitting it better than doxastic accounts). Crucially, as noted before, these accounts do not seem to *explain* intuitions and their features but only re-describe them.

So why think that the comparisons will shed light on all the relevant dimensions of intuitions? Isn't it highly plausible to assume that if intuitions are in fact a mental kind of their own, that there will be features to intuitions that have no counterpart in the compared states? And that these unparalleled characteristics might be ultimately the most essential?

As soon as one considers the features of intuition experiences, doxastic states do not easily fit the profile and perceptual experiences merely allow for a superficial analogy that does not deliver a more substantial grasp on intuitions. Thus, in order to retain an open mind on the nature of intuitions a premature restriction to doxastic and perceptual states is to be avoided. Perhaps the approach to intuitions has been a deal too narrow so far, presenting us with a forced choice that might conceal important characteristics of intuitions that are either unparalleled in our mental ontology altogether or just have no corresponding dimension in doxastic and perceptual states. This leaves open the possibility that intuitions can be assimilated to kinds of states that we already know but that are distinct from the "usual suspects" of intuition theorists. We might find out

by looking beyond for a class of states that fit the profile already outlined by intuition theorists. I will look to such a class in the next chapter.

Looking for Intuitions Elsewhere: Appealing to Feelings

4.1. Introduction: Where to go from here?

Intuition experiences do not seem straightforwardly identified with doxastic states, nor are they a kind of perceptual experience. It seems, then, that they are a *sui generis* psychological kind. However, as we have seen in the previous chapter, this is not a foregone conclusion. At least not, if one is willing to give intuition experiences a ride outside the realm of doxastic and perceptual states.

One might now wonder, where shall we go from here? Or even: where *can* we go from here? Here is a consideration: intuition experiences are phenomenal states. This is why a Doxasticist analysis appears unappealing: doxastic states are not paradigmatically phenomenal. This is also why a perceptualist analysis is initially appealing: perceptual experiences are phenomenal states—in fact, they are *canonical* phenomenal states. So an idea might be to try to look for other kinds of states that are canonical phenomenal states. Surely, perceptual experiences are not the only kinds of canonical phenomenal states. Two other ideas come to mind: imaginings and affective experiences. Now it does, in fact, seem that imaginings appear to stand sometimes in intimate relation with the occurrence of intuitions. As suggested by Chudnoff, imaginings can lead up to intuitions or even partly constitute them. Furthermore, imaginings, as intuitions and perceptual experiences, appear to be intentional states and have content-gradeability. In

contrast to perceptual experiences and similar to intuitions, imaginings can also be about abstract subject matters (see also section 7.4). Ultimately, however, it does not seem plausible that imaginings *are* intuitions. Imaginings starkly contrast with intuitions in their phenomenal profile: they are neither assertive, pushy, valenced or fundamentally nonvoluntary.⁷⁶ This leaves us with affective experiences (or feelings) of which paradigmatic examples are orgasms (bodily feeling) or fears (emotional feeling). You might think now: Intuitions — feelings? Really? I can see why the idea that intuitions are (some kind of) feelings can appear odd, especially coming from philosophy. However, it does get some support from what seems to be a folk theory of intuitions.

Note that I do not mean (a theory of) the various ways we use the term “intuition” and its cognates in everyday life. Neither do I mean “intuitions” that laypeople have *about something*, say, “freedom”. I do really mean a folk theory in the sense of what laypeople would say intuitions are when asked: “What do you think intuitions are?”⁷⁷ If asked the same question for, say, freedom, laypeople would presumably not (primarily) go on about the ways how they use the word “freedom” and “free” in everyday life. They’d tell me what freedom is for them.

Now the folk theory of intuitions that I have often encountered in casual conversations (of course, unbiased by me) is: intuitions are feelings. You don’t have to rely on my testimony. I am sure you can find your own anecdotal evidence, even if you’re a philosopher. Just ask some of your (non-philosopher) friends. Or you can look it up in a dictionary. In the Oxford Online Dictionary you would find this:

⁷⁶As far as I can tell only Chalmers has once suggested something along the lines that (modal or conceivability) intuitions *are* imaginings (Chalmers 2002). I, together with Chudnoff, think that imaginings can, in fact, be important in bringing about intuition experiences, especially when it comes to abstract subject matters (see section 7.4). I conjecture that it is this — at times tight — relation between imaginings and intuitions that explains Chalmers’ suggestion. Alternatively, we might simply add the imaginings mentioned by Chalmers to the mental states that some people sometimes call “intuitions”, without them picking out intuition experiences.

⁷⁷So perhaps I *do* mean laypeople’s “intuitions” about something, namely intuitions.

intuition, NOUN

1 [*mass noun*] The ability to understand something instinctively, without the need for conscious reasoning: *‘we shall allow our intuition to guide us’*

1.1 [*count noun*] A thing that one knows or considers likely from instinctive feeling rather than conscious reasoning: *‘your insights and intuitions as a native speaker are positively sought’*

intuitive, ADJECTIVE

1 Using or based on what one feels to be true even without conscious reasoning; instinctive: *‘his intuitive understanding of the readers’ real needs’*

1.1 (chiefly of computer software) easy to use and understand.

I think this hints quite clearly on the mentioned folk-psychological “feeling theory of intuitions”.

Indeed, it is a curious datum about intuitions that philosophers appear negligent of a readily available folk theory. In the philosophy of mind, it is rather common to use existing folk theories at least as a starting point. Consider, for instance, Levy contemplating explanations for what he calls neurotic anxiety (NA):

Explaining NA cases does not require the postulation of any mental states or processes other than those countenanced by folk psychology. I do not doubt that folk psychology is unable to account for the full range of [...] cases, but, given that we ought to avoid multiplying mental states unnecessarily, we need to ensure that we are postulating exotic states and processes only when they are truly needed. (Levy 2016, p. 9)

Why not for intuitions?⁷⁸⁷⁹ I conjecture that it is due to the rather narrow metaphilosophical and epistemological perspective most philosophers take on the subject. Especially when it comes to epistemic matters, many philosophers tend to put feelings in

⁷⁸But see McGahhey and Van Leeuwen 2018 for one exception.

⁷⁹It also does not appear plausible that folk psychology can be reduced to beliefs and desires, as some philosophers sometimes seem to suggest. Surely, at the very least, perception, imaginings and feelings are part of folk psychology as well.

second place or even treat them as epistemically pernicious (Brun et al. 2008).⁸⁰ As Michael Stocker polemically states:

Much of contemporary philosophical psychology is inadequate, and pathologically so. It omits, denies, or radically misunderstands affectivity or feeling. (Stocker 1983, p. 5)

Meanwhile, psychologists *did* pick up on the folk theory and are, as we will see in the next chapters, busy trying to understand what is behind these feelings addressed in the folk theory of intuitions.⁸¹

So let's examine affective experiences closer and try to square the idea that intuition experiences might be some kind of feelings. What are affective experiences? A first extensional stab is to point towards bodily and emotional feelings, two prominent subclasses of affective experiences. However, what do intuitions have in common with bodily feelings such as wrist pain or emotional feelings such as sadness about one's wrist pain? More than is apparent on first sight – but to see this, we need to identify some of the crucial commonalities that unite various affective experiences into one class. These commonalities are found in the specific phenomenal and intentional features of affective experiences. This will be the topic of the present chapter.

In the course of this chapter, I will introduce the conceptual resources used to analyse affective experiences. In the next chapter, these will then be put to work to analyse epistemic feelings which I claim to be a class of affective experiences. In chapter 5 I hope to show that we can identify intuitions, as characterised in chapter 1, as a type of epistemic feelings and thus as affective experiences as well.

4.2. Clarifying Affective Experiences

Some clarifications at the outset: If not stated otherwise I use “feeling” and “affective experience” synonymously here, but some use “feelings” more broadly as synonymous

⁸⁰Work on *moral* intuitions and *moral* epistemology should be mentioned as a possible exception here (e.g. Prinz 2007; Greene 2008; Kauppinen 2013; Dancy 2014; Clavien and FitzGerald 2018).

⁸¹One has to distinguish here. Not all psychologists working on “intuition” take conscious states as their point of departure. Many psychologists are not primarily concerned with intuitions understood as specific conscious states but rather with broadly construed, (typically) unconscious intuitive *processes* marked by their automaticity and working-memory-independence (e.g. Evans and Stanovich 2013). However, some (but, of course, not all) of the *outputs* of these processes are dubbed feelings and intuitions (e.g. Glöckner and Wittman 2010; De Neys 2012).

with “qualia”, i.e. with phenomenal qualities in general. Note that in calling epistemic feelings epistemic *feelings* I follow a conventional practice. However, I do not intend it as a matter of definition that “feelings” in epistemic feelings are to be understood in my sense and synonymous with affective experience. I will *argue* that epistemic feelings are affective experiences in chapter 6.

To complicate matters, others use “feelings” synonymously with “bodily sensations” understood as (conscious) somatic representations, i.e. representations of bodily events. The prominent feeling theory of emotions claims that emotions are essentially “feelings” in this sense (James 1884; Lange 1887; Damasio 1994; Prinz 2004b). This sense of feelings risks to lump together two things I want to keep apart: bodily feelings such as wrist pain, hunger or relaxation and bodily sensations proper such as just feeling one’s heartbeat, stomach at work or breathing in and out. Only the former are uncontroversially affective experiences or feelings in my sense.⁸²

That there is a distinction to be drawn is indicated by empirical work on bodily pain. Pain research has long recognised that bodily pain, a bodily feeling and paradigmatic affective experience, is decomposable into a “sensory-discriminative” and an “affective-motivational” component (Price 2000; Auvray et al. 2010). Crucially, we observe a double dissociation between these components (e.g. Rubins and Friedman 1948; Berthier et al. 1988; Corder et al. 2019).

In pain asymbolia, subjects show a striking indifference to pain. Importantly, they do this while fully aware of the pain, or at least of its sensory-discriminative aspect. They have the sensation of pain, but they are not affected by it in any way—the pain doesn’t hurt (Bain 2014; Klein 2015). On the other hand, pain affect without pain sensation has been documented in humans and rats (Ploner 1999; Uhelski et al. 2012). Ploner and colleagues report on a patient with lesions in the right somatosensory cortex following a stroke:

Following cutaneous laser stimulation “pain sensations [for right hand and both feet] were characterized as ‘pinprick-like’ and were well localized within

⁸²The feeling theory of emotions thus makes commitments when it claims that certain affective experiences — emotional feelings — are reducible to bodily sensations. If this point is made for emotional feelings, it would be surprising not to see it extended to bodily *feelings*. If one generalizes this point over to affective experiences, this would be to say that all affective experiences (in my sense) might turn out to be bodily sensations of some kind or the other. Though I want to leave open this possibility, I also want to note that it is controversial. The remarks I am about to make in the main text might be thought to make it also *implausible*.

2–3 cm. For left hand [...] the patient spontaneously described a ‘clearly unpleasant’ intensity dependent feeling emerging from an ill-localized and extended area ‘somewhere between fingertips and shoulder’, that he wanted to avoid. The fully cooperative and eloquent patient was completely unable to further describe quality, localization and intensity of the perceived stimulus. Suggestions from a given word list containing ‘warm’, ‘hot’, ‘cold’, ‘touch’, ‘burning’, ‘pinprick-like’, ‘slight pain’, ‘moderate pain’ and ‘intense pain’ were denied nor did the patient report any kind of paraesthesias” (Ploner 1999, p. 213).

Generalizing from pain to other bodily experiences one can distinguish between cases of “hot” affective bodily feelings and “cold” non-affective bodily sensations (cf. Proust 2015, pp. 20 sq.; Gerrans 2015). As Frédérique de Vignemont aptly points out: “self-regulatory interoceptive feelings such as thirst and hunger [...] provide a narcissistic colouring to our bodily sensations. All sensations are experienced not only on the spatial background of the protective body map but also on the affective background of our interoceptive feelings” (de Vignemont 2018).

It is obvious that bodily feelings *typically* come with bodily sensations. What I am saying is that the affective and the bodily component appear to be in principle dissociable and that the former seems not wholly reducible to the latter. In other words, if one subtracts all the bodily sensations from a bodily feeling, there might be a non-somatic remainder—a genuinely *affective* component (see section 4.3).

You might wonder: Why do I belabour this point? That’s because the topic of the present thesis is intuition. Now, *if*, as I argue, intuitions are affective experiences, then they appear to be affective experiences which, at least consciously, do not have a close liaison with the body. In other words, on a personal level, they appear rather “disembodied”, “psychic” (Stocker 1983), felt “as though they were in the soul itself” (Descartes 1649, section 25). Insofar I see myself under pressure to establish some initial plausibility for an affective experience that appears that way and still is genuinely affective. So let me belabour this point some more.

In fact, the point made seems to carry over from bodily feelings to affective experiences in general (cf. Stocker 1983; Mitchell 2018). The Cotard syndrome is a point in case. It is known to produce the Cotard delusion where patients come to believe that they are dead or do not exist. There is something in the experience of Cotard patients that gives rise and explains this delusion: Cotard patients display a general loss of affect.

At the same time, there appears to be no impairment in their somatosensory awareness (Michal et al. 2014). In other words, Cotard patients' bodily sensations are intact but their conscious life has been purged of affective experiences:

The delusion is a personal level response to an intractable and impenetrable loss of affective response to the world [...] In the Cotard syndrome something is amiss with the mechanisms that appraise perceptual and interoceptive information for self-relevance [...] Thus, felt significance disappears. When the depressive patient then focuses on her experience she feels alienated from the world and depersonalised. (Gerrans 2015, pp. 4, 14)

To be sure, there are good reasons that make bodily states (that are represented by bodily sensations) of particular significance to an organism. After all, they are *literally* tied to an organism's biological integrity. Nevertheless, the upshot here is that what is needed for affective experiences are not just (and not necessarily) bodily sensations but something that adds significance to these (and other) states. So while bodily sensations are usually involved in many affective experiences, they appear to be neither necessary nor sufficient to make an experience an affective experience. There is something over and above mere bodily sensations when it comes to affect.

The Jamesians will now probably flock to the streets in protest. What can I say to them when they knock on my door? Well, not much, I presume—except, perhaps, this: I think that William James has been misunderstood in his treatment of emotions. Many interpreters have taken him to say that emotions are reducible to bodily sensations. But here comes an eye-opening passage:

I should say first of all that the only emotions I propose expressly to consider here are those that have a distinct bodily expression. That there are feelings of pleasure and displeasure, of interest and excitement, bound up with mental operations, but having no obvious bodily expression for their consequence, would, I suppose, be held true by most readers. Certain arrangements of sounds, of lines, of colours, are agreeable, and others the reverse, without the degree of the feeling being sufficient to quicken the pulse or breathing, or to prompt to movements of either the body or the face. Certain sequences of ideas charm us as much as others tire us. It is a real intellectual delight to get a problem solved, and a real intellectual torment to have to leave it unfinished. The first set of examples, the sounds, lines, and colours, are either bodily sensations, or the images of such. The second set seem to

depend on processes in the ideational centres exclusively. Taken together, they appear to prove that there are pleasures and pains inherent in certain forms of nerve-action as such, wherever that action occur. The case of these feelings we will at present leave entirely aside, and confine our attention to the more complicated cases in which a wave of bodily disturbance of some kind accompanies the perception of the interesting sights or sounds, or the passage of the exciting train of ideas. Surprise, curiosity, rapture, fear, anger, lust, greed, and the like, become then the names of the mental states with which the person is possessed. The bodily disturbances are said to be the “manifestation” of these several emotions, their “expression” or “natural language”; and these emotions themselves, being so strongly characterized both from within and without, may be called the *standard* emotions. (James 1884, p. 189)

So it turns out that James’ feeling theory was intended as a theory of what he calls “standard emotions”, not of *all* emotions. James grants that among emotions we also find those that have “no obvious bodily expression” and “depend on processes in the ideational centres exclusively”. It stands to reason to assume that intuition experiences as characterised in chapter 1 will be best construed as such “nonstandard emotions” (Stocker 2009). In any case, there is something independent of bodily sensations that nonstandard emotions have in common with standard emotions, something that warrants them both being *emotions* (or affective experiences): This something is plausibly the remainder that also distinguishes bodily feelings from bodily sensations. Note that when I talk of affective experiences here I do not mean to *only* refer to this remainder but to the whole experience, possibly including a bodily sensation and/or other sensory and non-sensory components. I stress the remainder because it makes the experience an *affective* experience and what makes a nonstandard emotion still an emotion in James’ view. I will take up the issue of what this remainder consists in shortly.⁸³

But first some further clarifications: I chose to talk of “affective” instead of “emotional” experiences. It is true that when it comes to affective experiences or feelings most theorizing is found in the literature on emotions. However, for now, I do not want to take a stance on the nature of emotions *per se*. Rather, I am interested in emotions insofar as they help shed light on emotional *feelings* and thereby on affective experiences. I don’t

⁸³To anticipate: This remainder should not be associated with something like the non-intentional phenomenal “add-on” component postulated by the “add-on emotion theories” criticised by Peter Goldie (Goldie 2000, Ch. 3).

deny that there can be more to emotions than “just” emotional feelings; emotional dispositions, for instance. When I use terms such as “bodily pain”, “joy” or “pride” here, I mean to refer to the corresponding *feelings of* bodily pain, joy or pride. I take the class of affective experiences or feelings to comprise all occurrent conscious affective states. At the same time, the whole class of phenomena that some call “affective” or “emotions” might comprise more than only affective *experiences*. It might be taken, for example, to also include affective dispositions, sentiments, unconscious appraisals and action tendencies (e.g. Shand 1914; Broad 1954; Arnold 1960; Frijda 1986; Lazarus 1991).⁸⁴ As such the class of affective experiences is larger than the class of emotional feelings but smaller than the class of affective phenomena. *Among other things*, emotional feelings such as joy, anger or sadness do not comprise the mentioned bodily feelings such as joint pain, hunger, thirst or relaxation.

To sum up: affective experiences or feelings are understood here as forming a more restrictive category than the general category of phenomenal states but as encompassing more than just emotional feelings. It is, furthermore, an open question whether this subset has bodily sensations as a constitutive feature (cf. Stocker 1983; Mitchell 2018). As dissociations between the sensory and the affective component of bodily pain demonstrate, bodily sensations might be neither necessary nor sufficient for a mental state to qualify as an affective experience. I will now turn towards a positive characterization of affective experiences.

4.3. Affective Phenomenology

As should be clear by now, by distinguishing affective experiences or feelings from qualia and bodily sensations I by no means want to negate that there is something it is like to have an affective experience or that, at times, *part* of what it is like to have an affective experience is to have bodily sensations. Seeing blue or feeling one’s heartbeat has a phenomenal quality, and so do feeling pain in one’s wrist and feeling sad about it. All four have something important in common: they are all essentially phenomenally conscious. Yet, the former two are not affective experiences while the latter two are. What distinguishes non-affective from affective experiences? In other words, what is this remainder I was talking about?

⁸⁴One might want to call these “affective” phenomena derivatively affective in my sense due to their constitutive relation to affective experience in giving rise to or being the result of affective experiences (cf. ‘connection principles’ Kriegel 2014, pp. 429 sqq.).

4.3.1. The Hallmark of Affect: Phenomenal Valence

Perhaps the most central aspect is phenomenal valence, i.e. the felt positivity or negativity of certain experiences (e.g. Charland 2005; Weiss 2016; Carruthers 2017b). This basic positivity or negativity is often made sense of in hedonic terms (“hedonic tone”) as pleasantness or unpleasantness or in value terms as a representation of value or disvalue or “seeming” goodness or badness (Carruthers 2017b). Affective experiences are essentially *valenced* experiences.⁸⁵ As Frijda points out, the property of valence is essential to affective experiences (that is a broader category than emotional feelings) and contrasts them to other experiences:

[T]he essence of emotions is feeling, notably that of pleasure or pain. Affective valence is commonly regarded as a criterial aspect. [...] Affects, pleasure and pain, certainly set the experiences in which they occur apart from all other kinds of experience—if only because, as feelings, they cannot be readily reduced to something else, such as cognitions or body sensations [...] Yet, conversely, many valenced reactions are not usually classed as emotions. Tasting sweet substances merely produces a pleasant sensation that usually is not regarded as an emotion. (Frijda 2008, p. 71)

Neither the visual experience of something blue nor the bodily sensation of one’s heart-beat are felt as intrinsically positive or negative. A lack of intrinsic valence seems to apply in general to exteroceptive modalities such as visual, auditory, gustatory, olfactory or tactile experiences (Fulkerson 2019, for an important qualification see, see also section 6.4). Similar points apply to proprioception. In interoception, it becomes trickier to keep the non-affective and affective realm apart. This is likely because some of the brain areas that are important for interoception are — together with other brain areas — also the locus of valence (e.g. Craig 2009; Damasio and Carvalho 2013; Carruthers 2017b).

However, exteroceptive experiences and bodily sensations naturally prompt or co-occur with affective experiences such as pain, sadness, joy or fear which *do* feel positive or negative (see section 4.4).

Suppose you step into a bath that, being too hot, causes an unpleasant pain in your foot. This experience will be bad for you; and it will also motivate you to act, for example to lift your foot from the scalding water.

⁸⁵I use the terms “valence” and “affect” synonymously.

[...] A table’s looking square to you, we might say, is neither bad for you nor good for you, nor motivational (independently of further desires). If that’s right, what makes unpleasant pain different? [...] An obvious answer is: its unpleasantness. (Bain 2013, S69 sq.)

It is a remarkable datum about consciousness that dissociable non-affective and affective experiences are unified in one single total phenomenal state of the subject (Bayne 2004; Koksvik 2014). It is because of this unity of consciousness that it tends to be hard to tell seeing, hearing, imagining or judging from being afraid, angry or happy. Nevertheless, that these are distinct components blended together in a unified phenomenal state is indicated by the fact that these experiential components come apart (see the dissociation cases in section 4.2). This, in turn, implies that we might sometimes get the psychological units wrong and misattribute features of one state to another.

Some clarifications: First, when I talk of valence here I mean valence as a *phenomenal* property of affective experiences (see also Weiss 2016). Such phenomenal valence has been called elsewhere “affect valence” and needs to be distinguished from closely associated but ultimately non-phenomenal properties such as emotion- or object-valence (Charland 2005; Colombetti 2005). Emotion valence refers to whether an emotion is negative or positive as such, regardless of whether the emotion is felt or not. Against this background, fear or sadness are considered ‘negative’ emotions while pride and joy are considered ‘positive’ emotions.⁸⁶ Analogously, object valence refers to whether an object or stimulus is negative or positive as such — objectively as it were —, regardless of whether it elicits emotions or feelings. One might, for instance, consider angry and sad faces, snakes, guns, crimes as ‘negative’ objects and happy and attractive faces, tasty food and erotica as ‘positive’ objects.⁸⁷

However, their functional profiles and the behaviours they facilitate seem to be sufficiently similar to merit calling both valence.⁸⁸ One might, of course, still attempt to

⁸⁶One might want to argue that they are so considered *because* the related emotional *feelings* have positive or negative phenomenal valence. Though this is to already provide a theory of emotion valence in terms of affect valence or phenomenal valence, claiming that the former is e.g. derivative on the latter (see also footnote 84).

⁸⁷Again, one might want to argue that they are so considered *because* they trigger positive and negative feelings that are more fundamentally positive and negative in that they have a phenomenal valence and represent the objects in questions as positive or negative through their phenomenal valence (see section 4.4). Nevertheless, this again is to provide a theory of object valence as derivative on other kinds of valence.

⁸⁸Taking one’s lead from Carruthers, one can say something instructive on the functional difference between unconscious affective responses and affective experiences, i.e. on what the first-order consciousness of valence adds to the mix:

deny that there is such a thing as non-phenomenal valence. One could, for instance, draw attention to the important distinction between something being unconscious, conscious and reflexively conscious.⁸⁹ Then one could point out that the only way employed to experimentally test whether something is conscious or not is to simply ask the subjects (e.g. Winkielman, Berridge, and Wilbarger 2005). Taking this as a point of departure one could argue that what such experiments really test is not *experience* (i.e. consciousness) but the awareness of experience (i.e. reflexive consciousness) (for a similar point on interoceptive feelings vs. interoceptive awareness see de Vignemont 2018, pp. 261 sq.).

Consequently, what such experiments show is *not* that subjects lack valenced *experiences* but that they lack *awareness of* valenced experiences, perhaps due to their briefness or other factors. Naturally, the subject will be unaware of them – after all, the experiments quoted in support usually use *subliminal*, i.e. *very* brief, presentations of affect-eliciting stimuli (e.g. Winkielman, Berridge, and Wilbarger 2005; Pessiglione et al. 2007; Childress et al. 2008). Unsurprisingly, the subject will be unaware of the briefly occurring valenced feeling in response to an even briefer encounter with an invisible stimulus. There is no consciously perceived stimulus that could direct a portion of her attention inwards at the time of the valenced experience’s occurrence.⁹⁰ It seems unwarranted to assume that the

[A]ll affective states can issue in action via either one of two partly separate routes. They can enter into practical decision-making, interacting with one’s perceptions and beliefs to issue in decisions about whether and how to act. But they also give rise to behavioral dispositions *directly*, independently of anything resembling belief-desire practical reasoning. In fact, all affective states tend to activate related motor plans directly, in such a way that these will need to be inhibited if one is *not* to act. [...] consider the fear caused by the sight of a bear on the trail ahead. This might interact with one’s knowledge that black bears will normally retreat if one is noisy and looks large, leading one to shout and stretch oneself fully upright. But it will also give rise to an immediate and unplanned impulse to run away. This impulse (and the motor-system activation that accompanies it) will need to be *inhibited* if one is to do the sensible thing and stand one’s ground. (Carruthers 2018, p. 7)

Now, it is plausible that unconscious affective responses and affective experiences both influence behaviour via the *direct route*, interacting with specialised behavioural systems triggering habitual emotional behaviour. In order to influence behaviour via the more *indirect route* of practical decision-making, however, the affect will need to be simultaneously accessible to a large number of sub-systems. This is what consciousness brings to the table: In becoming conscious, the valence is made “accessible by virtue of being globally broadcast”, engendering practical decision-making and more informed behaviour. This is what Carruthers’ example illustrates.

⁸⁹For simplicity, I omit breaking down first-order consciousness into access and phenomenal consciousness. Though interesting points could be made there as well, I presume.

⁹⁰One might argue that affective experiences qua being affective automatically draw attention to themselves. Nonetheless, this does not seem to be the case. It is true that affective experiences determine patterns of salience among the contents of our representations of the world, body and mind. However, they make us usually attend to *these* contents, not the affective experiences themselves. This makes good practical sense since what needs to be done is usually something about the object of the affective experience and not the affective experience itself (cf. Bain 2013, 2019). It is in virtue

valenced experience lasts substantially longer than the inducing subliminal presentation itself, so as to be detected and reported.⁹¹ So by the time the subject's attention is pointed inwards by being questioned about her feelings, the valenced experience has long subsided.⁹² However, that a valenced experience has occurred is evidenced by the trace it left: characteristic behavioural dispositions that lead the subject to engage in the observed affective behaviours – characteristic products of affective experiences. Thus, the experimental observations can be explained more parsimoniously via the workings of good old phenomenal valence. There is no need to postulate the existence of some dubious non-phenomenal valence.

Now the proponent of non-phenomenal valence might counter: It is true that awareness is not an *infallible* guide to consciousness, but it is *a* – perhaps the only – guide to consciousness. A lack of awareness can thus be taken as nondemonstrative evidence for a lack of consciousness.⁹³ And, on second thought, is non-phenomenal valence really such an outlandish construct? Aren't we happy to grant that perception can be unconscious (cf. Prinz 2005)? Why not valence? After all, we could deny that the perception of subliminally presented stimuli is unconscious following the same logic applied in the denial of non-phenomenal valence.⁹⁴ I think we reached a venerable impasse here. Whether something is first-order phenomenally conscious or not is a notoriously vexed question. And I won't attempt to solve it here. I am happy to grant that non-phenomenal valence exists as long as its existence is not at odds with the existence of *phenomenal* valence which *is* necessarily conscious (cf. Prinz 2004b, pp. 176-178).⁹⁵

of doing something about these objects that we do something about the affective experiences. It is running away from the frightening bear or tending to one's aching wound that does away with one's fear or one's pain, not attending to (the negative valence of) one's fear or pain.

⁹¹After all, something that can be barely unconsciously perceived for a few milliseconds can't be that important for the survival of an organism in natural circumstances. And *these* are the circumstances that forged our affective apparatus.

⁹²It is also unlikely that, at the time of the occurrence of the valenced experience, the subject forms a memory of the occurrence that she could later consult when reporting.

⁹³One might question this, though, by pointing out that while it might well be that awareness is a good indicator *for* consciousness, it is not a good indicator for its absence. That is, one can easily grant that if someone is *aware of* a certain state, this is good evidence *for* the state being conscious. One can deny, however, that not being aware of a certain state is good evidence *against* the state being conscious.

⁹⁴One could further try to make the point that we need non-phenomenal valence in order to account for affective reactions in non-human animals (Berridge and Kringelbach 2008, pp. 459 sq.; Winkielman and Berridge 2004, p. 122). Why would one try that? Because the cortical regions that are by some believed to add consciousness to the more basic unconscious affective reactions are more developed in humans. However, this point can be countered by pointing to extensive evidence for consciousness *without a cerebral cortex* (Merker 2007).

⁹⁵Is phenomenal valence necessarily conscious only by definition? Only to the degree that perceptual experience is necessarily conscious only by definition, I submit.

Back to questions about phenomenal valence and affective experiences: note that the claim here is that all affective experiences intrinsically have *a* valence, not that they intrinsically have a *specific* valence. Surprise by itself, for instance, might be indeterminate: there are positive and negative surprises (but see Knight et al. 2013; Noordewier and Breugelmans 2013; Noordewier, Topolinski, et al. 2016). This is compatible with surprise being an affective experience as long as it always comes with positive *or* negative valence and not wholly *without* valence. Third, I want to leave open the possibility that some feelings such as surprise or nostalgia can feel (to different extents) positive *and* negative at the same time, i.e. they can be ambivalent.

[A]ll emotions are valent. Some emotions may be intrinsically negative (sadness, fear), some may be intrinsically positive (joy), and some may have variable valence markers (surprise). [...] It also turns out that some emotions are *intrinsically* both negatively and positively valent. Some emotions are intrinsically mixed. (Prinz 2004b, p. 164)

This implies that valence is not a bipolar negative/positive continuum but that negative and positive valence are two independent dimensions (see e.g. Cacioppo et al. 1999; but see Russell and Carroll 1999).⁹⁶ Note, what I grant here is that a *single* feeling can feel positive and negative. This is different from cases where we have mixed feelings, i.e. concurrent or sequences of individual feelings some of which are positive and some of which are negative (Larsen and McGraw 2014; Schneider and Schwarz 2017).

4.3.2. Other Hallmarks of Affect: Felt Arousal and Motivation

Another characteristic phenomenal aspect of affective experiences is felt arousal: during an affective experience, the subject feels a more or less localised increase or decrease (i.e. *change*) in the level of activation, energy or excitement (Barrett and Bliss-Moreau 2009; Proust 2015; Colombetti and Harrison 2018).

Such felt arousal co-varies with but is distinct from actual *non-phenomenal* physiological arousal states (Colombetti and Harrison 2018) that, among other things, consist “of a variety of autonomic changes in heart rate, blood pressure, activity in the sweat glands, and levels of adrenaline and other chemicals in the bloodstream, as well as behavioural changes in posture, muscle tension, breathing rate, and so on” (Carruthers 2011, p. 127)

⁹⁶These reflections should defuse some criticisms that have been advanced against the idea of valence (Solomon and Stone 2002; Solomon 2003a).

and which are functionally associated with the allocation of processing resources to task-relevant sites (Dawson et al. 1989; Filion et al. 1991). It is often thought that the degree of felt arousal determines the “intensity” of the affective experience while the valence determines its “sign”, i.e. whether a feeling is positive or negative. Note that there are several possible ways of understanding the relationship between arousal and valence. A natural way is to see valence and arousal as essentially united, being two sides of the same coin: the arousal determines (perhaps among other things) how intensely positive or negative an affective experience is, just like the amount by which one increases or decreases the sound volume determines how much louder or quieter the emitted sound gets. The net change in volume would be similar to a change in arousal and the direction of the change would be similar to valence. Call this idea valence/arousal unity:

Together, valence and arousal form a unified state, so although it is possible to focus on one property or the other, people cannot feel pleasant or unpleasant in a way that is isolated from their degree of arousal. (Barrett and Bliss-Moreau 2009, p. 171)

Another possibility is that valence and arousal are ultimately dissociable phenomenal properties (Anderson et al. 2003; Lewis et al. 2006; Kuhbandner and Zehetleitner 2011) like, say, shape and colour are dissociable in visual experience as is e.g. evidenced by the condition of achromatopsia. Call this idea valence/arousal disunity.⁹⁷

On this picture, it could turn out that valence has only two discrete values, positive or negative, or that it is gradable independently of arousal which is essentially continuous. That arousal independent valence has only two values appears little plausible since it would be unable to explain the observation that, say, some pains are worse than others. Valence/arousal unity can explain this datum since it implies that a discrete positive/negative notion of valence is complemented with an intensity by an integral arousal dimension. So a valence/arousal disunity view should assume that both, valence and, obviously, arousal are gradable. Consequently, the degree to which an affective experience is positive or/and negative is independent of the degree to which it is arousing. For instance, we could have mildly positive but highly arousing feelings. Note that valence/arousal disunity fits more naturally with intrinsically ambivalent feelings and a two-dimensional account of positive and negative valence.

⁹⁷This idea implies that valence and arousal are ultimately independent in the sense of being dissociable. It should be not taken to imply that they do not typically interact with each other. To use the comparison from the main text again: Although colour and shape are independent properties, colour nevertheless facilitates shape discrimination.

Independence between valence and arousal is also better compatible with the possibility that valence and arousal are not just distinct and dissociable phenomenal properties but that they are distinct *kinds* of phenomenal properties. In principle, each of these properties might be a somatic (being a subcategory of sensory properties), a non-somatic but sensory or a wholly non-sensory phenomenal property.⁹⁸ If valence and arousal are independent, then it would be natural to grant that e.g. arousal is understood as a somatic property while valence is, say, some kind of distinctive non-sensory affective property (Carruthers 2011, 2017b). At this point, however, I want to remain neutral on these questions.

What seems clear is that both, valence and arousal, are closely associated with another crucial feature of affective experiences: they are directly motivating and so move us to behave or act in more or less stereotypical ways (e.g. Zeelenberg et al. 2008; Scarantino 2014, 2017; Proust 2015). Jonathan Cohen and Matthew Fulkerson vividly illustrate this feature:

Consider poor Lucy, a machinist who’s having a bit of bad luck. While working at the bench, she is momentarily distracted and strikes her thumb hard with a ball-peen hammer. This, as you might expect, results in a great deal of pain. [...] [F]eeling it typically moves us to act, and often in predictable ways. Thus, for example, we can imagine Lucy acting in various ways to this unfortunate incident. Screaming would be understandable, though not the sort of thing stoic Lucy would do. Still, even she would wince heavily and pull her arm away, immediately taking ameliorative actions on her throbbing thumb—holding it, squeezing it, shaking it. And it seems plausible that Lucy’s pain motivates these actions. Without the pain she wouldn’t have winced or grabbed her thumb, and her performing these actions is a direct consequence of her experience of pain. [...] There is nothing special about our choice of example here. Pain [...] is merely a paradigm instance of what we shall call affect. Affective experiences have some quality or character that drives or motivates us to act in various ways. They often seem to accomplish this by being in some sense bad or unpleasant (in the negative cases) or

⁹⁸Note that hedonic theories of valence might understand hedonic qualities of pain and pleasure to be somatic properties, i.e. some kind of bodily sensations. If this is correct and if arousal is somatic as well, then there would be *no* remainder if we subtract all the bodily sensations from a bodily feeling and an emotional feelings, for that matter. Thus, affective experiences would appear to be exhausted by bodily sensations of different sorts. What would distinguish “cold” bodily sensations from bodily and emotional feelings would be some special somatic component (e.g. valence) or components (e.g. valence and arousal) that the latter possess and the former lack.

pleasant and enjoyable (in the positive cases). (Cohen and Fulkerson 2014, p. 104)

Their motivational force plausibly co-varies with the valence (and its intensity) and the degree of felt arousal (which, remember, is a proxy for the rallying of resources to implement behaviour) of the affective experience (but see Corns 2014). In the case of valence, Carruthers goes so far as to argue “that valence might provide the motivational component underlying all intentional action, either directly or indirectly.” (Carruthers 2017b, p. 3)

A specific set of behaviours worth mentioning is that affective experiences tend to establish or modify patterns of significance or salience among the things in the world by motivating attentional shifts and investments (de Sousa 1987; Brady 2013). In other words, affective experiences point us towards what matters or is relevant by exerting an effect on attention (Sander et al. 2003; Vuilleumier 2005; Pessoa 2008; Pessoa and Adolphs 2010). Here is how Michael Brady illustrates and explains this fact:

[E]motions have two subsidiary functions: (a) they alert us to the presence of significant objects or events in our environment, and thereafter facilitate cognitive processing of such objects and events; and (b) they enable us to act appropriately with respect to these objects or events. [...] Emotions alert us to significant matters by capturing and directing our attention on to important objects and events. [...] The need for capture and direction of attentional focus stems from the fact that human beings are presented with vast amounts of information about the state of the world and the state of themselves, only some of which will be important to them. Given that human beings have limited mental resources, they thus face a problem of efficiently locating or identifying which information in their environment is important. Our emotional systems, at least in part, are thought to have evolved in order to solve this problem, and they do so by capturing and focusing the subject’s attention. In other words, important or significant events in our environment need “preferential perceptual processing”, and “[o]ne means of achieving this is by emotion enhancing attention, leading to increased detection of emotional events.” (Brady 2009, p. 422)

To wrap up on phenomenology: affective experiences are valenced, arousing and motivational. Furthermore, they are gradable along these (and probably other) dimensions:

they can be more or less arousing as well as more or less positive or negative and motivating.

4.4. Affective Intentionality

Although not wholly uncontroversial, there is fairly large agreement that affective experiences are intentional states (see e.g. Tye 1995, 2008; Goldie 2000, 2002; Crane 2009; Kriegel 2014; Mendelovici 2014; Weiss 2016).⁹⁹ There is, of course, more disagreement about how exactly to understand the intentionality of affective experiences. Is it, for instance, that *all* aspects of affective experiences are intentional (in one way or another), including but perhaps not restricted to their phenomenal qualities? Or are there some non-intentional bits, some kinds of “raw feels” or quale in the mix of components that constitute affective experiences? Do all kinds of affective experiences exhibit intentionality or are there some that don’t? Or instead, are different kinds of affective experiences characterised by distinct kinds of intentionality? These and many other are substantial and important questions. I will set them aside for now, however, and will shortly proceed with some observations that motivate the assumption that affective experiences are intentional.

But first, a terminological note: when it comes to the properties that feelings (phenomenally) represent, I will use response-dependent terms such as “unpleasant”, “painful”, “frightening”, “disgusting”, “pleasant”, “joyous” etc. in order to stay as noncommittal as possible concerning the nature of the properties that the feelings in question represent.¹⁰⁰ I could (and perhaps should) even go more noncommittal and use somehow “good/bad” or “positive/negative”.¹⁰¹ In some cases, such as bodily pain, fear or disgust perhaps relatively uncontroversial response-independent properties such as “tissue-damage” or “harmful”, “danger” and “contamination” might be available. All these response-independent construals are controversial, however. This just points to

⁹⁹Sometimes this assumption is defended against the background of a larger project, such as arguing for general representationalism or intentionalism about phenomenal consciousness (e.g. Tye 1995, 2008; Crane 2009).

¹⁰⁰For standard Jamesians about emotions, affective experiences as bodily sensations would (at least in part) represent bodily events (Damasio 1994; Tye 1995). At the same time, more sophisticated feeling theorists understand at least some affective experiences — emotional feelings — as representing non-somatic properties *via* registering bodily events (Prinz 2004b). These non-somatic properties are what I refer to with the response-dependent terms.

¹⁰¹Those who feel more at ease with other noncommittal terms, are invited to mentally substitute them for the terms I use. As far as I can tell, it shouldn’t affect the general line of thought.

the fact that it is often not clear at all what response-independent properties are involved when it comes to feelings. This seems to undermine the idea that these are the properties that are represented, making it questionable that these are really the properties in question that after all, everybody should agree on by consulting their affective experiences (Dokic and Lemaire 2013). Nevertheless, I think everyone can agree that there is something response-dependent terms capture about what affective experiences represent, and be it only something trivial.¹⁰² In any case, this usage is not to commit to a response-dependent understanding of these properties but to leave questions about them open.

4.4.1. The Particular Objects of Affective Experiences

After having clarified this, here come the promised observations: imagine that Jake accidentally cuts his left index finger while cutting vegetables. In this event, assuming Jake is not pain asymbolic etc., This is different from, say, when Jake just focuses his attention on his finger during mindfulness meditation. He might be directly aware of his finger and certain sensations and events, but neither his finger nor the events that are seemingly taking place in it, are represented in some specific way. Imagine now, that as a result of Jake's meditation practice, he feels his tensed neck relax. The events in his neck are represented as pleasant (in some determinate way). Note that the bodily pain and relaxation Jake feels is precisely localized.

Nevertheless, this does not have to be the case. Jake might feel tired, relaxed or energised after his meditation, without these bodily feelings being confined to any single body part. Instead, Jake's body and, consequently, Jake as a whole — with his whole being as it were — feels unpleasantly tired or pleasantly relaxed or energized. In fact, it does not seem out of the question that Jake's global bodily feelings come to overflow his body, colouring his experience of the world and perhaps even making him lose the experienced connection between his body and his altered experience of the world (cf. Ratcliffe 2005). This is not very surprising since neither the bodily sensations nor the pleasantness or

¹⁰²Response-dependent properties seem to be not very popular as an account of the feeling-specific properties because they are usually couched in terms of dispositional properties (Prinz 2006a, p. 149; Dokic and Lemaire 2013; Mendelovici 2014, p. 143). I think more promising accounts are available, however. These accounts would substitute or complement the dispositional analysis of response-dependent properties with Edenic phenomenal properties that are tailor-made for the phenomenal properties of specific feelings (Mendelovici 2014). Another proposal worth closer consideration is Cowan's account of 'phenomenally present as absent' properties (Cowan 2015). Having said that, to discuss these now would take me too far afield.

unpleasantness are *really* in Jake’s finger, neck or body to begin with. Bodily events are one thing, the registering of those in the form of bodily sensations and feelings another (Sizer 2006). Both can come apart. As Weiss notes:

We have known since Descartes (or at least since experimental psychology really got going in the nineteenth century) that a “pain in one’s hand” is not really in one’s hand. There can be tissue damage in one’s hand, for example, without pain and there can be “pain in one’s hand” without one’s even having a hand (as we see e.g. in “phantom limb” cases.) Pain (or pleasure) in one’s hand should, rather, be understood as pain in one’s mind/brain directed towards (putative) goings-on in one’s hand. (Weiss 2016, p. 39)

Bodily feelings represent bodily parts or events as, using generic and clearly determinable terms here, pleasant or unpleasant. And, perhaps, sometimes they can do so without being experientially bodily localised at all (cf. Mendelovici 2014). In fact, if they cease to be *experienced* as being *in* or *about* the body, one might start to wonder whether they haven’t ceased to be *bodily* feelings. Surely, they would continue to be brought about by bodily events and to be feelings, but they would no longer be experienced as representations of *bodily events* as pleasant or unpleasant. In other words, they would cease to have a *bodily phenomenology* while retaining a *bodily aetiology* and an *affective phenomenology*.

In a similar vein, emotional feelings often “transcend” the body and represent all kinds of things in characteristic ways, be they localised in the body, the external world or solely “in the mind”. Linda’s fear of an approaching bear represents the bear as frightening (or dangerous) while her joy upon winning the *Iron Man* represents the victory as joyous (or, perhaps, as an achievement).¹⁰³ Linda might also be sad that *her memory is degrading*, or relieved that *the US-Democrats got a majority in the House of Representatives*. Equally, she might be happy that her surgical wound is healing promptly, regret that she didn’t learn the drums as a child or hopeful that she (and not Smith!) will get the job. At other times, Linda is amused by the concept of a Flying Spaghetti Monster or feels infatuated with Sherlock Holmes. Linda might also be pleased by her quickness of the mind, frustrated with her anxiety about speaking in public or proud of the way her body looks. Emotional feelings seem to be very versatile in what they

¹⁰³It is interesting that seemingly response-independent properties that stand in some intimate connection to negative feelings such as pain (harmfulness), fear (dangerousness), disgust (contamination) and sadness (loss) come easier to mind than for positive feelings. What are the response-independent properties connected to bodily pleasure, joy, relief and pride, for instance? For the latter two it’s not even clear to me what the response-*dependent* properties are if one sticks to common language.

are about, i.e. in what they represent and how they represent it as being. They can be about objects, (past, present and future) states of affairs, other mental states and about propositions.¹⁰⁴

In the case of propositions, one can distinguish between feelings with propositions as *contents* and those with propositions as *objects* (Grzankowski 2016). Contrast the case of fearing that *Santa Claus does not exist* and fearing the proposition *Santa Claus does not exist*. Only the former is a propositional attitude while the latter is a non-propositional attitude, and only the former has truth- or accuracy-conditions.¹⁰⁵ Usually, when we talk of propositional attitudes we refer to the propositions in question in terms of their contents. A belief represents the contents of a proposition as true while a desire represents the contents of a proposition as to be realized. If not stated otherwise, when I will talk of propositions as intentional objects, I will mean it in the sense of propositions as *contents*, not as objects.

So emotional feelings can be about the body, but also about the external world, “the realm of ideas” as well as about the past, present and future. At extremes, they are about non-existent or abstract entities such as Swampsmans, Orfolei, Flying Spaghetti Monsters or propositions. Based on these observations it is easy to see that affective experiences typically have intentional objects, be they body part, entities in the external world, abstract and fictional objects, (past, present and future) states of affairs, mental states, propositional contents or propositions. Jake’s finger is the intentional object of his

¹⁰⁴There seems also to be no reason to exclude activities and actions from what feelings can be about. One can enjoy focusing on one’s breathing or be angry about someone’s offensive action. The same goes for properties and relations. One might feel comfortable wearing black and happy about one’s blooming friendship. Depending on one’s metaphysics, one can presumably make sense of some of these things in terms of objects, states of affairs or propositions. It doesn’t really matter here.

¹⁰⁵Attitudes with propositions as objects appear odd but not impossible. In fact, in some constellations, they might be not odd. Consider Grzankowski’s tongue-in-cheek remark:

If one had a certain view about propositions, one might come to think that they could cause one serious harm (for example). Similarly for other attitudes. It is odd to love a proposition, but not impossible. In fact, maybe it isn’t even that odd. People often claim to have a favorite number. Is it so outrageous that someone might also love a number? If a number, why not a proposition? (Grzankowski 2016, p. 325)

On a more serious note: If one takes propositions to be just sentence-like mental representations (e.g. Thagard 2008, p. 172), attitudes that take these as objects will be metacognitive in some sense or even metarepresentational but not necessarily odd. Perhaps, in fact, such attitudes are ubiquitous. Also: there might be *propositional* attitudes whose propositional contents include propositions as objects. Consider suspecting that *the proposition [Santa Claus does not exist] is ambiguous* or fearing that *proposition [A] entails proposition [B]*. This, all of a sudden, does not sound that odd at all but rather like kinds of states philosophers are in all the time. Perhaps philosophers are odd and so it is unsurprising that they have odd states. Or perhaps I misunderstand the objects involved in the above propositions as propositions—perhaps they are states of affairs and all is well.

pain, Jake's neck is the intentional object of his pleasure. The nature of the intentional object (at least, if they are consciously represented) can be thought to make the feelings in question *bodily* pain and *bodily* pleasure, i.e. *bodily* feelings.¹⁰⁶

The bear is the intentional object of Linda's fear and the propositional content (or state of affairs) that *the US-Democrats got a majority in the House of Representatives* is the intentional object of her relief. Though such intentional objects, often called their *particular objects*, are not the only things that feelings seem to be about. Feelings represent their particular objects as being *a certain way*, as bearing some feeling-specific property such as being pleasant, unpleasant, frightening or joyous.

4.4.2. The Formal Objects of Affective Experiences

These feeling-specific properties are variously called the formal objects (Kenny 1963; de Sousa 1987; Teroni 2007), core relational themes (Lazarus 1991; Prinz 2004b) or concerns (Prinz 2006a) of the feelings.¹⁰⁷ In short: A feeling represents its intentional object, i.e. its particular object, as having a property characteristic to the feeling, i.e. its formal object. My terminological note on the noncommittal usage of response-dependent terms was addressing these properties. That such a note is necessary foreshadows that one needs to tread carefully here. There is widespread disagreement on how to understand the nature of these properties and how they figure into different aspects of affective experiences such as their metaphysics, intentionality, phenomenology and epistemology. I think one can say that much: it is obvious that these properties are of central importance to an understanding of affective experiences. It seems that there is some constitutive relationship between the property and the feeling with its different dimensions.

That's where the common ground ends. Let me flag some issues. I was already courting controversy by using representational language in saying that a feeling *represents* its formal object. Using such terminology suggests that the feeling-specific property is represented on the level of the feeling's content. Most agree with this assessment (e.g.

¹⁰⁶Are there other pains and pleasures than bodily? On the face of it, there are: social pain (Eisenberger 2012) and cognitive pleasure (Goldstein 1981), for instance (cf. Helm 2001).

¹⁰⁷As a concept, formal objects are the most general ones and can also be applied to other kinds of mental states than affective experiences such as beliefs or perceptions (Kenny 1963; Deonna and Teroni 2015). Among other things, formal objects are often taken to specify the correctness conditions of the respective states. Fear of a dog is correct if the dog is *actually fearsome*, a perception of a dog is correct if a dog is actually there, i.e. if what the perception represents *actually obtains*, believing that *P* is correct if *P* is *actually true*. Core relational themes and concerns, on the other hand, are terms that are typically reserved for the formal objects of affective experiences.

Goldie 2000; Mendelovici 2014; Mitchell 2019).¹⁰⁸ If one assumes that formal objects are to be found on the level of content, then we can understand formal objects as similar to *aspectual shapes*, which *together* with the intentional object, make up the *contents* of mental states.¹⁰⁹ That is, feelings have not only an intentional object (particular object) but also a feeling-specific way in which they represent this object, i.e. a feeling-specific aspectual shape—a formal object. Here is how Crane illustrates the idea:

The basic idea of aspectual shape is very simple: in any intentional state, the objects on which the mind is directed are presented in a certain way. Suppose that you are thinking of St Petersburg—with its elegant baroque buildings and its harsh climate. You are thinking about it in a particular way: maybe you are visualizing it in the imagination, on the basis of pictures you have seen or on the basis of experience. Or maybe you are just thinking about it as *St Petersburg*—that is, a thought which you would express by using the name ‘St Petersburg’. You may just think to yourself, ‘Vladimir is in St Petersburg; I wonder what the weather is like there?’ When you think about St Petersburg as St Petersburg, the aspectual shape of your thought is different from when you think about St Petersburg as Leningrad, or when you think of it while listening to Shostakovich’s *Leningrad Symphony*. Similarly, when you visually perceive St Petersburg, you see it from some particular place, in certain particular conditions of illumination, and so on. You see it under a certain aspect. Your experience, like your thought, has a certain aspectual shape. These truisms are just ways of expressing the simple idea that one cannot think of something without thinking of it in some way.

However, when it comes to formal objects, there is disagreement on the nature of the content of feelings, i.e. on whether it is conceptual, nonconceptual or both and whether it is propositional, non-propositional or both. There is equal disagreement on how feelings represent, i.e. phenomenally, functionally or both. Then there are those who insist that formal objects (together with other properties of the feeling¹¹⁰) are not part of

¹⁰⁸As far as I can tell all (quasi-)judgmentalists and (quasi-)Perceptualists about emotions at least agree with this when it comes to emotional feelings.

¹⁰⁹While “formal object” is a term that has settled in the philosophy of emotions, aspectual shape is the more common term in other areas of philosophy of mind, expressing essentially the same idea. The notion of a concern-based or evaluative “construal” as used by Robert Roberts resembles an emotion-specific version of aspectual shape (Roberts 2003, 2009).

¹¹⁰Note that one can ask the same questions as one can ask for formal objects about their other intentional aspects such as their particular object and the phenomenal properties of valence and arousal introduced in the former section as well (for the case of valence see especially Weiss 2016, see also

the representational content of feelings but that they somehow reside in the attitude component of the feeling (Crane 2009; Deonna and Teroni 2012, 2015; Weiss 2016). Since I want to remain neutral between the content and attitude view for now, please read (pro- and retrospectively) in sentences involving parts like “a feeling represents its formal object” as “a feeling represents *or what the attitude-view theorist says the feeling does with* its formal object”. In any case, here is how Julien Deonna and Fabrice Teroni motivate their move towards the attitude-view:

If it makes sense to say that what frightens Julianne is what John is amused by [...] then we have reason enough to think that the difference between their two emotions is not to be located at the level of their respective contents. If this difference were located at the level of the content, that would imply that Julianne is frightened by (a specific instance of) dangerousness, whereas John is amused by (a specific instance of) funniness; their respective emotions would then be about different things. The contrast here is analogous to that between attitude and content in the case of other mental states. Consider belief. If it makes sense to say that what John believes is what Julianne doubts, then the reason why truth enters only into [...] John’s mental state has to do with the fact that only he takes the specific attitude of belief towards it. (Deonna and Teroni 2012, p. 77)

There are ways to resist this move. Assume that it is a dog Julianne is frightened and John amused by. In this argument, it seems implicit that intentional objects are all there is to representational contents. If we, however, employ the concept of aspectual shape, it becomes clear that the same dog can very well be the *intentional object* of Julianne’s fear and John’s amusement while the emotional aspectual shape under which the same dog is presented differs. In one case the dog appears as frightening while in the other the same dog appears as amusing. So it *is* the case that the difference between the two emotions can be located at the level of their respective contents, but not in the part that makes up the intentional object but in the part that makes up the aspectual shape.¹¹¹ In

the clarifications on valence below).

¹¹¹This will seem all the more plausible when we encounter the concept of *base* in section 4.4.4. The base of a feeling is a mental state (or a set of states) that supplies a feeling with its particular object. In our example it is not Julianne’s fear or John’s amusement themselves that establish a representation of the dog in the first place, rather it is, say, their visual experience of the dog. However, this visual experience delivers the dog with its *own* aspectual shape. And it is this visual experience and its intentional object under a specific visual aspectual shape which might be the *same* for Julianne and John. The difference lies in their respective affective experiences that differ in the way they *affectively* represent the *equally visually* represented dog.

other words, they are not about different things when it comes to the intentional object but about different things when it comes to how they represent this object as being. Of course, one can counter by saying that aspectual shapes are parts of the attitude, not the content.¹¹² But at least on first sight, this seems to lack independent motivation (cf. Bourget 2017, pp. 11 sq.). Moves like this prompt Goldie to remark that “it always seems possible for an opponent to force all the difference into the attitude, so that the debate degenerates into a matter of competing intuitions” (Goldie 2000, p. 60; for a systematic elaboration of Goldie’s point see Mitchell 2019).

Don’t get me wrong: the attitude theorist certainly has a point in that formal objects *covary* with the type of feeling (i.e. attitude). That is, if one changes the formal object, one changes the feeling. If the formal object is fearsomeness, the feeling is fear. If we change the formal object of this feeling to, say, funniness, then the feeling effectively ceases to be fear and turns into amusement. That is, formal objects *are* attitude-specific. It remains, however, unclear why feelings (or any mental state, really) cannot have contents or aspectual shapes that are attitude-specific.

Apart from questions regarding where to best locate the formal object in a feeling, it is also not clear what sort of theoretical work formal objects are supposed to do for feelings: is it metaphysical in that they individuate feelings? Is it phenomenological in that they are part of the phenomenal character of feelings? Is it intentional by being part of their representational content? Is it epistemic in that they determine the correctness conditions of feelings or make them somehow intelligible? All of them? Some of them? None? And as already mentioned above and intimately related to all these issues, it is not clear what sort of properties these feelings-specific properties are.

Be it as it may, one can, uncontroversially I think, say that the intentionality of feelings has two parts: a particular object and a formal object where the feeling represents the former as bearing the latter.

¹¹²In fact, the term “intentional mode” seems to be sometimes used synonymously with content-wise aspectual shape and sometimes with attitude (cf. Weiss 2016, p. 49). One might distinguish between: 1) representing a dog: that’s the intentional object (content) of the fear (provided by its base, more on that below), 2) representing a dog as frightening: that’s the intentional object and the characteristic aspectual shape (content) of the fear and 3) frighteningly representing a dog: that’s taking the fearful *attitude* towards the intentional object of the fear (content). It is not clear how to adjudicate between 2) and 3).

4.4.3. Interjecting Remarks on Valence and Intentionality

I think it is a good time to clarify a point about phenomenal valence which will reinforce the assumption of affective experiences being intentional. Recall: Phenomenal valence describes the felt positivity or negativity of affective experiences. I noted that this basic positivity or negativity is usually made sense of in hedonic terms as pleasantness or unpleasantness or in value terms as a representation of goodness (value) or badness (disvalue). This leaves a free parameter: the hedonic view of valence seems to suggest that phenomenal valence is a property of the experience (or attitude) itself while the representational view of valence seems to suggest that phenomenal valence is a property of what the experience is about, i.e. it is encoded in its representational content.

I do not necessarily have to take a stance here between the two views or better still: I can be pluralist and assume that they are complementary rather than mutually exclusive. Sometimes it is the experience itself that feels positive or¹¹³ negative, sometimes it is what the experience represents that feels positive or negative and sometimes, perhaps, both are possible at the same time: the experience itself feels positive or negative *and* what the experience represents feels positive or negative.

Having noted that, my sympathies lie with the representational view, however. The hedonic view seems to flirt with the notoriously mysterious idea of non-intentional qualia or raw feels. In fact, I think that the representational view can accommodate the hedonic view: An experience can feel positive or negative *in virtue* of representing something as positive or negative. On the other hand, it seems less natural that something is represented as positive or negative because one has an experience that feels positive or negative. It appears less phenomenally sound to posit such a kind of seemingly “explicit egocentricity” for *all* affective experiences (see also 6.4). As Carruthers explains:

One initial strike against the hedonic account is that [it] accords less well with our affective phenomenology than does the representational one. When a bear looms out of the bushes while one is hiking (causing fear) it is the threatening aspect of the bear (its size, its claws) that seems bad. All of one’s focus when afraid is generally outward-directed, targeted on the object of one’s fear. And that is what (according to the representational account) seems bad as a result. [...] In contrast, the hedonic account gives a much less natural, more self-focused, treatment of these cases. The sight of the

¹¹³The “or” here is to be read as an inclusive one, accounting for the possibility of ambivalent feelings.

bear causes an intrinsically-bad quality to become a component of one's experience, and the presence of the bear is only taken to be bad because one believes that it is causing this quality (Carruthers 2017b, pp. 7 sq.).

Fear of a bear feels bad *because* it represents the bear as bad. The bear is not represented as bad *because* the fear feels bad (for more arguments against the hedonic view see *ibid.*). The representational view also leaves open that *sometimes* affective experiences *themselves* can feel positive or negative. If affective experiences can be *about* mental states, affective experiences among them, why can't they sometimes be (partly) about a specific affective experience: themselves? If that seems odd and overly theoretical, here is a more natural alternative: we can sometimes have an affective experience and an affective experience about this affective experience that is similar to the one it is about. Just think back to Linda's frustration about her anxiety to speak in public (it could have been also anxiety about her anxiety). So the representational view can accommodate instances where affective experiences themselves feel positive or negative: this is because they are represented by (another) affective experience as positive or negative.

In the end, I think it is natural to assume that if valence is a felt positivity or negativity, *something* needs to be felt as positive or negative: "No mental states are valenced in their own right, so to speak: they must be valenced attitudes or orientations or stances towards or about something" (Weiss 2016, p. 34). In the end, phenomenal valence is an *evaluation of* something that is represented. The bear is evaluated as bad (in the determinate form of being fearsome) within the content of one's experience. Valence is plausibly the phenomenal ground for this evaluation, so that without it, there would be no *phenomenal* representation of fearsomeness, joyfulness or any other feeling-specific property.¹¹⁴

This does, however, not mean that we cannot be sometimes *unaware* of what this something is which is positively or negatively evaluated: "we might have utterly no idea, sometimes, what our feelings are about, even though they are always conscious and are always about something" (*ibid.*, p. 40). In such cases, the felt positivity or negativity might appear to be "on its own"—suggesting that it is the experience itself that is positive or negative. However, that this cannot be the full story is suggested by the

¹¹⁴This is not to say that we cannot *somehow* represent the feeling-specific properties or their response-independent analogues when we e.g. judge something to be, say, fearsome and do so without being afraid. However, this would be best construed as a *non-phenomenal* representation of the property in question. There seems to be no property-specific phenomenal contrast between judging something to be fearsome (without being affectively afraid) and judging something to be amusing (without being affectively amused).

observation that we do not usually rest content with having *just* a positive or negative experience. Typically, when we are unaware of a feeling's particular object, we go looking and try to figure it out.

4.4.4. The Bases of Affective Experiences

The point about valence just made will become clearer when we now talk about a subject of special importance for the particular object. A characteristic feature of affective experiences is that they rely for parts of their intentionality on other mental states such as bodily sensations, perceptions, judgments, memories, imaginings etc. These mental states provide affective experiences with their particular or intentional object (Herzberg 2012).

[Emotions] differ from perceptions in that they cannot be seen as independent ways of accessing the objects that exemplify these properties. For instance, while the injustice of Jonas's remark is perceived by Mary through her indignation, the remark itself is not. Mary must access it by some other means (perception, memory, belief, etc.) [...] she hears the remark and feels it is unjust. The evaluative apprehension [...] is grounded in such non-evaluative bases [...] there is no such comparable distinction between two psychological levels exemplifying causal and epistemic relations within the field of perception proper. (Deonna and Teroni 2012, p. 69)

Where perceptions get their inputs from organs and transducers, emotions get theirs from their cognitive bases. (Oliver-Skuse 2016, p. 28)

Such mental states are usefully called the *base* of a feeling (Mulligan 1998; Deonna and Teroni 2012, 2015; Bain 2013; Oliver-Skuse 2016). Affective experiences are flexible in that they can take different kinds of states (or sets of those) with different kinds of contents (e.g. propositional/non-propositional, conceptual/nonconceptual, iconic/non-iconic, conscious/unconscious) as their bases.¹¹⁵

This explains how affective experiences can be so rich in the kinds of intentional objects they admit.

¹¹⁵Relatedly, Prinz notes that, for instance, fear “is triggered when the auditory system detects a loud sudden noise, or when the visual system detects a looming object, or when we proprioceptively detect a sudden loss of support” (Prinz 2004a, p. 55).

As a consequence, there might be affective experiences that are not clearly related to the contents of other attendant *conscious* states because their base is likely (as yet) unconscious (Weiss 2016, pp. 37 sqq.). This is presumably what explains the observed double dissociations in section 4.2 (see also “blindfright” cases Scarantino 2010, pp. 734 sqq.). Ultimately, however, there really are no affective experiences *without* other mental states. Well, of course not: There are no feelings without *at least* some causes that might serve as their base. However, as illustrated by the double dissociation, this does not entail that affective experiences are reducible to these other mental states, i.e. the bases that supply their particular object.

Now, if Rita is afraid of an approaching bear, her fear relies on the multi-modal perceptual experiences of the bear in order to represent *the bear* as frightening. The feeling will so to say “encode”, “localize” or “embed” the badness of the bear in the determinate form of fearsomeness within the content of Rita’s multi-modal experience of the bear. Note that the base will not only provide the feeling with *just* a particular object but that the base itself will represent its intentional object under a determinate aspectual shape (this is what I mentioned in footnote 111). So Rita’s visual experience will not only represent an approaching bear *simpliciter* but it will represent the bear as approaching quickly, approaching over a specific path in the forest, being massive, having sharp, flashing teeth, having brown fur etc. The feeling, in turn, will modify or complement this content of the base with its feeling-specific property.

Depending on the kind of the base and on the content of the base, Rita’s feeling will appear justified or unjustified (Deonna and Teroni 2012; Echeverri 2017). That Rita is afraid of an approaching bear is in principle neutral between the approaching bear being the dreadful beast just described or a little bear cub in the zoo. While being afraid of the former seems like an eminently justified reaction, being afraid of a cute bear cub might be less so.

In fact, it appears plausible that the base and its content do not only help decide whether a feeling is justified or not but that it helps decide whether such assessments make sense in the first place. Is Jake justified in feeling pain upon cutting his finger? Is he justified to feel pleasure upon the relaxation of his neck? Is it justified to feel pain in one’s phantom limb? I am not sure whether these questions have yes/no answers. It seems not implausible that in the case of bodily feelings, i.e. where bodily sensations are the bases, the feeling will appear as not properly assessable in terms of justification, i.e. they are beyond justification (for considerations to the contrary see e.g. Siegel 2013; Martínez

2015). Before a base of a feeling can lend justification to a feeling, it must endow the feeling with the ability to be justified in the first place. And not all bases seem to do that.¹¹⁶

Note that the base does not have to be itself subject to justification assessments in order to make the feeling that it serves as a base assessable. For instance, many consider visual experiences to be not assessable in terms of justification. At the same time, the beliefs they give rise to are considered to be so assessable. It is fairly common practice to apply similar considerations to emotional feelings: Rita’s fear of a bear *can* be justified. And whether it is justified depends on the contents of her perceptual experiences that are beyond justification.¹¹⁷

Note further that while some bases and feelings might not appear to be assessable in terms of justification, that does not mean that they do not have correctness-, accuracy- or truth-conditions (or other conditions of satisfaction) (cf. Bain 2013). Pain in one’s phantom limb is a problem not only because it is pain but because it is *illusory* pain. There is no “real” bodily basis for it: there is a representation of bodily events (as painful) where there are no bodily events. Similarly, “pain” in someone with pain asymbolia is a problem not because it fails to represent bodily events — in fact, it doesn’t — but because it fails to *correctly* represent these bodily events *as painful*. Thus, feelings are evaluable in terms of their correctness although they might not always be evaluable in terms of justification. This they appear to have in common with perceptual experiences about which most agree that they are never subject to justification (but see Siegel 2017).

Another significant relationship between the base and the feeling is what I call “base property mirroring”. The idea behind it is that an affective experience will come to mirror certain dynamical, intentional and phenomenal properties of its base. Take Rita’s fear of the bear. Her visual experience of the bear will, given good eyesight, be crystal clear. On the other hand, if Rita needs glasses to correct her vision and is not wearing them, her visual experience will be not so crystal clear. Consequently, the aspectual shape of the content will be different, one determinate and one more blurry. The ensuing fear based on these experiences will perhaps vary in some aspects as well. In one case it

¹¹⁶My conjecture is that whether a base is justification-endowing or not has to do with whether the content of the base is in principle publicly accessible or not. If the bear is not a hallucination, then everybody in Rita’s entourage can see for themselves whether being afraid is a good idea. On the other hand, it is difficult to access the contents of someone else’s bodily sensations. This implies that whether some kind of base is justification-endowing or not is not set in stone but subject to change. If, for instance, the contents of the bodily sensations of others would be made (reliably) accessible by some technological innovation, then bodily sensations might as well become justification-endowing.

¹¹⁷Alternatively it might depend on states that *are* rationally assessable such as beliefs or judgments.

will be perhaps fear of a massive approaching bear with sharp teeth while in the other case it will be fear of a massive approaching bear (but she can't see the sharp teeth) or even fear of a massive approaching figure, without identifying it as a bear.¹¹⁸ In any case, it seems that the content of perceptual experiences is gradable and this will have a phenomenal impact on feelings based on them. Contrast this with Linda's relief upon reading in a newspaper that *the US-Democrats got a majority in the House of Representatives*. Upon entertaining the thought that *the US-Democrats got a majority in the House of Representatives* she is relieved. This thought does not seem to be content-gradable in anything like the way as Rita's visual experience of the bear or Linda's visual experience of the graphemes in the newspaper, for that matter.

Of course, there are also deep open questions when it comes to the base. For now, I want to flag the most significant: What is the relationship between base and feeling? Of that, there are, of course, different conceptions. And these conceptions might plausibly vary with the kind of feeling. In the case of pain, the connection seems very tight—perhaps constitutively blended, at least phenomenally and in normal conditions (see the double dissociation). In general, one can adopt a causal understanding: the base is what causes the feeling (e.g. James 1884; Goldstein 2002). One can adopt a non-causal interactive Componentialism: the base and the feeling co-occur (and appear related or bound together in experience) (Herzberg 2012). One can adopt Blenderism: the feeling and the base are constitutively unified or fused in one blended state (Goldie 2002; Fulkerson 2019; Mitchell 2019; for the same point about intuitions see Chudnoff 2011b and section 3.4.2.¹¹⁹

Note that the different views are not necessarily mutually exclusive but can be complementary: For instance, it seems plausible that sometimes the feeling is caused by something to which it will also appear related or bound. One can distinguish here at least between *causal base* and *phenomenal base*. In the mentioned case, the causal and phenomenal base coincide. However, on other occasions, where the feeling is caused by

¹¹⁸Perhaps the approaching bear even fails to elicit fear in poor Rita due to her poor eyesight.

¹¹⁹Depending on how one fixes the relation between feelings and bases, complications for the common practice to take the justification-conferring relation between base and feeling to be similar to the one between base and beliefs (cf. Echeverri 2017). If the base and the affective experience is in fact inextricably tied together in a Goldie-style blend, then the relationship between feelings and the mental states that are their bases is of a markedly different sort than the relationship between beliefs and the mental states that are their bases. It might well be that in such a case the connection is so tight that either the feeling must lie outside of the realm of justification or the base must be inside of it. Alternatively, one could posit qualitatively different rational status for different aspects of the contents of the (blended) emotional state (cf. Mitchell 2019).

something but appears related, or is bound to something else, the causal and phenomenal base can come apart. Just consider the last time you were “hangry” or when you fell victim to someone else’s “hangriness”, i.e. someone got angry at *you* (because you were a readily identifiable particular object) due to his or her hunger (the cause for the irritability). Now, what’s the base here? The exteroceptive experience of you? Is this the phenomenal base? The hunger? Is this the causal base? Perhaps it makes the most sense to say that both are the bases.

Taking a staunch stance on the non-correspondence between causal and phenomenal base, Peter Carruthers points out:

Affect can be transparently accessible by virtue of being globally broadcast. But it doesn’t get tied to the representations involved in the cognitive appraisals that produce it, in such a way that affective representations and representations of those properties are unified together for purposes of global broadcast. On the contrary, affect from different sources tends to combine to form a single evaluation of whatever happens to be the object of attention, or to be the most relevant among objects of current attention. [...] Most of the objects or events that we react to affectively are highly complex, with many different properties that are potentially evaluatively relevant. [...] With time and learning, of course, we may develop theories about the properties of objects and people that influence us the most, and sometimes these theories may be correct. But there is no reason to think that the sources of affect are, as such, transparently accessible. I grant that in most cases [...] one can know the object of one’s affective state. But the phrase “object of one’s affective state” here needs to be read as involving a particular thing or event [...] One knows that one likes *this person* or that one is disgusted at *that action*, but there will be many different perceptually-embedded judgments occurring while one attends to an object or event, and many different aspects of it may be part of the content of the resulting perceptual state. One has no introspective access to which subset of these aspects provides the fine-grained propositional object of one’s affective state. [...] Although perception of the stimulus will give rise to numerous perceptual judgments [...] the resulting affect isn’t bound to any one, nor any subset, of these in particular (despite being caused by one or more in particular). (Carruthers 2011, pp. 146-147, 150)

Of course, one could try to adjudicate what the *real* base here is, settling for either the causal or the componential view. Having said that, adopting pluralism seems viable as well (cf. Prinz 2014; Fulkerson 2019). Finally, one can adopt nihilism: the feeling and the base are not actually related in any significant way. In each case, one will have to distinguish between how things phenomenally appear and how things metaphysically and aetiologically relate (Herzberg 2012).

Let me briefly remind you of the phenomenal features of affective experiences before I sum up the intentional features: affective experiences are phenomenally valenced, arousing and motivational—and gradeable in that (sections 4.3.1 and 4.3.2). Now, here is a summary of the intentional features. The intentionality of feelings has two parts: a *particular object* and a *formal object* where the feeling represents the former as bearing the latter. In this context, the formal object refers to feeling-specific properties such as the painfulness of pain, fearsomeness of fear, funniness of amusement etc. (section 4.4.2) Representing something to be painful or amusing constitutes an evaluation which is *phenomenally grounded* in the *valence* of affective experiences (see section 4.3.1 and 4.4.3). Concerning their particular objects, affective experiences can be in principle about all kinds of things, ranging from bodily events and parts over objects in the physical and imaginary realm (including fictions and propositions themselves), states of affairs in past, present and future over to propositional contents (section 4.4.1). To accommodate this broad range of intentional objects feelings rely on a division of representational labour with other mental states such as bodily sensations, perceptions, judgments, memories, imaginings etc. This *base* of a feeling is a mental state (or a set of states) that supplies a feeling with its particular object: the feeling does not access the particular object directly but through other mental states (this section).

4.5. Affective Experiences and Intuitions

In this section, I want to juxtapose the features of affective experiences with those of intuitions. If they admit of significant similarities, then, I think, this should give us some reason to take the idea that intuitions might be (some kinds of) feelings seriously, meriting further exploration. Recall that intuitions are 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced, and 7) nonvoluntary. Now, what about feelings? As we have seen, affective experiences are *intentional*. They can have all kinds of things as

their intentional or particular objects, including particulars and propositional contents. Affective experiences are, in a specific way, also *assertive* (cf. Helm 2001, pp. 43 sqq.). Let me explain: As we have seen feelings come with bases that can but do not have to be assertive. For instance, they can take assertive perceptions just as well as non-assertive imaginings as their bases (Lamarque 1981; Tye 2008). However, affective experiences seem assertive independently of whether their base is assertive or not: they are assertive in that they “claim” that their base-provided particular object exhibits the feeling-specific formal object. The affective experiences in which you feel a cramp in your calf, are afraid of a bear or are sad that Bambi’s mother died leave little doubt that the cramp *is* unpleasant, the bear dangerous and the fate of Bambi’s mother sad.

Suppose you see a spider and are afraid. Your fear [...] involves its striking you that your situation is dangerous [...] Again, your fear is not a mere *consequence* of things so striking you; nor does fear merely cause things to strike you that way; rather, fear is a state *in which* things so strike you. (Bain 2013, S87)¹²⁰

In fact, this kind of emotional assertiveness in the face of fictive scenarios is what gives rise to the “paradox of fiction” (Radford and Weston 1975). In this context, let me spend some more time on non-assertive imaginings as bases of feelings.¹²¹ If I imaginatively relive the death of Bambi’s mother, what is it that my soaring sadness is “claiming”? Is it that the death of Bambi’s mother constitutes a loss? I am inclined to accept this account as by and large accurate. Seriously, aren’t you saddened by the death of Bambi’s mother? Doesn’t your sadness lead to a flurry of loss-related thoughts? Don’t you want to cry? At the very least, my sadness quite assertively tells me that the death of Bambi’s mother is, well, sad. This it does despite the fact that, somehow, in the back of my mind, I am perfectly aware that the death of Bambi’s mother is a fiction.¹²²

¹²⁰Juxtapose this with how Bengson describes what happens when one has an intuition:

When reading Gettier’s paper [...] Professor Typical considers whether Smith knows that the man who will get the raise has ten coins in his pockets. It *strikes* her that, even though Smith is justified in believing that this is so, Smith does not know it. (Bengson 2015, 711 p., my emphasis)

¹²¹Note that not all imaginings are about fictions, at least not in the sense that the fictionality of the imaginative content is obvious to or acknowledged by the imaginer. Equally, not all fictive scenarios are part of an imagining — they might well be represented by perceptual experiences when we, say, watch a play.

¹²²Alternatively, what my sadness tells me assertively is perhaps something along the lines that *in the fictional world of Bambi, the death of Bambi’s mother is sad/a loss* (cf. Lewis 1978). This would make my sadness very accurate, though perhaps on the expense of overintellectualizing it.

It is presumably this belief that leads me to regulate my (initial) behavioural response accordingly (Sperduti et al. 2016): “Come on, don’t cry! It’s not real!” But that there is something that needs to be regulated is due to my sincere, perhaps somewhat gullible sadness. Usually, we don’t take our affective experiences to fictions at face value—but that does not mean that they have no face value. To the chagrin of people suffering from idiopathic pain, phobias, jealousy or just less than optimal affective responses, the processes behind our affective experiences are to some degree insulated from our (explicit) beliefs. This kind of belief-independence affective experiences have in common with illusory perceptual experiences that we know to be illusory. In the case of affective experiences, this is all the more unnerving since, in contrast to perceptual experiences, they usually single-handedly carry motivational force, potentiating often unwanted bodily and mental behaviours. Thus, while perceptual illusions are easily overpowered by better knowledge, unwanted “affective illusions” usually need wrestling. As Brian McLaughlin points out: “We easily handle visual illusions, but we may be suckers for affective illusions.” (McLaughlin 2010, p. 157)

There is another important contrast: In the moment of their occurrence, affective experiences can usually be merely challenged on the level of their behavioural insinuations, and then also only given a good amount of willpower. However, if one looks beyond the very moment of their occurrence, we find ourselves capable of affective learning in the sense that, given the right circumstances, practical knowledge and behavioural engagement, we can change the very way affective experiences make things appear. Jealousy might lessen in strength, dissipate and, if one believes practitioners of polyamory, even turn into compersion. In other words: You can easily behaviourally override the way things look but you can’t change the way things look. In contrast, it is not easy to behaviourally override the way things feel but you can change the way things feel.

To sum up the former point: affective experiences assertively mark their particular objects as bearers of properties that are their formal objects. In Bennett Helm’s words “emotions are evaluative feelings in the sense that they are a distinctive kind of passive assent to their targets as having the import defined by their formal objects” (Helm 2001, p. 59).

And these formal objects can plausibly be all kinds of properties that merit our concern and are of import to us. These properties, in turn, can be (causally) grounded in *inter alia* biological, bodily, social, linguistic, moral, aesthetic and *epistemic* imperatives. Quite clearly, feelings are *motivational* and their motivational force is *gradable*. They can

more or less strongly move us to behave or act in specific ways. What is important in the present context: they plausibly not only motivate physical but also *mental* behaviour and action, such as attending, remembering, reflecting or assenting (Clare and Huntsinger 2007). As Nico Frijda notes:

Many reactions that are considered “emotional” do not include much overt action. States of action readiness may remain just that: states of readiness. However, the actions for which action readiness is readiness include cognitive actions: changes in beliefs. Emotions have been defined by Aristotle as well as by Spinoza as inclinations to think one way rather than another. (Frijda 2008, p. 72)

Notwithstanding the fact that affective experiences are highly potent engines of persuasion, they do not amount to full commitment to what they would want us to believe or judge. Feelings occasionally stand in conflict with our (better or worse) judgments. Insofar they are *noncommittal*. It seems also obvious that we are patients with respect to affective experiences, i.e. feelings are *nonvoluntary*. Brady nicely summarises some of these aspects for emotional feelings:

It is clear that emotional experience is typically passive: [...] [W]hen we feel an emotion the import of our situation impresses or thrusts itself upon us. [...] [T]o say that the import of a situation impresses itself upon S is to say, roughly, that S is inclined to assent to or endorse this view of the situation. In other words, when S experiences an emotion she is subject to some kind of pressure to accept the relevant appearance (Brady 2009, pp. 420 sq.).

One of the central features of affective experiences is their *phenomenal valence*. It is a general phenomenal property specific to and shared among affective experiences (Weiss 2016). Now, phenomenal epistemic valence seems to resemble “general” phenomenal valence. Positive intuitions presumably feel somehow positive in comparison to negative intuitions and they do so in relation to truth; while negative intuitions feel somehow negative in comparison to positive intuitions and they do so in relation to falsity. This is indeed reminiscent of affective experiences: specific feelings are not only characterised by their valence but also by a formal object, i.e. feelings do not only feel somehow positive or negative, but also they let things appear in a certain way, for instance as offensive, sad, funny or pleasant (cf. Teroni 2018). Fear (or shame) feels somehow negative in comparison to, say, amusement (or an orgasm) and it does so in relation to danger (or self-directed inadequacy). Happiness (or pride) feels somehow positive in comparison to,

say, sadness (or migraines) and it does so in relation to success (or achievement). It seems plausible that phenomenal valence is part of the story that explains the phenomenal epistemic valence of intuitions. Phenomenal valence seems like a reasonable candidate to *phenomenally* ground and evaluation of a proposition being true or false.

From there it is not too far-fetched to see phenomenal valence also as a possible ground for the assertiveness of intuitions. As we have noted earlier, the assertiveness (and pushiness) of intuitions covaries with their phenomenal epistemic valence (see section 2.2.6). Helm links the assertiveness of emotions, something he calls “distinctively emotional assent”, directly to their valence which he couches in terms of pleasure and pain:

The idea of such distinctively emotional assent is implicit in the idea that emotions are evaluative feelings: being pleased or pained by things being thus and so is a kind of acceptance that things really are that way, an acceptance that falls short of full-blown judgment. (Helm 2001, pp. 45-46)

[E]motions are evaluative feelings in the sense that they are a distinctive kind of passive assent to their targets as having the import defined by their formal objects, and I have described such passive assents as feelings of pleasure and pain: to be afraid is to be pained by danger, to feel hope is to be pleased by the prospects for success, to feel frustration is to be pained by repeated failure to attain some good, etc. This appeal to pleasure and pain is intended to make three points: first concerning the way in which emotions, as passive assents to import, are evaluations of a kind that differs from evaluative judgment; second concerning the way in which emotions motivate subsequent action; and third concerning how emotions feel—their phenomenology. (ibid., p. 59)

In fact, it becomes hard to see which other class of mental states has anything resembling phenomenal epistemic valence. Perceptual experiences do not have valence¹²³ and while beliefs might be thought to have an epistemic valence, this valence does not seem to be phenomenal (Koksvik 2011; Martin and Dokic 2013).¹²⁴ So either intuitions are their own class of mental states characterised by the primitive property of phenomenal epistemic valence or they belong to the familiar class of affective experiences. In the latter case, their phenomenal epistemic valence could be partly explained by the general phenomenal valence of feelings. This does not appear implausible considering how Clore

¹²³But see Dokic and Martin 2015a.

¹²⁴Though those subscribing to cognitive phenomenology might disagree.

characterizes the impact of affect on judgment and thought in a way reminiscent of intuitions:

[A]ffect serves as information about the value of whatever comes to mind. Thus, when a person makes evaluative judgments or engages in a task, positive affect can enhance evaluations and empower potential responses. Rather than affect itself, the information conveyed by affect is crucial. [...] In tasks, positive affect validates and negative affect invalidates accessible cognition (Clore and Huntsinger 2007, p. 393).

A related similarity comes to the fore if we consider that positive and negative intuitions are phenomenal polar opposites. Now affective experiences harbour many pairs of phenomenal polar opposites: pain and pleasure, sadness and joy, disgust and delight, shame and pride, contempt and admiration, hate and love. On second thought, it seems that phenomenal polar opposites can be found *only* among affective experiences. A plausible explanation for this observation lies in the hallmark of affective experiences: the property that makes and breaks phenomenal polar opposite pairs is phenomenal valence. After all, phenomenal valence by itself embodies the most fundamental polar opposition: positive and negative, good and bad, “yin and yang”. Add determinate positivity or negativity along the same or a related property dimension, such as saddening and joyful, and you get a pair of affective experiences that appear to be phenomenal polar opposites.

This is presumably why it is hard to find phenomenal polar opposites among perceptual experiences and imaginings: the visual experience or imagining of red is not the polar opposite of the visual experience or imagining of blue. One might now argue that one example shows nothing since one can find also examples of affective experiences that have no polar opposites such as...anger?¹²⁵ If so, it should be enough to find just *some* phenomenal polar opposites among, say, perceptual experiences. The best I can come up with is the thermoceptive experience of coolness and warmth. Nevertheless, I think that this appearance can be most naturally explained by positing that thermoception often comes together with affective components, namely a positive or negative valence. It seems that the positive experience of warmth is the polar opposite of the negative experience of coolness. However, if both coolness and warmth feel good, are they then still phenomenal polar opposites? I doubt that they are. Thus, phenomenal polar opposites appear as exclusively affective phenomena.

¹²⁵But what about feeling flattered?

Affective experiences are *content-gradable* in a particularly interesting way. First: feelings come with other mental states as their bases. Some of these bases such as perceptual experiences and imaginings will have content-gradability, lending this feature to the affective experiences that take them as their base. Having said that, there is more to the story: it is not only that the content of the feeling which is provided by the base will appear more or less vivid because the base has the feature of content-gradability. Crucially, the fact that it is the content of a feeling will make a characteristic *additional* contribution to the vividness of the content. This is suggested by the ample empirical evidence that feelings facilitate content encoding in virtue of their affective properties such as valence and arousal (e.g. Lang et al. 1998; Schupp et al. 2003; Kensinger and Corkin 2004; Sharot et al. 2004; Phelps et al. 2006; Talarico and Rubin 2007) as well as their effect on attention:

Emotions “also function to enhance the *quality* of a subject’s representation of her emotional situation, by focusing attention onto important events and by keeping attention focused there. This results in an increased *sensitivity* to emotionally-relevant features in the subject’s environment. [...] Emotions thus function to enhance a subject’s evaluative construal of her situation, by making her additionally focused on and sensitive to features which constitute *reasons* for that construal” (Brady 2009, 422-423, his emphasis).

The upshot is that if paired with an affective experience, representational content appears as more *vivid*. Note also that the intensity or motivational force of a feeling does not have to co-vary with the clarity and distinctness of the content of their base. The perception of a blurry snake-like shape can trigger fear as intense as the clear and distinct perception of a snake. Something similar about intuitions is noted by Koksvik:

One possibility is to place variability in the content of the intuition [...] as variance in its determinacy. [...] However, [...] it might clearly but weakly seem to someone that torturing one innocent person to save 20 others is permissible (while it clearly and strongly seems to that person that torturing an innocent person to save a thousand is). [...] My own preference is to say that the pushiness of intuitional experience is itself gradable. (Koksvik 2017, p. 6)

So what we get is this: On a closer look, intuition experiences appear quite similar to affective experiences. This is reminiscent of the reason why Perceptualists exploit perceptual experiences as an analogy to elaborate features of intuitions: intuition experiences

are similar to perceptual experiences. There they stop, however. What explains this similarity is not that intuition experiences *are* perceptual experiences. Even Perceptualists dismiss such Literal Perceptualism as counter-intuitive. In contrast, I submit that the similarity between intuitions and affective experiences is straightforwardly explained: it's because intuitions just *are* affective experiences. Thus:

Affectivism: Intuitions are affective experiences.

Is this as counter-intuitive as Literal Perceptualism? Perhaps for some philosophers. However, as we have seen in section 4.1, to the “man on the street” and to psychologists this conception seems presumably quite natural. Of course, we cannot simply *jump* to the conclusion that intuitions are affective experiences from here and leave it at that. What I take these reflections to motivate, however, is that one can *reasonably* formulate the hypothesis that intuitions are *some* kind of affective experiences. This is the central idea behind Affectivism. To this, a sceptic might object that intuitions appear very different from classical examples of feelings such as hunger, fear or joy. Indeed, casting intuitions as affective experiences against *this* background might seem counter-intuitive at first, especially if one takes *only* bodily and basic or paradigmatic emotional feelings to be representatives of the class of affective experiences. However, as even James is willing to admit, the class of feelings comprises more: Chapter 5 will introduce *epistemic feelings* and chapter 6 will argue that they are affective experiences, emerging as the ideal candidate class for intuitions (chapter 7). I am not going to jump to the conclusion that intuitions are (some kind of) affective experiences, I am going to work up to it slowly.

Epistemic Feelings

5.1. Introducing Epistemic Feelings

Epistemic feelings have been broadly described as “feelings that enter into the epistemic processes of inquiry, knowledge and metacognition” (de Sousa 2008, p. 189).¹²⁶ Among the phenomena that populate our mind, epistemic feelings have only recently begun to attract the attention of philosophers (e.g. de Sousa 2008; Arango-Muñoz 2014).¹²⁷ However, as Mangan and Prinz note, epistemic feelings have made appearances in the writings of James and Descartes (Mangan 2001; Prinz 2011).¹²⁸ Looking back, de Sousa goes as far as to locate epistemic feelings in decisive moments of philosophical history such as “when Meno’s slave boy recognizes the correctness of Socrates’ demonstration of the way to double the square” in Plato’s Meno or when in the course of the Meditations, Descartes arrives at his *cogito* which finally seems to him *clear and distinct*, rescuing him from drowning in radical doubt (de Sousa 2008). We also find less well-known but exceptionally vivid mentions of epistemic feelings in Dewey’s writings:

¹²⁶Epistemic feelings, with possibly some variation in extension, are also sometimes called noetic, cognitive or metacognitive feelings as well as epistemic or intellectual emotions (e.g. Clore 1992; Koriat 2000; Stocker 2004; Prinz 2007, p. 350, 2011, pp. 190 sq.; Morton 2009; Dokic 2012; Goldie 2012; Carruthers 2017b).

¹²⁷That is to say, a *larger portion* of the class of epistemic feelings has recently become topic of interest. Some epistemic feelings such as surprise, doubt and curiosity (but rather rarely construed as feelings) have been around for a while (cf. Prinz 2011, p. 190; for some exceptions for the case of doubt see Hookway 1998; Thagard 2004).

¹²⁸Although he does not speak of “epistemic feelings”, Hookway identifies states that are best construed as such in the work of James, Peirce and Quine (e.g. Hookway 1998, 2002, 2003).

[T]hese “feelings” have an efficiency of operation which it is impossible for thought to match. Even our most highly intellectualized operations depend upon them as a “fringe” by which to guide our inferential movements. They give us our sense of rightness and wrongness, of what to select and emphasize and follow up, and what to drop, slur over and ignore (Dewey 1929, p. 299).

In contrast to their relatively new (re)appearance on the philosophical stage, epistemic feelings have been a topic of intensive research in cognitive science for more than 50 years (e.g. Hart 1965). There, epistemic feelings have been identified as playing important roles for e.g. memory, perceptual decision-making and reasoning (e.g. Koriat 2000; Volz, RübSamen, et al. 2008; Thompson, Prowse Turner, et al. 2011). Inspired by these empirical findings, philosophers have proposed innovative accounts of puzzling phenomena such as animal metacognition, hallucinations and delusions but also of more familiar topics such as moral judgement and aesthetic experiences (e.g. McLaughlin 2010; Dokic 2012b; Dub 2015; Dokic 2016; Proust 2013; Clavien and FitzGerald 2018).

I propose that given what we know about epistemic feelings, they appear as an ideal psychological class in which to locate intuitions. What would be the gain in so doing? One of them, I submit, is that identifying intuitions as epistemic feelings elucidates the nature of intuitions further than doxastic or perceptualist alternatives. How? Perceptualism succeeds in describing the features that intuitions appear to have. This sets it favourably apart from Doxasticism. A major drawback of Perceptualism is, however, that it either posits intuitions as a *sui generis* kind of states or puts intuitions into a class of mental states (e.g. seemings or presentations) whose postulation is contested. Such a move does not *explain* the features assigned to intuitions in a satisfying way.¹²⁹ At least some versions of Doxasticism (were they otherwise sound) would not suffer from this drawback since (part of) the specific nature of intuitions could be explained by appeal to the rather well-established kind of state that intuitions are (e.g. judgments or beliefs), about which we arguably know more than about intuitions.¹³⁰ Against this background, a theory that 1) allows intuitions to have the right features (as Perceptualism does) *and*

¹²⁹Perhaps, this could be considered not too much of a problem for the purposes of most Perceptualists. What they are primarily after is not explaining why intuitions have the features they have but showing how it comes that intuitions justify beliefs. The idea is that intuitions qualify as justifiers qua having the described features. Why they have the features they have does not matter that much for Perceptualists. However, I (and perhaps also some Doxasticists) think it should matter. So I want it to matter for Affectivism.

¹³⁰...or do we? I think philosophers tend to overstate our grasp on judgments and beliefs. Be it as it may, this leaves the point about our poor(er) grasp on intuitions and our relatively good grasp on affective experiences unaffected.

2) *explains* these features by appeal to a kind of state that is relatively well-established and about which we know more than about intuitions (as Doxasticism does) would be preferable to theories that only manage one of the two.

I will argue that Affectivism is such a theory: it acknowledges the features and accommodates them by appeal to a fairly well-established kind of state. The claim that intuitions are epistemic feelings will become instrumental in accounting for 1) and 2). In the present chapter I, want to show that we know quite a bit about epistemic feelings specifically. It seems to me that this by itself renders the identification of intuitions with epistemic feelings *informative*. I mean, even Plato and Descartes were talking about them, weren't they? Notwithstanding the mechanistic understanding of epistemic feelings I am going to sketch and the mentioned historical pedigree a sceptic would probably still object that epistemic feelings do not seem to be "a fairly well-established kind of state". In the next chapter, I am going to counter this objection with more than "just" talk about mechanisms and historical pedigree. I am going to show that they, actually, are affective experiences—and affective experiences are a fairly well-established kind of state. So in accounting for 1) and 2) I do not only want to exploit the fact that we know quite a bit about epistemic feelings specifically but to additionally exploit the fact that we know even quite a bit more about the kind of state of which epistemic feelings are a specific subclass, namely affective experiences.¹³¹ But this is for later. Now we are going to get a better grasp on epistemic feelings (section 5.2) and then we will dive deeper into their mechanisms (section 5.3).

5.2. Getting a Feel for Epistemic Feelings

Epistemic feelings constitute a class of experiences which pervades our mental lives. A good way to get a grasp on these familiar phenomena is to consider some typical situations where one would experience certain epistemic feelings. Remember the last time you encountered a person seemingly for the first time but had the impression that you had seen him or her before. In such cases, you might discover that, in fact, you have seen the person before, e.g. somewhere in passing or on an occasion where you did not interact with the person. And sometimes you might find out that you have seen someone who looks similar, e.g. a similarly looking actor. Often, however, the reasons

¹³¹In case of a superclass-subclass relation all true statements about the superclass typically extend to the subclass. This is why identifying the superclass to which epistemic feelings belong facilitates the project of characterising them and by extension intuitions.

for why the person “looks” familiar remain opaque. Such *feelings of familiarity* (FOF) (e.g. Jacoby and Whitehouse 1989; Whittlesea and Williams 1998) happen to everyone from time to time: In having a FOF all kinds of particular objects might appear as familiar. As a consequence, we often take FOFs to mean that we have encountered a certain stimulus before. Clearly, FOFs are not restricted to the visual modality. Think about FOFs when confronted with the beginning of a song, the sound of a name, the smell, or taste of some food ingredient or the texture of some surface. That is, things cannot only “look” familiar but they can also “sound”, “smell”, “taste” and (in the tactile sense) “feel” familiar (Plailly et al. 2007).

[W]hat is the strange difference between an experience tasted for the first time and the same experience recognized as familiar, as having been enjoyed before, though we cannot name it or say where or when? A tune, an odor, a flavor sometimes carry this inarticulate feeling of their familiarity so deep into our consciousness that we are fairly shaken by its mysterious emotional power. But strong and characteristic as this psychosis is [...] the only name we have for all its shadings is ‘sense of familiarity’. (James 1890, p. 252)

In fact, abstract or “mental” objects such as ideas, stories or propositions can feel familiar as well. Just recount the last time you exclaimed, “I had the exact same idea!”, after hearing what a friend had to say. It is also plausible that it is not only the content of a mental state that can feel familiar: certain mental states or events can feel familiar *themselves*, such as when we respond emotionally in the “good old” or the “same old” way. A closely related feeling is the puzzling *déjà-vu experience* (Brown 2003) where, against your better knowledge, it seems to you as if you have already been in the situation you find yourself in now. For an instant, it appears as if you’re reliving a moment from the past, that what happens has already happened and that you can predict what will happen next.¹³² Note that, plausibly, familiarity is not encoded in the content of one’s perceptual experience. It is not a property that one sees, hears or tastes (with the “mind’s sensory organs”) such as (imagined) colours or shapes. Rather, it is the FOF that accompanies or follows the experience that makes some content *feel* familiar to us.

¹³²Despite their similarity, a curious difference between feelings of familiarity and *déjà-vu* experiences is that the former occurs in connection with individual objects such as persons, physical objects, and situations while the latter typically occurs in connection with whole situations only, not with specific aspects of them. Maybe this is because in the case of a *déjà-vu* experience, a FOF is caused by an opaque aspect of the situation so that it fails to be bound to a specific aspect of it. This makes the familiarity permeate the whole representation of the situation, putting it obviously at odds with one’s knowledge.

Everyone knows what it is like to feel that something is familiar, and I doubt anyone would claim that familiarity, in itself, is a sensory experience. The feeling of familiarity is not a color, not an aroma, not a taste, not a sound. It is possible for the feeling of familiarity to merge with, or be absent from, virtually any sensory content found on any sensory dimension. (Mangan 2001, p. 3)

Capgras delusion is a striking demonstration of the dissociation between perceptual experiences and FOFs. Capgras patients have the delusion that loved ones are replaced by identically *looking* imposters of some kind. This delusion is believed to be caused by an unusual phenomenology: The Capgras patient notes that the person or pet looks right. The person or pet, however, does not *feel* right to the patient.¹³³ In fact, Capgras phenomenology does presumably not only consist of an *absence* of a FOF but the presence of an intense feeling of *unfamiliarity* or *strangeness*, the negative twin of FOF (Ratcliffe 2007a; Young 2007b,a, 2009).¹³⁴ This, presumably with some other factors, leads to the delusion that the loved one was replaced by an imposter of some kind.¹³⁵

It is worth mentioning that in everyday life FOFs are most salient when they occur (or fail to occur) in situations where one does not expect (or expect) them to occur (Whittlesea and Williams 1998; Whittlesea and Williams 2000, 2001a; Whittlesea and Williams 2001b). In these situations, the presence (or absence) of a FOF comes together with an absence of apparent and available reasons to account for its presence (or absence). In such instances (the absence of) FOFs constitute unexpected ingredients in one's overall phenomenological state, making them (or their absence) "stand out" or "pop-up" phenomenally. This usually calls for an explanation that we seek in the case of FOFs with

¹³³Remember what I noted above: mental states themselves can feel familiar. In fact, de Sousa suggests an explanation of Capgras syndrome along these lines:

The sight of a close relative – a parent, in the case of Ramachandran's patient Arthur – normally triggers an affective response, which is itself subject to a 'familiarity' evaluation. In Arthur's case, the direct link to the area in charge of generating the affective response is missing. As a result, the affective response to his father is not produced. This sets up an incongruity between the strictly cognitive familiarity check that applies to the face and the missing familiarity check effected by the expected affective response. (de Sousa 2008, pp. 198-199)

¹³⁴I think that such feelings of unfamiliarity are nothing intrinsically pathological. I have them sometimes when e.g. visiting a place where I spent a lot of time in the past but haven't been for a while.

¹³⁵The "other factors" have to explain the additional feature that Capgras patients not only have an unusual experience but furthermore respond to it by adopting the impostor belief which appears to be demonstrably false. Researchers have pointed towards further features of the *phenomenology* of the Capgras experience or reasoning deficits (e.g. Davies et al. 2001).

apparent strangers. In fact, it is the unexpected absence of a FOF vis-à-vis a loved one that leads a Capgras patient to adopt the absurd but explanatory belief that the loved one has been replaced by an imposter. Ordinarily, however, the fact that something feels familiar (a close friend) or unfamiliar (some stranger) does not draw our attention since this is what one would expect.¹³⁶

Another epistemic feeling might turn out to be familiar to you¹³⁷ when you think back of your time in school. In situations when a teacher was probing the knowledge of one of your classmates asking her questions such as “When did the French Revolution start?” or “What is the capital of Australia?” it might have occurred to you that you know the answer.

Importantly, this feeling struck you *before* you had the chance to retrieve the relevant information from memory. This *feeling of knowing* (FOK) (e.g. Hart 1965; Koriat 1994, 2000; Thomas, Lee, et al. 2016) was perhaps quickly followed by disappointment that the teacher has not picked you to demonstrate your knowledge: You would have known that the answer is “1789!” or “Canberra!”. On other occasions, you might have considered yourself fortunate that the teacher did not pick you since the requisite FOK was absent. Or perhaps you had an even more pronounced experience of the negative twin of the FOK: the feeling of *not knowing* (Koriat and Lieblich 1974; Kolers and Palef 1976; Glucksberg and McCloskey 1981; Klin et al. 1997; Bar-Anan et al. 2009). Plausibly the ability to exploit FOKs makes a successful player of quiz shows such as *Jeopardy*. In the TV show, it is essential to buzz as fast as possible if one knows the answer to a general knowledge question (Michaelian and Arango-Muñoz 2014, p. 99).

One might worry that in the described cases one does not have a FOK before retrieving some memory item—one just retrieves it and that is how one knows that one knows.¹³⁸ I admit that this is well possible. And even when a FOK occurs: most of the time the

¹³⁶An alternative to the present-but-not-salient interpretation is the possibility that FOFs are simply absent in cases where familiarity is expected. That is, FOFs are only present (and not just salient) if familiarity or the absence of it is unexpected. Note that this expectation does not have to be present on a personal level (one might, however, become conscious of one’s hitherto unconscious expectation through the FOF or its unexpected absence).

¹³⁷Pun intended.

¹³⁸That is, one *infers* that one knows the answer from the fact that it comes to mind and perhaps from the way it comes to mind, say, quickly (see footnote 200). Alternatively, the inference might be more indirect: you recall that you diligently studied the Oceanic capitals yesterday and from that you infer that you know the capital of Australia. In such cases, the judgment “inferentially derives from independent beliefs based on memory” (Dokic 2012, p. 304). This is what Koriat has termed a theory- or information-based judgment and contrasted with experience-based judgments such as the ones based on FOKs (Koriat and Levy-Sadot 1999, see section 5.3.1.2).

succession between the FOK and the retrieval of the information is so swift and seamless that there is no awareness of a temporal succession or of the fact that a FOK occurred at all.¹³⁹ So in the cases described, the answers to the questions are short and thus quickly retrieved, perhaps so quickly that one might miss one's antecedent FOK, if any. Though what if I ask you whether you can recite your favourite song, poem or quote? That is, I ask you for a longer string of information that is not as quickly retrieved as a single word or number. In such cases, you will surely sometimes feel that you know the information before fully retrieving it. In fact, a major line of research on FOK comes from metamemory studies on the ability of subjects to judge whether they have mastered some studied material, i.e. substantial amounts of information.

Suppose now, you would, encouraged by your FOK, go on to retrieve the relevant information. Although a FOK might be a relatively reliable predictor of retrieval success, it does not infallibly guarantee it. And so in some cases you might run into a state of consciousness masterfully captured by James:

Suppose we try to recall a forgotten name. The state of our consciousness is peculiar. There is a gap therein; but no mere gap. It is a gap that is intensely active. A sort of wraith of the name is in it, beckoning us in a given direction, making us at moments tingle with the sense of our closeness, and then letting us sink back without the longed-for term. If wrong names are proposed to us, this singularly definite gap acts immediately so as to negate them. They do not fit into its mould. And the gap of one word does not feel like the gap of another, all empty of content as both might seem necessarily to be when described as gaps. [...] The rhythm of a lost word may be there without a sound to clothe it ; or the evanescent sense of something which is the initial vowel or consonant may mock us fit fully, without growing more distinct. Every one must know the tantalizing effect of the blank rhythm of some forgotten verse, restlessly dancing in one's mind, striving to be filled out with words. (James 1890, pp. 251-252)

This exasperating phenomenon is commonly known as the *tip-of-the-tongue experience* (TOT) (e.g. Brown and McNeill 1966; Schwartz and Metcalfe 2011; Schwartz and Brown 2014), the unpleasant feeling that the relevant information is (stuck) on the tip of your tongue. This is, you are in possession of the relevant information but are currently unable to produce it. It is typically in instances of TOTs where the occurrence of component

¹³⁹ *Awareness of* a FOK and *having a FOK* are not the same. This amounts to the distinction between reflexive or second-order consciousness and first-order consciousness.

FOKs is most salient. Similar to FOFs, FOKs and TOTs appear to be transmodal: They can be had not only in response to linguistic representations but also to e.g. odours (Jönsson and Olsson 2003; Jönsson and Stevenson 2014; Cleary et al. 2010). Furthermore, practitioners of sign language seem to have “tip-of-the-finger experiences” where one feels that one knows the visual sign (component FOK) for a concept but cannot recall the movements of the hand that express it (Thompson, Emmorey, et al. 2005). Finally, Chinese speakers appear to have “tip-of-the-pen experiences” where they felt they know a spoken word (component FOK) and feel as if they are about to recall how to make the written Chinese character for it (Sun et al. 1998).

Here are some brief descriptions of other epistemic feelings (more examples of epistemic feelings will follow in the remainder of this thesis):

- *Curiosity*: a state that motivates exploration and knowledge acquisition (e.g. Carruthers 2018).
- *Surprise*: a state resulting from something conflicting with one’s (implicit) expectations (e.g. Meyer, Niepel, et al. 1991; Meyer, Reisenzein, et al. 1997).
- *Doubt*: the feeling that the epistemic status of a proposition is unclear (e.g. Hookway 1998).
- *Feeling of performance (un)certainty or confidence*: the feeling that an activity is (un)likely to be successful (e.g. Bach and Dolan 2012; Koriat 2012, see also section 7.3.2).
- *Feeling of perceptual uncertainty*: the feeling that a perceptual representation is imprecise (e.g. Dokic 2014b).
- *Feeling of forgetting*: the feeling that one is forgetting something (Halamish et al. 2011; Arango-Muñoz 2013).
- *Feeling of meaningfulness*: the feeling that some representation (e.g. written or spoken word) is meaningful (e.g. Rapp and Goldrick 2000; Hicks et al. 2010; Dodd 2014, see also sections 6.3, 6.5, 8.3 and footnote 208)

After these illustrations, one might say: “Indeed, I got a pre-theoretical grasp of epistemic feelings. I have encountered such phenomena in my life. But, *what are epistemic feelings?*” In analogy with my guiding question about intuitions one could take this question to ask for the kind of state that epistemic feelings are: “What kind of state

are epistemic feelings?” (kind-question) In analogy to the proposed answers to the kind-question about intuitions, possible answers to this question would be that, say, epistemic feelings are beliefs or judgments or that they are perceptions or that they are similar to beliefs or perceptions. On the other hand, in a more empirically minded manner, one could take this question to ask for the mechanisms underlying epistemic feelings: “What are the mechanisms behind epistemic feelings?” (mechanism-question).¹⁴⁰ Of course, the answers to these questions are intimately intertwined. If you answer one of the two (main) questions, at least parts of the answer to the other question will probably fall out of it. If, for instance, you propose epistemic feelings to be perceptions, then it is plausible to assume that the mechanisms behind epistemic feelings will be, at least on some level (e.g. a functional one), importantly alike to the ones behind our perceptual modalities.¹⁴¹

Now, the way from the mechanism-question to the kind-question might seem not as straightforward. That is, answers to the mechanism-question can be significantly *underdetermined* in regard to an answer to the kind question.¹⁴² One will have to make a

¹⁴⁰This question can itself be subdivided into a proximal, an ontogenetic distal, and a phylogenetic distal version. The proximal version is: “In the moment epistemic feelings occur, what are the mechanisms that bring them about?”. The ontogenetic distal version is: “In the course of an individual’s life, how do epistemic feelings develop?”. The phylogenetic distal version is: “In the course of evolution, how did epistemic feelings develop?”. I will focus on the proximal version here. Note that there is a further ambiguity when it comes to the proximal version: are we talking about the mechanisms that diachronically *cause* or are we talking about the mechanisms that synchronically *ground* epistemic feelings? That is, are we talking about the proximal aetiology of epistemic feelings (diachronic proximal version) or are we talking about, well, the mechanisms that (assuming a naturalistic metaphysics) somehow *are* the epistemic feelings (synchronic proximal version)? This way of putting it makes clear that the synchronic proximal version is actually (a layer of) a specific answer to the kind-question that embeds a kind-proposal into a physicalist-functionalist framework à la “epistemic feelings are kind X because they are instantiated by (physical/brain) mechanisms that serve the function to X”. In turn, the plausibility of such a mechanistic proposal, i.e. an answer to the question “What are the mechanisms that (synchronically) ground/instantiate epistemic feelings?” to the kind-question, will to a large extent depend on the specific proposal to the diachronic proximal question “What are the mechanisms that (diachronically) cause epistemic feelings?”. Thus, I will have something to say about both. Note further that a mechanism can be more or less proximal. Consider the example of visual experience of colour. There will be something “maximally proximal” that grounds it (hard problem of consciousness ahead), say, some kind of brain activity responsible for colour consciousness; something that proximally causes this brain activity such as antecedent photoreceptor activity; and something that causes colour experience/colour experience brain activity more distally (but still “proximally”) by causing the antecedent photoreceptor activity such as electromagnetic radiation in a visible frequency spectrum. It seems that all these variously proximal components are important parts of an answer to the proximal question.

¹⁴¹In fact, looking at the mechanisms behind epistemic feelings can be thought of as a way of *testing* specific proposals to the kind-question.

¹⁴²The same, is, of course, true for the underdeterminacy of a kind-proposal relative to a mechanism-proposal. It will be rare that a kind-proposal made not based on findings about the specific mechanisms, will specify the mechanisms *in sufficient detail*. However, it could fruitfully guide the inquiry

point (e.g. a functional one) why the specific mechanisms underlying epistemic feelings should or should not be likened to, say, perception or belief mechanisms. Ideally, specific observations about the mechanisms should lend support to one's answer to the kind-question (and vice versa). When it comes to epistemic feelings, the mechanism-question has received relatively more attention. This is likely due to the predominantly empirical angle on the matter.

In the next section, I will focus on the different proposals for the mechanisms of epistemic feelings. After having provided the reader with a feel for the psychological mechanisms behind epistemic feelings, I will present a proposal to the kind-question in chapter 6. Drawing on the hallmarks of affective experiences introduced in chapter 4, I will argue that the empirical evidence points in favour of conceiving epistemic feelings as affective experiences. In section 6.5 I will back this claim further with some phenomenological and theoretical points and will apply the conceptual resources introduced for affective experiences to analyse epistemic feelings.

5.3. The Epistemic Feeling Machinery

In order to sketch a picture of the mechanisms assumed to be behind epistemic feelings, I will focus here on two well-researched epistemic feelings: the feeling of knowing (FOK, section 5.3.1) and the feeling of familiarity (FOF, section 5.3.2). Extrapolating from findings on FOKs and FOFs to epistemic feelings in general is fairly common for work that tries to give general accounts of epistemic feelings (e.g. Dokic 2012; Arango-Muñoz 2014; Proust 2014). In other words, we might get *general* models of the mechanisms behind epistemic feelings based on findings concerning the mechanisms behind *specific* epistemic feelings. In section 5.3.3 we will touch upon the most influential general framework along these lines, which understands (some) epistemic feelings as a form of metacognition. It is an important endeavour to identify generalities in how epistemic feelings work and for that matter to extrapolate from specific findings, especially if one wants to establish epistemic feelings as a coherent class of experiences. However, one has to be cautious not to overextend specific findings, sometimes even when it comes to one and the same kind of feeling. For instance, if some specific property tends to lead to an experience, it does not mean *ipso facto* that no other properties will lead to the

into these mechanisms.

same experience (or another experience from the same class) as well. With this caveat in mind, I begin sketching the proposals.

5.3.1. The FOK (and Co.) Machinery

A great deal of what we know about the causal mechanisms behind epistemic feelings comes from metamemory research on FOKs and TOTs (for reviews see Schwartz 1994; Schwartz and Metcalfe 2011, 2014; Thomas, Lee, et al. 2016). Over the course of years, different models have been proposed to account for the workings behind these feelings. I will first briefly introduce the Direct Access model which in recent years has become somewhat unfashionable. Then I will focus on the contemporary majority view: the Heuristic Inference model that posits an inferential process that produces FOKs on the basis of cues such as question familiarity, information accessibility and processing fluency. Furthermore, we will learn that these cues are not mutually exclusive in determining epistemic feelings but that epistemic feelings are multiply determined. In the course of this section, we will also encounter feelings such as the TOT and feelings of information availability.

5.3.1.1. Direct Access

Originally, an internal memory-monitoring mechanism that is independent of memory search and retrieval has been proposed to produce the FOK (and TOT) (Hart 1965; Brown and McNeill 1966). Importantly, this monitoring mechanism itself was assumed to have direct access to the memory states representing the sought after information (memory traces). In virtue of this direct access, the mechanism would be able to monitor the presence of the memory trace as well as its strength and, based on that, assess the likelihood of future retrieval success—even though the information is inaccessible to memory search and retrieval. If the estimated likelihood would reach a threshold, the mechanism would trigger a FOK which would, in turn, lead to a judgment that one “knows” or is in possession of a specific piece of information (“FOK-judgment”).

5.3.1.2. Direct Access, Denied: Heuristic Inferences

The direct access model was highly influential but has meanwhile largely fallen from favour in the light of more recent empirical findings. The direct access view has been

superseded by *heuristic models* of the mechanisms behind FOKs. These heuristic mechanisms have been proposed to operate over various accessible non-target cues. These cues, in turn, can either be independent of search and retrieval processes or pertain to the search and retrieval processes themselves, the cues being features or outputs of these very processes. All these views have in common that they reject the idea of a monitoring mechanism that has independent access to memory traces in favour of a heuristic mechanism that has access to accessible non-target cues.

The mechanisms in question are heuristic because they operate according to heuristic rules or, shortly, heuristics. Heuristics in this context are understood as simple if-then rules used to infer whether something is the case (e.g. a criterion satisfied, a distal property present) based on accessible, proximal information that (imperfectly) correlates with the target dimension. So heuristic rules will take the following form: if A (e.g. a certain condition obtains), then X (e.g. “trigger FOK” or “information is present”). Note that one needs to distinguish between implicit and explicit heuristics. Implicit heuristics are utilised by processes that operate below the level of consciousness, i.e. the application of the heuristic rule takes place on a sub-personal level. In other words, the (unconscious) heuristic inference is effected by unconscious processes. The product of this inference are states with a meaning specified by the then-bit of the heuristic rule. Explicit heuristics, on the other hand, are heuristic rules that are consciously applied: based on the conscious observation that something is the case, the subject makes the conscious inference that something else is the case. In other words, explicit heuristics are employed in conscious heuristic reasoning. Now, the heuristic mechanisms assumed to trigger FOKs are unconscious and the heuristics deployed by them implicit.

In the literature judgments that one “knows” are often somewhat suboptimally called “FOK-judgments”. However, plausibly, judgments that one “knows” do not have to be based on FOKs (and thus mediately on the application of an implicit heuristic that generates the FOK) but can be the result of using 1) explicit heuristics, 2) explicit non-heuristic rules or 3) something else entirely. In case 2) the “know-judgment” might be based on some kind of direct access to the target dimension (i.e. one’s memory or knowledge) as suggested by the direct access model.¹⁴³ In case 3) it might, in fact, get even more direct than direct access: the target dimension (i.e. knowledge) itself might *directly*

¹⁴³Note that although the direct access model is operating according to an implicit rule (“if memory trace is present (and strong enough) then trigger FOK”), it would not be heuristic since it operates over the target dimension itself: the memory trace.

cause the know-judgment, without any intermediate inferential steps whatsoever.¹⁴⁴ Insofar know-judgments can but do not have to be FOK-judgments. In fact, they can be based on an ensemble of FOKs and other factors. This is why Koriat and Levy-Sadot helpfully distinguish between experience-based and information-based judgments (Koriat and Levy-Sadot 1999).¹⁴⁵ Experience-based judgments are based on immediate experiences such as a FOK, leading to an experience-based judgment that is properly called a FOK-judgment. Information-based judgments are judgments that are based on (various) other (additional) factors than experiences, such as conscious inferences from observation etc.

These reflections on the FOK-*judgment* highlight something about the FOK itself (and epistemic feelings in general). As the direct access model suggests, it is at least *conceptually* possible that FOKs do not have to be the result of implicit *heuristics*. Conceptually, they can also be the result of, say, implicit non-heuristic rules or even some kind of unconscious direct causation. While these possibilities seem poorly empirically supported in the case of FOKs, they are worth keeping in mind since the jury on other epistemic feelings is still out.

5.3.1.3. Exhibit 1: Cue Familiarity

Now, what cues do the heuristic mechanisms behind the FOK take as their input? Over the years, different cues have been proposed to trigger the mechanisms generating FOKs. One such is *cue familiarity* (e.g. Reder and Ritter 1992; Metcalfe et al. 1993). Definitions of cue familiarity are hard to come by but I presume it can roughly be understood as a subject's acquaintance with elements of a question or information query, i.e. the cue or target pointer that prompts memory search. Whether some element is familiar or not is, in turn, determined by success or failure to recognise it. The unconscious metamnemonic mechanism will then apply the cue familiarity heuristic (something along the lines "if

¹⁴⁴Note that the direct access model still assumes an internal monitor that makes an *inference* to knowledge via triggering the FOK after directly accessing the memory representation (if any). In a hypothetical "direct causation model" there is no such inferential step. To illustrate this, imagine there would be an experience of a specific neuron firing. We would undergo this specific experience immediately whenever and only when the specified neuron fires. The mechanism involved would not be properly called to have direct access to the firing of the specific neuron, nor would there be any heuristics involved. The mechanism in question would rather simply be direct causation. That is, the relation between the specific neuron firing and the experience of this specific neuron firing is wholly non-inferential—it is a relation of direct causation. (This is perhaps akin to what direct realists have in mind for perceptual experiences.) Of course, there is still some logical space between different degrees of "heuristicity", direct access and direct causation.

¹⁴⁵In the literature it is, alas, often not clear on which side of the distinction one stands.

a cue is familiar, then generate FOK”) and (not) give rise to a FOK, which leads to a judgment that one knows the target (FOK-judgment). In fact, in different circumstances such as another task, the very same recognition would have likely given rise to a FOF (which can be expected or unexpected) instead of a FOK.¹⁴⁶ Cue familiarity emerges as a cue supplied by recognition rather than memory search and retrieval processes. Insofar, it is a cue that is independent of the latter.

5.3.1.4. Exhibit 2: Accessibility

Another cue that has been proposed and that, in contrast to cue familiarity, *is* dependent on search and retrieval is *accessibility* (e.g. Koriat 1993, 1994). In Koriat’s own words:

FOK is based on the overall accessibility of pertinent information regarding the solicited target [...] This account assumes that monitoring does not precede retrieval but follows it: It is by trying to retrieve a target from memory that a person can appreciate whether the target is “there” and worth continuing to search for. This occurs because even when retrieval fails, people may still access a variety of partial clues and activations, such as fragments of the target, semantic and episodic attributes, and so on [...] These partial clues may give rise to a sheer feeling that one knows the answer. An important assumption of the accessibility account is that participants have no direct access to the accuracy of the partial clues that come to mind, and therefore both correct and wrong partial clues contribute to the FOK. (Koriat 2007, p. 299)

At the heart of this account lies the implicit accessibility heuristic: a FOK is triggered if a sufficient amount of target-related information is retrieved or accessible. At the same time, Koriat assumed that the subject does not know whether the accessible information is *really* related to the target or not.¹⁴⁷ However, recent findings cast some doubt on this assumption, showing that sometimes actual target relevance does matter: if the information is conceptual or semantic in nature, then actual target relatedness has an

¹⁴⁶This relationship between recognition, FOF, FOK and FOK-judgment is at least the most plausible construal according to me. There are *complementary* alternatives: 1) Recognition might trigger a FOF which is used like a FOK, leading to a FOK-judgment. 2) Recognition might trigger a FOF which might in turn trigger a FOK that leads to a FOK-judgment. 3) Recognition leads to a FOK-judgment. 4) Recognition leads to a FOF *and* a FOK that leads to a FOK-judgment.

¹⁴⁷Koriat, however, argues that by their very nature memory mechanisms triggered by the information query will produce more related than unrelated information (“commission errors”) (Koriat 1993, pp. 614 sqq.; see also Koriat and Goldsmith 1994).

effect on the magnitude of the FOK (Thomas, Bulevich, et al. 2011, 2012). Be it as it may, the *amount* of accessible information will contribute to the magnitude of the FOK, accounting for its gradeability.

5.3.1.5. Main Exhibit 3: Processing Fluency

Finally, the cue that has received by far the most attention in FOK research and research on epistemic feelings in general is *processing fluency* (Alter and Oppenheimer 2009; Unkelbach and Greifeneder 2013a). Processing fluency describes a process property that refers to the “ease”, usually understood as the speed, with which a given cognitive process is executed. The dominant experimental proxy for the fluency of a cognitive process is the measure of response time or latency, i.e. how fast a subject responds in performing an experimental task.¹⁴⁸ Naturally, given that there are many kinds of cognitive processes, there are many kinds of fluencies: perceptual fluency, retrieval fluency, encoding fluency, answer fluency, conceptual fluency, to name a few (Alter and Oppenheimer 2009). Now, manipulations of fluency have been proven to be a reliable way of inducing epistemic feelings. In fact, it has become common in the literature to distinguish between “objective fluency” and “subjective fluency” (e.g. Reber, Wurtz, et al. 2004; Forster et al. 2013). Objective fluency simply denotes processing fluency as described above while “subjective fluency” is being used as a generic term for the subjective experiences that result from objective fluency, i.e. typically for fluency-based epistemic feelings.

However, when it comes to “subjective fluency” it is important to make a distinction that is seldom made in the literature. For that we need the following background: Some researchers have called the tendency to infer the presence of certain properties on the basis of fluency the “fluency heuristic” (you might have guessed it) (Jacoby and Dallas 1981; Kelley and Jacoby 1998; Hertwig et al. 2008). As mentioned earlier,

¹⁴⁸In descriptions of processing fluency one sometimes also finds talk of “efficiency”, “simplicity”, “effortlessness” next to speed. However, in contrast to speed, the usage of these terms either requires additional assumptions that go beyond the observations made (usually, fast response times) or they are metaphorical. As such, the notion of processing fluency, can appear quite shifty if not fixed. Just consider: one can almost always talk about some process being in one way or another more efficient or “simple”—*even* in the absence of any observation of fast response times. This threatens to trivialize the notion of processing fluency. That’s why process speed seems as the only “hard currency” of processing fluency to me: as a property it is tangible and can stand on its own without requiring further assumptions. At the same time it can serve as a proxy for other properties such as efficiency, simplicity or effortless—if they can be meaningfully defined in a way that goes beyond processing speed. These remarks will become more clear in the course of this section and will be taken up again in section 8.2.

though, one should distinguish between implicit and explicit versions of heuristics, in this case between an implicit and an explicit fluency heuristic. It should be clear that the epistemic feeling kind of subjective fluency denotes the application of an *implicit* fluency heuristic: The products of an unconscious heuristic inference based on fluency are immediate experiences with a specific meaning, i.e. epistemic feelings (Unkelbach and Greifeneder 2013b).

Yet, there is room for another kind of “subjective fluency” based on an application of the explicit fluency heuristic. For that one needs an idea of how fluency could look like on a personal level *without* being an epistemic feeling. So far we have only been talking about an *unconscious* process property of speed or about *conscious* (fluency-based) epistemic feelings. But what is *just* fluency on a personal level? I suppose something like that: consciously, we do sometimes observe that e.g. a response to a question or a request comes to mind immediately or quickly (as we do sometimes observe that it comes to mind sluggishly or not at all). This looks like a kind of conscious “introspective (subjective) fluency”, which is the introspective observation that some content comes to mind quickly (that, in turn, might be based on objective fluency).

Based on introspective fluency (or something similar) we might sometimes consciously apply an (explicit) fluency heuristic and infer that the reason why something came to mind quickly is because of the presence of some distal property (perhaps instantiated by the content coming to mind). These kinds of inferences are made based on explicit “naïve theories” people happen to hold about what introspective fluency means (Schwarz 2004). It is conceivable that the application of these naïve theories becomes at some point automatised and might “go implicit” so that the associated inferences are henceforth (sometimes) carried out unconsciously (cf. Mandik 2006). That is, a specific explicit (fluency or other) heuristic becomes implicit. As a consequence, instead of making conscious inferences on the basis of observed states of affairs, epistemic feelings might start to emerge (cf. Carruthers 2009, see also section 6.4).

To illustrate the difference between implicit and explicit heuristic inferences in this context: we might, for instance, have an immediate feeling of knowing based on accessibility or cue familiarity, or we might recall some related fact to a question very quickly and based on this observed quickness of recall (i.e. fluency) consciously infer that we know the answer to the question. The latter is not a FOK but — analogously to the FOK — it might easily serve as the basis for a know-judgment, echoing the distinction between experience- and information-based judgment. In fact, there is no reason to preclude

that sometimes *both*, a FOK and the introspective observation of quick conscious recall of some related information (introspective fluency) *jointly* inform a know-judgment. In some cases, the two subjective fluencies might even interact with each other: on the one hand, a FOK might be a cue for how to interpret the introspective fluency and, on the other hand, introspective fluency might help the subject to recognise the feeling she is experiencing as a FOK. Recognizing this complicates the picture and introduces potential confounds for research on proper FOKs and other fluency-based epistemic feelings. At the same time, it paints a more adequate and nuanced picture of our psychological life.

Coming back to the FOK: In the context of situations where FOKs occur, processing fluency can be exhibited by perceptual recognition processes as mentioned for the cue familiarity heuristic as well as by the search and retrieval processes whose non-target products prominently figure in the accessibility heuristic. In the former case it makes sense to talk of perceptual fluency (since fluent perceptual processing is taken as a recognition signal) while in the latter case retrieval fluency is at work (Jacoby and Dallas 1981; Schwarz, Bless, et al. 1991). On the one hand, processing fluency can be a cue *independent* of target search and retrieval processes as in the case of perceptual fluency. On the other hand, it can be a cue *dependent* on the target search and retrieval processes where the search and retrieval of information can be fast.

In fact, a general processing fluency account fully accommodates the cue familiarity account and supplements the accessibility account: First, that a cue is familiar is signalled by perceptual fluency which serves as an early recognition signal, a proxy for the distal (and directly inaccessible) property of actual acquaintance with the cue due to prior encounter. In other words, the recognition mechanism at work in cue familiarity is the detection of perceptual fluency in processing the cue.

Second, in the accessibility account of the FOK, one determining factor is the amount of non-target information retrieved in a failed attempt to come up with the target. Now processing fluency brings another factor to the table: the retrieval processes implicated in mustering the non-target information can be more or less fluent (Koriat 1993). This retrieval fluency contributes to the generation and strength of a FOK as well. In fact, retrieval fluency in this context is plausibly construed as a signal of accessibility or availability of information that has not yet been retrieved or accessed, including the target itself. Some information is actually retrieved or accessed, some information is *retrievable* or *accessible* but not retrieved or accessed and, finally, some information is somewhere in

our memory but not accessible. But how do we know that some information *is* accessible without accessing it? Based on prior reflections four possibilities come to mind: The accessibility of information is inferred 1) based on retrieval-independent direct access to the sub-threshold activation (assuming that supra-threshold activation equals retrieval) of this information, 2) based on perceptual fluency with the target pointer (cue familiarity), 3) based on the amount of accessed information, or 4) based on retrieval fluency. Now, the FOK signals that the target is accessible. That this FOK is based on direct access to the sub-threshold activated target information appears only as a theoretical possibility.¹⁴⁹ Empirically it seems that the FOK is based on 1) cue familiarity, 2) the amount of accessed non-target information, and 3) on retrieval fluency. All these properties (*including* direct access to sub-threshold activated information) can be taken as proxies for likely future retrieval or accessibility of the target.

5.3.1.6. FOK's Cousins: Feelings of Information Availability

Note that all these features are not *per se* proxies for the accessibility of the specific target but for the accessibility or availability of not yet accessed information *in general* that is constrained by the prompting stimulus. Thus, in using these proxies, the FOK is a wager that the latter contains the former. The mentioned proxies are thus not only proxies for specific “knowledge” but also for the sheer accessibility or availability of (relevant) information. Thus, it is conceivable that in other contexts some of these cues are used to generate feelings with similar but ultimately distinct meanings than “knowledge” or possession of a *specific* piece of information.

I suggest that there might be something like a “feeling of information availability” (FOA), the simple feeling that more information is available for access, as well as its negative twin in the form a “feeling of information unavailability” (FOU). . What distinguishes the FOK from the FOA is that it concerns a specific target for retrieval and marks an active attempt to retrieve *this* target. In contrast, the FOA concerns a wider net of information (which might e.g. take the form of stimulus-specific mental files, Recanati 2012, 2016) and does not necessarily involve an ongoing attempt to retrieve it (rather it only registers its presence or availability). At the same time, similar to the FOK, the

¹⁴⁹A corollary theoretical possibility involving direct access is that the sub-threshold activation of *non-target* information feeds into the amount of accessed information that is used to estimate future target retrieval according to the accessibility heuristic.

FOA results from the products (amount of information) and properties (fluency) of a spontaneous retrieval triggered by a stimulus.¹⁵⁰

This retrieval will sometimes produce a FOU instead of a FOA if the spontaneous retrieval attempt does produce only a little or no information and/or is disfluent (cf. Schwarz, Bless, et al. 1991). The FOU seems similar to (but more general than) what has been described as the “blank-in-the-mind experience”, “the awareness that one cannot retrieve the goal that initiated a course of action (why did I do this?) or the goal following a course of action that has been completed (what should I do now?)” (Efklides 2014, p. 233). The FOA is very similar to what Mangan has termed “the feeling of immanence” which is “the feeling that much more detailed information is available on the periphery for retrieval if needed” (Mangan 2001, p. 3).¹⁵¹

The FOA is also related to the “feeling of imminence” that is (sometimes) a phenomenal component of the TOT, describing the “sense of imminence to recall” or “feeling of being on the verge of recall” (Schwartz, Travis, et al. 2000, p. 19; see also Smith 1994). I would add that the feeling of imminence does not have to concern only information that is remembered but can also concern information that is newly constructed but not yet consciously accessed.¹⁵² The more generalised version of the feeling of imminence describes then the feeling of being on the verge of accessing some information. As such, the feeling of imminence seems to be located in between the FOK or the FOA and actual information access.

5.3.1.7. Main Exhibit 4: Multiple Realizability

Coming back to the FOK it is important to emphasise another aspect that recent research has revealed: “FOKs are not governed by one underlying mechanism. Rather, a host of [...] mechanisms seems to work in concert to produce the single subjective state” (Thomas, Lee, et al. 2016, pp. 88-89). That is, the relation between the various described

¹⁵⁰In contrast to the case of the FOK, this stimulus does not act as a pointer or cue that targets a *specific* bit of information to be retrieved. So while FOKs are typically triggered by questions, FOAs can also be triggered by physical objects, people, propositions etc.

¹⁵¹I presume that with “periphery” Mangan means sensory periphery. Thus, the feeling of immanence seems to be specifically related to sensory states. By being constrained to the sensory domain the feeling of immanence is more specific than the FOA which can refer to accessible information in the sensory periphery but also in memory etc.

¹⁵²In fact, we seem also to have something like a feeling of imminence for worldly events, such as when you get a feeling that your favourite soccer player is about to score a goal or that the person you’re flirting with is about to kiss you. Thanks to Pablo Fernandez Velasco for pointing that out.

cues is not mutually exclusive but complementary. Cue familiarity, accessibility and fluency work in tandem to generate the unitary subjective state of FOK. In other words, the FOK is *multiply determined*. In fact, sometimes different cues play out at different temporal stages of the FOK. For instance, Asher Koriat and Ravit Levy-Sadot point out that:

Both cue familiarity and accessibility [...] contribute asynchronously to FOK, but whereas the effects of familiarity occur early, those of accessibility occur later and only when cue familiarity is sufficiently high to drive the interrogation of memory for potential answers. (Koriat and Levy-Sadot 2001, p. 34)

5.3.1.8. FOK's Brother: TOT

Similar things can be said about the TOT, the feeling that one is in possession of a target information but is currently unable to produce it.

Bennett Schwartz and Janet Metcalfe summarise the current perspective on the TOT which is largely analogous to the one on FOK:

[T]he heuristic-metacognitive account [...] hypothesizes that retrieved semantic, syntactic, and phonological information combined with other related information converges to inform a metacognitive monitor that retrieval of the word is likely, stimulating the TOT phenomenology. [...] The monitor examines the amount of related and partial information recalled, the familiarity of the cue, and even the recent history of retrieving the particular word [...] That is, it is not sub-threshold activation of the actual word that triggers the TOT, but a host of accessible cues and clues, which may include partial activation or whole activation of a bit of the target, or even the whole target itself, if the person is not sure that the answer is correct (Schwartz and Metcalfe 2014, pp. 19-20).

The metacognitive monitor in this case can be understood as “a simple accumulator of feature matches and mismatches” or of “retrieved information from all sources”, producing “a TOT response at a certain criterion” (ibid., p. 21). These are essentially the same two factors as encountered in the accessibility model of the FOK. There the sheer amount of retrieved information (as the accumulated information from all sources for

the TOT) possibly together with the amount of actually relevant information (as the accumulated amount of matches for the TOT) contribute to the character of the FOK.

5.3.1.9. Section Recap

To sum up: the FOK was previously assumed to be the result of an independent memory monitor with direct access to memory traces (section 5.3.1.1). This view has been largely superseded in favour of a heuristic mechanism that operates over cues correlating with the actual accessibility of a target information in memory (section 5.3.1.2, but see the end of the section). These cues are the familiarity with elements of the information query (cue familiarity, section 5.3.1.3), the amount of retrieved non-target information (accessibility, section 5.3.1.4), as well as the speed with which the implicated perceptual and memory processes operate (processing fluency, section 5.3.1.5). Furthermore, these cues are not mutually exclusive but complimentary when it comes to bringing about the FOK, making it multiply determined (section 5.3.1.7). In the course of this section we have also made our acquaintance with a functional analogue of FOK in the form of introspective fluency (the introspective observation of some content or response coming quickly to mind, see section 5.3.1.5) and with experiential relatives of the FOK: the FOA and FOU, i.e. the feeling that more information is (un)available for access (section 5.3.1.6), and the TOT (the feeling that one is in possession of a relevant information but is currently unable to produce it, section 5.3.1.8) whose mechanisms can be understood largely analogously to FOK-mechanisms.

5.3.2. The FOF Machinery

These reflections on the FOK and TOT foreshadow the mechanisms behind the FOF. The dominant account of the FOF is largely analogous to the processing fluency account of the FOK, proposing perceptual fluency as an initial recognition signal. The idea is that a previous encounter with a stimulus will facilitate its perceptual processing upon a renewed encounter, resulting in perceptual fluency. This fluency, in turn, will be attributed to the stimulus being familiar, triggering a FOF.

The ability to perceive and process something quickly and easily is evidence that the cognitive system has processed the stimulus in the past. It is the attribution of processing fluency to prior exposure that gives rise to familiarity (Verde et al. 2010, p. 142).

5.3.2.1. Main Exhibit 5: Discrepant Processing Fluency

One interesting extension that work on the FOF has brought to the fore is that fluency cannot only pertain to different processes but that the fluency pertaining to a specific process can take different forms. Specifically, it has become apparent that the effects of objective fluency, as well as the phenomenology of fluency-based epistemic feelings, tend to be more pronounced and salient if a process is not only fluent but *discrepantly* fluent, i.e. if a process is not “just fast” (absolute fluency) but fast relative to something else (relative fluency) (Whittlesea and Leboe 2003, see also section 5.2).¹⁵³ A process can be discrepantly fluent in several complementary ways: It can be discrepantly fluent relative to 1) the rest of the currently ongoing processing background, i.e. a specific process can be fast and therefore fluent relative to *other* ongoing processes, or 2) relative to an expectation concerning the speed of the specific processing, whereas this expectation is formed 2.1) outside or 2.2) within the present context (Whittlesea and Leboe 2003; Hansen and Wänke 2013; Unkelbach and Greifeneder 2013b; Garcia-Marques, Silva, Mello, and Hansen 2019).¹⁵⁴ In the special case of FOF it has been argued that the FOF is most pronounced when the perceptual processing of a stimulus is fast in the sense of being faster than expected for comparable stimuli. That is, the processing is fast relative to a stored standard or expectation for how much time perceptual processing usually takes for stimuli of a similar kind. In relation to strong FOFs Whittlesea and Williams explain:

[S]trong feelings of familiarity are seldom elicited by faces one knows well (e.g., spouse, friends, etc.) or faces one does not know at all (e.g., strangers in a crowd). Instead, strong feelings of familiarity seem to occur when one has limited knowledge, for example, having met the person only a few times (true familiarity) or when a complete stranger resembles someone one knows (false familiarity). We suspected that these outcomes occur because people expect known faces to be fluently processed and nonfluently processed faces to be unknown; that is, they expect a match between the fluency of perceiving and the coming-to-mind of identity information. Fluent processing

¹⁵³The “absolute” in “absolute fluency” needs to be qualified though: an absolutely fluent process is hypothesised to be behind a response time that is (by some margin) below the average response time in a task (Whittlesea and Leboe 2003, p. 63). Insofar, “absolute fluency” is in a way clearly relative, namely relative to the average of the response latencies in a task.

¹⁵⁴A special case of 2) is when the speed of one and the same process is discrepant between differing stimulus or item categories, e.g. between low-frequency vs. high-frequency words (Jacoby and Dallas 1981). Assuming that the processes that take in low- and high-frequency words are the same, a high-frequency word encountered amidst low-frequency words will be processed relatively faster.

of a friend’s face meets this expectation; so does nonfluent processing of a stranger’s face. However, limited prior experience, or similarity to a friend, may augment the fluency of processing a current face without enabling retrieval of identity for that face. In that case, the fluency is surprising. We suspected that in such cases, people unconsciously attribute the unexplained fluency to a prior experience and consciously experience a feeling of familiarity. By this account, the basis of the feeling of familiarity is the perception of a discrepancy between the fluency of processing and failure to produce the person’s identity, rather than the fluency per se. (Whittlesea and Williams 2000, p. 548)

I suggest the following reconstruction of this line: an unexpectedly fluently processed stimulus paired with the unavailability of stimulus-related information leads to a FOF and, additionally, a FOU, that frames the FOF as salient and in need of explanation. In other words, perceptual fluency coupled with a poor (disfluent) recall of stimulus-related information (low amount) leads to a strong FOF which is rendered phenomenally salient by a (component) FOU.

Note that FOF research has been so far focusing on familiarity with perceptible material objects. As a consequence, the specific familiarity in question is in most cases *perceptual* familiarity. As I have noted in section 5.2, however, it seems plausible that also non-perceptible entities can be familiar, such as abstract objects, ideas or mental states, resulting in something like *intellectual*, *conceptual* or *introspective* familiarity.¹⁵⁵ As noted in section 5.2, FOFs are transmodal and can unite with (or be absent from) all kinds of conscious mental states.

My tentative suggestion is that the FOF can be more generally conceptualised as a form

¹⁵⁵In fact, the case of introspective familiarity is interesting in that something similar might be at the heart of what has been called the “feeling of pastness” (FOP) that accompanies episodic memory (e.g. James 1890, p. 605; Russell 1921, pp. 161 sq.; Klein 2013; Klein 2014; Dokic 2014a). The FOP phenomenally demarcates a specific content as being from the autobiographical past. James remarked that “we have a constant feeling *sui generis* of pastness, to which every one of our experiences in turn falls a prey” (James 1890, p. 605) and Russell noted its kinship with the feeling of familiarity:

There may be a specific feeling which could be called the feeling of “pastness,” especially where immediate memory is concerned. [...] [I]mages are regarded by us as more or less accurate copies of past occurrences because they come to us with two sorts of feelings: (1) Those that may be called feelings of familiarity; (2) those that may be collected together as feelings giving a sense of pastness. The first lead us to trust our memories, the second to assign places to them in the time-order. (Russell 1921, pp. 161 sq.)

I suggest that we can develop accounts of the feeling of pastness inspired by the FOK and FOF (which, in contrast, seem to be specific to semantic instead of episodic memory).

of *phenomenal* familiarity or acquaintance, including perceptual, intellectual, conceptual and introspective forms. And here (discrepant) processing fluency, taking various process-specific forms, emerges as an ideal determinant of these various kinds of phenomenal acquaintances: there are fluencies of the perceptual, conceptual, encoding, retrieval etc. kind. It is natural to assume that, in some contexts, conceptual, retrieval or encoding fluency will lead to a FOF with an idea. On second thought, one of the stimuli classes frequently utilised in FOF research bears actually some resemblance with immaterial objects: words (e.g. Whittlesea 1993). Plausibly it is not only perceptual fluency that can be triggered by words but also fluency of linguistic processing, resulting in conceptual fluency and familiarity (see also sections 6.3 and 8.2).

5.3.2.2. Exhibit 6: Matching Memory or Accessibility revisited

Notwithstanding the dominant (discrepant) processing fluency account of the FOF, there is a little-noticed but plausible alternative suggestion for explaining the FOF. Interestingly, this so-called mnemonic account bears significant resemblance to the accessibility account of the FOK:

[A]n item feels familiar because it matches a representation of a previous encounter. Of course, one would rarely expect such a match to be exact, because the surrounding context, the perceptual conditions, even integral parts of an item can change over time. Thus, it would make functional sense for familiarity to vary continuously with the degree of match. This allows familiarity to be informative, given the fluctuating environment, but it also means that familiarity is fundamentally ambiguous, because even novel items will bear some resemblance to things encountered in the past. [...] *[M]emory evidence* is a continuous dimension that represents some aggregate of the matches between a recognition probe and the memory images of recently studied items (Verde et al. 2010, p. 144).

Thus, encountering a stimulus prompts a spontaneous attempt to match the stimulus representation with memory content, i.e. to retrieve stimulus-related information (if any). If a sufficient amount of information is retrieved, a FOF is issued.

Note that matching memory for a presently encountered stimulus seems like a close psychological approximation of the distal property of actual familiarity. For how else could familiarity be psychologically realised and tracked? In other words, memory matches

are a proxy that is very closely related to the distal property of actual acquaintance. Thus, if the FOF is caused by memory matches, it would perhaps not be too far fetched to speak of a mechanism approximating *direct causation* or *direct access* (assuming a monitor accessing the matches and making an inference).

Be it as it may, as the researchers note themselves, the fluency and the mnemonic account are not mutually exclusive (Verde et al. 2010, p. 150). This makes sense: fluency *and* memory matches seem to be better evidence for actual familiarity than each of them alone. I thus suggest understanding the FOF analogously to the dynamically unfolding and multiply determined FOK: perceptual fluency gives rise to an early FOF and an ensuing retrieval attempt supplements available information about the stimulus (if any), possibly resulting in the recollection of some stimulus-related information and a FOA that “normalizes” the (expected) FOF. If the perceptual-fluency-prompted second retrieval stage fails, however, a FOU makes the FOF stand out and unexpected.

5.3.3. Epistemic Feelings and Metacognition

So much for the specific mechanisms for FOKs and FOFs. Now there is an overarching general framework that some researchers like to apply to certain epistemic feelings such as FOKs or feelings of confidence (FOC). We have found some of its overtones when discussing FOK-mechanisms. I am talking about the *metacognition* framework (Nelson and Narens 1994). It is due to this that FOKs and FOCs are often conceptualised as *metacognitive* feelings. Traditionally metacognition has been defined as “cognition about cognition” (Flavell 1979). Or using more precise terms: as cognition serving the acquisition of knowledge about (i.e. monitoring) and regulation of (i.e. control) cognition. Now, this idea has been variously fleshed out, understanding the terms involved (e.g. “cognition”) in different ways. Some understand metacognition relatively narrow as referring to metarepresentations, i.e. representations of one’s own representations (e.g. Carruthers 2009). Call this *metarepresentational metacognition*. Examples of metarepresentational metacognition are, for instance, beliefs about one’s desires. The belief is a representation that represents another representation, namely desire.

Others are more liberal and grant that there are also forms of metacognition that do not involve metarepresentation. Metacognitive feelings are claimed to be one such form of non-metarepresentational metacognition. They are assumed to be outputs of monitoring mechanisms that issue metacognitive feelings in order to control ongoing cognition directly and indirectly. As such, metacognitive feelings evaluate and regulate cognitive

processes and activities without explicitly metarepresenting them. Thus, metacognitive feelings are said to belong to *evaluative metacognition*: they evaluate our cognitive processes, activities and dispositions (e.g. Proust 2010, 2013). Now there is a debate about whether more permissive kinds of “metacognition” such as metacognitive feelings really qualify as *metacognitive* in any theoretically interesting way (e.g. Carruthers 2017a). What counts as metacognitive in a non-metarepresentational sense is an interesting and intricate question. However, concerning the present thesis, there is little at stake in this debate and so I want to leave it to others to decide.

What is more important in the present context is to point out what the talk of metacognition does *not* mean when it comes to epistemic feelings. First of all: not all epistemic feelings have to be metacognitive feelings, at least not in the same sense. Some epistemic feelings seem to indeed concern the subject’s mental constitution while others seem to concern things in the external world. For instance, while the FOK is usually construed as a metacognitive feeling, the FOF is not. The FOK is somehow about the “knowledge” or possession of information of a subject while the FOF is often about something that has been previously encountered in the external world. It seems intuitive to conceive of the former as metacognitive in some way while it is less straightforward for the latter. Of course, one could *also* find some sense of metacognitive to count the FOF in, such as: it is somehow about *previous* perception or cognitive processing; or: it is caused by properties of processes such as perceptual fluency, or: it motivates cognitive activity such as renewed recognition attempts or memory query, or: it might be used to make a metacognitive judgment. Such an elastic strategy towards the notion of metacognition might, however, trivialise it.

Second: Even if metacognitive feelings might be somehow about cognitive states and processes, metacognitive feelings do not have to be *explicitly* about these cognitive states or processes. As already noted: they are *not* metarepresentational.¹⁵⁶ They can, for instance, evaluate the states and processes in question by representing the first-order contents of these states and processes in a way that does not make an explicit reference to the carriers of these contents (i.e. the states and processes). In a similar vein, metacognitive feelings also do not have to be explicitly self-referential, i.e. consciously making reference to the subject who is the bearer of the evaluated states and processes (see also section 6.4 and 7.3.2).

¹⁵⁶On the other hand, it seems that sometimes epistemic feelings *can* be metarepresentational (for a more radical take on the prevalence of metarepresentation see Shea 2014). Consider a FOF with a specific mental state such as the “usual” emotional response to a loved one in the case of Capgras syndrome as proposed by de Sousa (see section 5.2).

However, metacognitive feelings might plausibly be all these things *implicitly*. Surely, from the ways metacognitive feelings evaluate, we can reconstruct implicit references to the evaluated and motivated cognitive processes as well as the subject and its mental constitution. In many ways, however, this does not seem to be highly specific to *metacognitive feelings*. Even perceptual experiences have (implicit) perspectival, relational and response-dependent aspects (cf. Deonna 2006; Cohen 2009, see also section 6.4).

Now, after we have acquired a pre-theoretical grasp on the phenomenology of epistemic feelings (in section 5.2) and after we have accumulated some knowledge about potential answers to the mechanism-question (in section 5.3), it is time to tackle the second question mentioned at the beginning of this chapter: What kind of state are epistemic feelings? This will be the topic of the next chapter.

To sum up: Some epistemic feelings have been proposed to be metacognitive in that they evaluate our cognitive processes, activities and dispositions. This kind of *evaluative* metacognition is said to contrast with *metarepresentational* metacognition referring to representations of one's own representations. It is important to keep in mind what the metacognitive framework does *not* imply for epistemic feelings. First, not all epistemic feelings have to be metacognitive feelings, at least not in the same sense. Some epistemic feelings seem to indeed concern the subject's mental constitution while others seem to concern things in the external world. Second, even if metacognitive feelings might be somehow about cognitive states and processes, metacognitive feelings do not have to be *explicitly* about these cognitive states or processes. As already noted: they are *not* metarepresentational. In a similar vein, metacognitive feelings also do not have to be explicitly self-referential, i.e. consciously making reference to the subject who is the bearer of the evaluated states and processes. On any count, metacognitive feelings might plausibly be all these things *implicitly*. Surely, from the ways metacognitive feelings evaluate, we can reconstruct implicit references to the evaluated and motivated cognitive processes as well as the subject and its mental constitution.

5.4. Chapter Recap

This chapter introduced a class of mental states that would appear as a rather comfortable home for intuition experiences: epistemic feelings. Epistemic feelings have been broadly described as “feelings that enter into the epistemic processes of inquiry, knowledge and metacognition” (de Sousa 2008, p. 189). I first gave the reader a pre-theoretical

grasp for some prominent instances of epistemic feelings such as the feeling of knowing (FOK) and the feeling of familiarity (FOF) (section 5.1 and 5.2).

Then I dived into the research of what brings about these feelings. In the course of discussing the FOK we learned about the nowadays unfashionable Direct Access model and then focused on the contemporary majority view: the Heuristic Inference model that posits an inferential process which produces FOKs on the basis of cues. These cues are the familiarity with elements of the information query (cue familiarity, section 5.3.1.3), the amount of retrieved non-target information (accessibility, section 5.3.1.4) as well as the speed with which the implicated perceptual and memory processes operate (processing fluency, section 5.3.1.5). Furthermore, these cues are not mutually exclusive but complementary when it comes to bringing about the FOK, making it multiply determined (section 5.3.1.7). In the course of this section we have also made our acquaintance with a functional analogue of FOK in the form of introspective fluency (the introspective observation of some content or response coming quickly to mind, see section 5.3.1.5) and with experiential relatives of the FOK: the feeling of information (un)availability (FOA/U), i.e. the feeling that more information is (un)available for access (section 5.3.1.6), and the TOT (the feeling that one is in possession of a relevant information but is currently unable to produce it, section 5.3.1.8) whose mechanisms can be understood largely analogously to FOK-mechanisms as being multiply determined by cues such as perceptual fluency and the accessibility of information related to a stimulus.

Concerning the FOF we saw mechanisms at work similar to the FOK: it appears to be multiply determined by cues such as perceptual fluency and the accessibility of information related to a stimulus. Furthermore, discussing the FOF enabled us to add discrepant processing fluency to our conceptual repertoire of mechanisms. Work on FOF has brought to the fore is that fluency cannot only pertain to different processes but that the fluency pertaining to a specific process can take different forms. Specifically, it has become apparent that the effects of objective fluency as well as the phenomenology of fluency-based epistemic feelings tend to be more pronounced and salient if a process is not only fluent but *discrepantly* fluent, i.e. if a process is not “just fast” (absolute fluency) but fast relative to something else (relative fluency).

We have seen that a process can be discrepantly fluent in several complementary ways: It can be discrepantly fluent relative to 1) the rest of the currently ongoing processing background, i.e. a specific process can be fast and therefore fluent relative to *other* ongoing processes, or 2) relative to an expectation concerning the speed of the specific

processing, whereas this expectation is formed 2.1) outside or 2.2) within the present context. In the special case of FOF it has been argued that the FOF is most pronounced when the perceptual processing of a stimulus is fast in the sense of being faster than expected for comparable stimuli. That is, the processing is fast relative to a stored standard or expectation for how much time perceptual processing usually takes for stimuli of a similar kind.

Finally, in section 5.3.3 we touched upon an influential *general* framework which understands (some) epistemic feelings as a form of metacognition. We learned that some epistemic feelings have been proposed to be metacognitive in that they evaluate our cognitive processes, activities and dispositions. This kind of *evaluative* metacognition contrasts with *metarepresentational* metacognition referring to representations of one's own representations. The emphasis of this section was to point out what the metacognitive framework does *not* imply for epistemic feelings. First, not all epistemic feelings have to be metacognitive feelings, at least not in the same sense. Some epistemic feelings seem to indeed concern the subject's mental constitution while others seem to concern things in the external world. Second, even if metacognitive feelings might be somehow about cognitive states and processes, metacognitive feelings do not have to be *explicitly* about these cognitive states or processes. As already noted: they are *not* metarepresentational. In a similar vein, metacognitive feelings also do not have to be explicitly self-referential, i.e. consciously making reference to the subject who is the bearer of the evaluated states and processes. On any count, metacognitive feelings might plausibly be all these things *implicitly*. Surely, from the ways metacognitive feelings evaluate, we can reconstruct implicit references to the evaluated and motivated cognitive processes as well as the subject and its mental constitution.

Now, after we have acquired a pre-theoretical grasp on the phenomenology of epistemic feelings (in section 5.2) and after we have accumulated some knowledge about potential answers to the mechanism-question (in section 5.3), it is time to tackle the second question mentioned at the beginning of this chapter: What kind of state are epistemic feelings? This will be the topic of the next chapter.

Epistemic Feelings are Affective Experiences

6.1. Introduction: Reinforcing a Case

After we've got a grasp on epistemic feelings and the mechanisms behind them, what about the kind-question? What kind of states are epistemic feelings? Interestingly, epistemic feelings seem to be poorly identified with beliefs or perceptions essentially for the same reasons as intuition experiences are not beliefs or perceptions. You might largely reread chapter 3, substituting, say, “feeling of familiarity” or “feeling of knowing” for “intuition” and most of the points, I presume, would still hold. I think this is some evidence that this thesis might be on the right track when identifying intuitions with epistemic feelings.

Still: what are epistemic feelings? Surprisingly, most researchers seem to agree on this point: they are experiences. So much is consensus. A bit less consensual — but still surprisingly consensual — is the idea that they are *affective* experiences (e.g. Prinz 2011; Dokic 2012; Arango-Muñoz 2014; Proust 2015; Carruthers 2017b). The grounds on which they do so does not always seem to be bulletproof, however. Some just assume that epistemic feelings are affective. Others employ an “affective by association” strategy by simply grouping epistemic feelings together with more established affective experiences. Yet others rely on some individual empirical finding or theoretical consideration that taken by itself appears rather inconclusive. Perhaps it is just *obvious* that

epistemic feelings are affective. I don't think it is *that obvious*—but I am sympathetic with the conclusion anyway. What I want to do in this chapter is simply to bring the evidence together and reinforce the pre-existing case that epistemic feelings are affective experiences.¹⁵⁷ If this is so, then my claim that intuitions experiences are epistemic feelings establishes intuitions as a fairly well-established kind of state: affective experiences. In other words, we would find ourselves in a Samuel-/Donald-situation.

In what follows I will establish my case by drawing on the features of affective experiences described in chapter 4 and looking for empirical (section 6.2 and 6.3), phenomenological and theoretical observations (section 6.4 and 6.5) that point towards epistemic feelings having these features. I will also outline how understanding epistemic feelings as affective experiences can enhance our understanding of epistemic feelings (section 6.4) and how the phenomenology and intentionality of epistemic feelings can be analysed with the conceptual resources we have established for affective experiences (section 6.5).

6.2. Affective Markers of Epistemic Feelings

As discussed in chapter 3, the hallmark of affective experiences is phenomenal valence. Importantly, valence is something that is distinctive of *affective* experiences and not shared by other mental states such as perceptual experiences or beliefs. Thus, if it can be shown that epistemic feelings have phenomenal valence, then this can be taken as evidence that they are affective experiences and *not* other kinds of states such as perceptual experiences or beliefs.

Now, as in general with phenomenal qualities, valence cannot be measured directly. One thus has to rely on indirect evidence by measuring observable variables assumed to be associated with valence. Luckily, a number of valence-associated variables have been identified in the form of psychophysiological and behavioural responses, including subjective ratings and reports (Mauss and Robinson 2009). One important fact about valence is

¹⁵⁷Here is some initial support from another angle: My claim is that intuitions are epistemic feelings. And you might be sympathetic with the idea after you got a grasp on these states. *Independently* of what kinds of states epistemic feelings are, they might indeed appear to you as quite similar to intuitions. At the same time, remember the reasonable fit between the features of intuitions and affective experiences that I outlined in section 4.5. So if it is not implausible that intuition experiences are affective experiences and if it is, independently, plausible that intuition experiences are epistemic feelings, then it seems do give additional plausibility to that epistemic feelings are affective experiences.

that it is commonly assumed to be bodily realised (see e.g. Craig 2009; Damasio and Carvalho 2013). Now, epistemic feelings co-vary with interoceptive changes as well as with specific activity of facial muscles indicative of positive or negative affect (Winkielman, Schwarz, et al. 2003; von Helversen et al. 2008; Topolinski and Strack 2009b, experiment 4**topolinski2009d**; Fiacconi, Peter, et al. 2016; Fiacconi, Kouptsova, et al. 2017; Forster et al. 2016).¹⁵⁸ Additionally, higher interoceptive awareness (i.e. the ability to pick up on one's interoceptive states such as heartbeat) in subjects predicts a stronger association between reports on epistemic feelings and the presence of interoceptive cues (Fiacconi, Kouptsova, et al. 2017).

Another widely accepted bodily proxy of the presence of affective experience is the skin conductance response (SCR) which is associated with another characteristic feature of affective states: arousal. Capgras patients display a similar SCR to familiar and unfamiliar faces indicating, among other things, the absence of a FOF on whose basis they could discriminate between familiar and unfamiliar stimuli (e.g. Ellis, Young, et al. 1997; Ellis and Lewis 2001). The patients recognise their loved one *visually* but the usual *affective* response ordinarily elicited by the sight of the individual in question (*inter alia* a FOF) is missing (Pacherie 2010).¹⁵⁹ This is hypothesised to be the consequence of a disrupted connection between the ventral visual pathway that is responsible for the overt visual component in recognition and regions (notably the amygdala) that are responsible for the covert affective component in recognition (Breen et al. 2000). While the former seems to be intact in Capgras patients, the latter is lacking as is indicated by measurements of unusually low skin SCR vis-à-vis loved ones.

Capgras delusion raises many interesting issues concerning the interaction between cognitive processes [...] and accompanying emotional responses, not least the basis for our sense of familiarity when encountering someone we know well, which can [...] involve an automatic concurrent 'glow'. The same may well be true in the parallel domain of object recognition: those items with which we are particularly familiar or to which we are in some way attached may enjoy special cognitive status (Ellis and Lewis 2001, pp. 154 sq.).

¹⁵⁸The interoceptive changes in question are changes in heart muscle activity. Note that these changes can also be understood as relating to arousal rather than valence, see further below. On the other hand, the facial muscle activity is a sure sign for valence: The facial muscles in question are the smiling muscle, zygomaticus major, for positive affect and the frowning muscle, corrugator supercilii, for negative affect.

¹⁵⁹Alternatively, there might be a pronounced alienating feeling of *unfamiliarity* (e.g. Bayne and Pacherie 2004, p. 4; Ratcliffe 2007b).

In healthy subjects, the occurrence of FOFs co-varies with a discriminatory SCR for familiar and unfamiliar stimuli (Morris et al. 2008).¹⁶⁰

Epistemic feelings have not only been shown to vary with SCR but also to lead to increased liking, a behavioural measure of positive affect (Reber, Winkielman, et al. 1998, experiment 3; Winkielman, Schwarz, et al. 2003; Forster et al. 2013; Trippas et al. 2016, experiment 1), additionally co-varying with ratings of apparently non-affective properties such as truth (see e.g. Reber and Schwarz 1999; Unkelbach 2007, experiment 1; Unkelbach, Bayer, et al. 2011a, experiment 1; Hansen, Dechêne, et al. 2008, see also section 7.3.4), perceptual semantic coherence (Bowers et al. 1990; Bolte and Goschke 2008; Volz and von Cramon 2006; Volz, RübSamen, et al. 2008; Topolinski and Strack 2009b, experiment 10; Horr et al. 2014), linguistic semantic coherence (Bolte and Goschke 2005; Topolinski and Strack 2009a, experiment 2, 2009c, experiment 4, 2009b, experiment 1-9), grammaticality (Topolinski and Strack 2009b, experiment 11) and logical validity (Morsanyi and Handley 2012; Trippas et al. 2016).

Interestingly, in the second and third experiment, Trippas et al. 2016 also find that epistemic feelings that are presumably triggered by logical validity lead to higher ratings of physical brightness as well. This appears puzzling at first but is actually construed as a plausible consequence of the affective nature of epistemic feelings: affective experiences have been shown to facilitate content-encoding, presumably in part via the influence they exert on attention (e.g. Schupp et al. 2003; Kensinger and Corkin 2004; Sharot et al. 2004; Phelps et al. 2006; Pessoa 2011). That is, if paired with an affective experience, representational content appears as more *vivid* or to use a theoretically better-established term: *salient*.

Furthermore, positive epistemic feelings can serve as affective primes (Topolinski and Strack 2009c, experiment 2). In affective priming subjects evaluate a target stimulus with an affective connotation, say, the word “snake” as positive or negative after being presented with an affectively-laden prime stimulus, say, the word “poison”.

If prime and target have the same affective connotation, the evaluation that is associated with the prime word typically facilitates the processing of affectively congruent target words [...] and thus leads to faster evaluations of the target [...]. [T]his should be understood as an *indirect affective priming*,

¹⁶⁰This reminds of the dissociable sensory and affective component of pain, only for recognition (see section 4.2).

because it is not the affective value of the triad constituents per se that is expected to prime subsequent evaluations but the [...] valence that results from the processing of the triad. Specifically, we assumed that the fluency-induced increase in [...] valence in processing a coherent triad should facilitate subsequent processing of positively evaluated concepts and hamper the processing of negatively evaluated concepts. (Topolinski and Strack 2009c, pp. 1474 sq.)

In the experiment in question previously presented coherent word triads were shown to lead to positive affective priming effects via triggering epistemic feelings. For incoherent word triads, no such effect occurred.

It appears furthermore that the induction of positive or negative mood enhances or impairs performance in tasks capitalizing on epistemic feelings (Baumann and Kuhl 2002; Bolte, Goschke, and Kuhl 2003; Balas et al. 2012; Sweklej et al. 2014; Remmers and Zander 2018). This point is further reinforced by the finding that said performance is impaired in major depression (Greifeneder and Bless 2008; Remmers, Topolinski, Dietrich, et al. 2015; Remmers, Topolinski, Buxton, et al. 2017; see also Goldie 2012).

In general, it has been broadly pointed out that epistemic and emotional feelings influence judgments in similar ways and that the reliance on these feelings in making judgments is mediated by similar factors (Schwarz and Clore 2007; Greifeneder, Bless, and Pham 2011). As Greifeneder and colleagues summarize:

The review revealed that moderators of the reliance on affective and cognitive feelings are remarkably similar and can be grouped into five major categories: (a) the salience of the feelings, (b) the representativeness of the feelings for the target, (c) the relevance of the feelings to the judgment, (d) the evaluative malleability of the judgment, and (e) the level of processing intensity. Based on the reviewed evidence, it is concluded that the use of feelings as information is a frequent event and a generally sensible judgmental strategy rather than a constant source of error. (Greifeneder, Bless, and Pham 2011, p. 1)

A natural way to explain these findings is to parsimoniously assume that they do so in virtue of both being *affective* experiences. Therefore, we do not need a strict division between affective feelings and epistemic feelings as upheld in the cited reviews because *both* are affective.

To sum up: Epistemic feelings covary with interoceptive changes, variations in SCR and facial muscle activity, all well-established bodily proxies for the affective properties of valence and arousal. Furthermore, positive epistemic feelings lead to increased liking and can serve as positive affective primes—behavioural proxies for the presence of valence. Additionally, performance capitalizing on positive epistemic feelings are enhanced in a positive mood. On the negative side, performance that relies on positive epistemic feelings is impaired under conditions of negative mood and depression. Further behavioural evidence comes from studies of the influences of emotional and epistemic feelings on judgments. They bring to light that both kinds of feelings influence judgments in a similar way and that similar factors moderate the reliance on them when making judgments.

Coming from the affective to the epistemic side: Induced affect co-varies with judgments of seemingly non-affective properties such as truth, grammaticality, logical validity and semantic coherence. I think this shows that epistemic feelings robustly co-vary with affective properties such as valence and arousal. This, in turn, makes a good case for the thesis that epistemic feelings *are* affective. Is what we observe in the experiments affective *experiences*, though? Perhaps the affective components in the studied cases are *unconscious* — non-phenomenal valence and physiological arousal —, giving out unconscious epistemic nudges (i.e. unconscious functional analogues of epistemic feelings) but not *phenomenally guiding* the behaviour of subjects (see section 4.3.1). This is well possible. I want to grant that this is what might sometimes happen in the observed cases. However, the next section will show that what at least *also* sometimes happens is that there are *conscious* epistemic feelings at work in these situations—in the form of affective experiences.

6.3. Misplacing Affect in Interesting Ways

In this section, I will first strengthen the case for the covariation between epistemic feelings and affective properties being not just a correlation but a constitution relationship (section 6.3.1). For that, I will present studies that observe false positives of epistemic properties on the basis of incidentally induced affect. That is, inducing non-diagnostic affect leads subjects to incorrectly judge that an epistemic property is present. The second part of this section will be dedicated to studies go the other way around: they make the subject believe that the affect they experience in a given epistemic task is *not* diagnostic for the presence of an epistemic property (section 6.3.2). This turns out to

strip the subject of her ability to accurately detect the epistemic property, indicating that epistemic properties are at times detected on the basis of affect, and, since the affect can be misattributed, that the affect in question is conscious.

6.3.1. Constitutive Affect in Epistemic Feelings

Particularly instructive evidence for the affective nature of epistemic feelings comes from two kinds of misattribution studies: The first kind of studies generates a misattribution of the seemingly non-affective properties of familiarity and semantic coherence on the basis of induced positive or negative affect. In the familiarity studies, affect is induced either via the contraction of facial muscles or the presentation of reliably affect-eliciting stimuli such as attractive or emotional (e.g. smiling or frowning) faces or affect-laden words. Novel stimuli are rated as more familiar (or unfamiliar) if the smiling muscle, zygomaticus major, (frowning muscle, corrugator supercilii) is contracted (Phaf and Rotteveel 2005, experiment 2) or the stimuli in question are either attractive (Monin 2003), display emotions (Baudouin et al. 2000; Garcia-Marques, Mackie, et al. 2004, experiment 1; Lander and Metcalfe 2007) or are preceded by subliminal primes in the form of happy vs. neutral faces (Garcia-Marques, Mackie, et al. 2004, experiment 2; Duke et al. 2014) as well as happy vs. sad words (Phaf and Rotteveel 2005, experiment 1). In the semantic coherence studies, affect is either induced via the contraction of said facial muscles (Topolinski and Strack 2009b, experiment 4) or the subliminal presentation of happy and sad faces (ibid., experiment 5). As a consequence of the affect manipulation, items are more (less) often judged as semantically coherent. Crucially, Duke et al. 2014 and Topolinski and Strack 2009b explicitly demonstrate that the effect of induced affect closely mirrors the effects of processing fluency (and *actual* familiarity and semantic coherence) on familiarity and semantic coherence judgments.

To understand the importance of this finding, recall the construct of processing (dis)fluency introduced in section 5.3 (Alter and Oppenheimer 2009). Processing fluency is a process property that refers to the “ease”, understood as the speed, with which a given cognitive process is executed. Naturally, given that there are many kinds of cognitive processes, there are many kinds of fluencies: perceptual fluency, retrieval fluency, encoding fluency, answer fluency, conceptual fluency, to name a few. Now, manipulations of fluency have been proven to be a reliable way of inducing epistemic feelings. The present finding is remarkable because many researchers assume that processing fluency causes

all epistemic feelings (see e.g. Unkelbach and Greifeneder 2013a).¹⁶¹ Note that even if this assumption is correct, such a proposal only answers the question about the *proximal* cause of epistemic feelings but leaves open the question about the *distal* causes of epistemic feelings. Crucially, it also leaves open the question about the phenomenology of epistemic feelings, i.e. the phenomenology of what is caused by this fluency. When it comes to the phenomenal aspect, most researchers use non-descript fluency-centred labels such as “subjective experience of fluency” (Alter and Oppenheimer 2009), “subjective experience of ease” (Oppenheimer 2008) or “feeling of fluency” (Forster et al. 2013). Now, the studies by Duke et al. 2014 and (Topolinski and Strack 2009b) suggest that what ultimately matters *phenomenally* is not processing fluency *per se* but its seemingly multiply-realizable product: positive affect.¹⁶²

This implies that, in principle, it does not matter whether it is processing fluency or something else that causes this positive affect. It seems that given a specific context, say, a task relying on the detection of an epistemic property such as familiarity or coherence, epistemic feelings can be triggered by *whatever triggers affect*, correctly or incorrectly signalling the presence of said property. Fluency is only one of many possible antecedents.¹⁶³ Against this background, it appears likely that the phenomenology of “subjective experiences of fluency” and of epistemic feelings in general (i.e. also potentially those not caused by fluency) is constituted by transient, context-specific positive or negative affect. However, perhaps this is too quick. An alternative explanation of these results is that both, fluency and positive affect (whether fluency-induced or not), are used in an analogous way to inform familiarity and semantic coherence judgments. In other words: if there is fluency then the subject uses this as a cue to make a familiarity or semantic coherence judgment and if there is positive affect then the subject uses

¹⁶¹The role of accessible information for the FOK as described in section 5.3 does not straightforwardly square with this, however (see also chapter 6).

¹⁶²It might also be taken to falsify the claim that fluency is the proximal cause of all epistemic feelings (see also chapter 6). For that one would have to rule out that stimuli that are known to trigger positive and negative affect do not trigger this affect *via* triggering processing (dis)fluency. This idea appears at first glance implausible, especially if generalised to all affective states: we would have to postulate links between fluency and affective experiences such as fear, sadness or joy. However, there is accumulating evidence for an affect-fluency—and not only a fluency-affect—link in certain domains (Unkelbach, Fiedler, et al. 2008; Alves et al. 2015; Koch, Alves, et al. 2016; see also Verde et al. 2010). The underlying idea is that bits of positive information are more alike than bits of negative information. Thus, positive information can be processed faster (i.e. more fluently) than negative information. On the other hand, accumulating recent findings begin to cast some doubt on the encompassing causal role of fluency for epistemic feelings (Erle, Reber, et al. 2017; Flavell et al. 2018; Erle and Topolinski 2018).

¹⁶³Or, perhaps, fluency is best understood as one manifestation of a more general property that it shares with other affective states, for instance predictive success in a predictive processing framework (Van De Cruys 2017).

this as a cue to make the judgment.¹⁶⁴ So the “subjective experience of fluency” and “positive affect” are not the same but have similar functional roles when it comes to the judgment of some properties. This proposal is a real possibility.

6.3.2. Conscious Affect in Epistemic Feelings

The second kind of misattribution studies helps settle precisely this issue (Topolinski and Strack 2009c,a). While in the first kind of studies the subjects misattribute non-affective properties on the basis of affect, in the second kind of studies the misattribution goes the other way around: informative affective reactions are discounted by being misattributed to an irrelevant source. In these studies, the experimenters ask subjects to make semantic coherence judgments by discriminating between word triads that either share a common remote associate (e.g. SALT, DEEP, FOAM implying SEA; coherent triad) or not (e.g. DREAM, BALL, BOOK; incoherent triad). In the fluency-retribution condition, the subjects are told that the “easiness of reading and the fluency with which the meaning of words is recognized” (Topolinski and Strack 2009a, p. 614) is due to a task-irrelevant source: background music. In the affect-retribution condition, the subjects are told that the positive affect that might arise in the course of the task is due to the background music. The authors show that while misattributing fluency has no effect on performance, misattributing affect essentially strips subjects of the ability to detect the seemingly non-affective property of semantic coherence (above chance level). This seems to imply that the positive affect belonging to an epistemic feeling is prompted by the presence of a non-affective property: semantic coherence.

Importantly, the aim of the researchers was to find out what is *felt* in the task: the increased processing fluency triggered by processing semantically coherent items itself or the transient positive affect that (in this context) is triggered by the processing fluency. The authors conclude that their “finding strongly suggests that it is not the fluency that is used as internal cue in intuitive judgments of semantic coherence, but rather the fluency-triggered positive affect” (ibid., p. 615). This is a crucial finding in two respects.

First, this strengthens the initial case made on the basis of the findings in Duke et al. 2014 and Topolinski and Strack 2009b by suggesting that the phenomenology of epistemic feelings essentially consists in context-specific, transient positive or negative affect and

¹⁶⁴And if there is both, fluency and positive affect, then the subject might use both in a (weighted) additive way to inform judgment (Topolinski and Strack 2009b).

that there are probably no non-affective epistemic feelings or “subjective experiences of fluency”.¹⁶⁵ Fluency is not a cue available in experience to use for judgment. What is available is the result of fluency: positive affect. The researchers additionally back this conclusion with the finding that coherent triads are liked more than incoherent triads, but are *not* rated as more fluent in processing (Topolinski and Strack 2009a, experiment 1). Commenting on this work, Winkielman and colleagues and Chetverikov independently note that:

[T]his work shows that participants cannot report and re-attribute changing levels of fluency (facilitation due to semantic coherence) but are only aware of affective (hedonic) consequences of changed fluency. (Winkielman, Ziembowicz, et al. 2015, p. 2)

[I]t is affect but not fluency that influences subsequent judgements. Thus, it is unlikely that such a kind of non-affective feeling exists or has any functional significance. (Chetverikov 2014, p. 409)

Reinforcing and extending this point, Balas and colleagues find that altering the semantic coherence task to include word triads that themselves are neutral but have an affect-laden common remote associate has a characteristic impact on judgments of semantic coherence:¹⁶⁶ there is an increase in accuracy and speed for triads with positive associates relative to those with neutral and negative ones. On this basis, the authors argue that “fluency-based positive affect can be strengthened or weakened by affective responses induced through partial activation of an affectively valenced memory content (i.e., solutions to triads).” (Balas et al. 2012, p. 318) This, in turn, makes the crucial point (in line with Duke et al. 2014 and Topolinski and Strack 2009b) that “fluency of processing is not the only source of affective response that can influence intuitive judgements” (Balas et al. 2012, p. 312). Together these findings imply that seemingly non-affective epistemic properties are (sometimes) detected on the basis of *affective* epistemic feelings.¹⁶⁷

¹⁶⁵There are “subjective experiences of fluency” if the term is meant to emphasise the (contingent) aetiology of the experience but not the phenomenology.

¹⁶⁶An example for a positive/negative coherent triad is: COMPETITION, FINISH, ROUND implying MEDAL; CANDLES, NOVEMBER, STONE implying GRAVE.

¹⁶⁷As in the case for what we mean when we talk of “intuitions” (in philosophy), nothing precludes that the mentioned non-affective epistemic properties are also sometimes assessed via judgments that are based on something else than (epistemic) feelings:

[I]t is worth keeping in mind that in some cases judgments of familiarity, rightness, tip-of-tongue, and imminence, do not reflect underlying feelings, but are simply judgments. Such a non-phenomenological approach would suggest that these judgments are like [...]

This is shown by the fact that in specific contexts (e.g. cognitive tasks) positive or negative affect correctly or incorrectly signals the presence or absence of the task-relevant property, be it affective or non-affective.

Indeed, although it is common to ask reasoners to express answers to logic or probability problems as judgments of validity or probability, it is possible to measure affective responses to such stimuli. (Thompson 2014, p. 62)

Second and perhaps even more important: The valence in epistemic feelings needs to be conscious in order to make them *affective experiences*. However, I discussed in section 4.3.1 that valence does not need to be conscious in order to bias (epistemic) behaviour. That is, the epistemic behaviours observed in the experiments might *not* be the result of conscious epistemic feelings but of some unconscious action-biasing *valenced* states that are functionally analogous to epistemic feelings, say, “epistemic nudges”. That such epistemic nudges occur is, I think, very plausible.

However, we cannot explain the present experimental findings by relying on them. On the contrary, the mentioned studies demonstrate that there *is* a phenomenology, i.e. that the affect integral to epistemic feelings is conscious. This is because the subjects are able to *misattribute* the conscious affective signals that they would usually use to make conscious judgments. This contrasts with e.g. their inability to misattribute and use the *unconscious* processing fluency directly. Subjects cannot misattribute something that is unconscious since there is nothing to (correctly or incorrectly) attribute in the first place. The present finding, thus, rules out something that might seem like a possible explanation when one considers *unconscious* valence. Instead, what we observe in the experiments appears to be the result of *affective experiences*—epistemic feelings. It seems that epistemic feelings are plausibly part of what James bracketed out in his famous analysis of emotions:

That there are feelings of pleasure and displeasure, of interest and excitement, bound up with mental operations, but having no obvious bodily expression for their consequence, would, I suppose, be held true by most readers. [...] Certain sequences of ideas charm us as much as others tire us. It is a real

inferences from rudimentary changes in processing fluency. (Reber, Fazendeiro, et al. 2002, p. 11)

However, even in many cases like these it is plausible that “rudimentary changes in processing fluency” are often not directly observed in a way as we e.g. visually observe a car driving by fast. Rather, we become aware that some sub-personal process is fast indirectly because it’s speed leads to a positive evaluation in the form of positive affect.

intellectual delight to get a problem solved, and a real intellectual torment to have to leave it unfinished. The [...] set seem to depend on processes in the ideational centres exclusively. Taken together, they appear to prove that there are pleasures and pains inherent in certain forms of nerve-action as such, wherever that action occur. (James 1884, p. 189)

To sum up: In this section, I first strengthened the case for the covariation between epistemic feelings and affective properties being not just a correlation but a constitution relationship (section 6.3.1). For that, I presented studies that observe false positives of epistemic properties on the basis of incidentally induced affect. That is, inducing non-diagnostic affect leads subjects to incorrectly judge that an epistemic property is present. This speaks in favour of an affective constitution of epistemic feelings. Secondly, in section 6.3.2 I made the case that the affect in question is *conscious*. It thus not only causally biases epistemic behaviour but phenomenally constitutes epistemic feelings that provide conscious guidance for the subject's epistemic behaviour. To make this idea plausible I recounted studies where the following happens: the experimenters make the subject believe that the affect they experience in a given epistemic task is *not* diagnostic for the presence of an epistemic property. As a consequence, the subject loses her ability to accurately detect the epistemic property. This does not only indicate that epistemic properties are at times detected on the basis of affect but also that the affect in question is conscious. The fact that the affect can be misattributed when making a *conscious* judgment points towards the possibility that the affect is, in fact, consciously available to the subject. On the basis of the reviewed empirical findings I conclude that epistemic feelings are affective experiences. In the remainder of this chapter, I will here and there provide additional theoretical and phenomenological considerations that further support this conclusion.

6.4. Understanding Epistemic Feelings better as Affective Experiences

One of the advantages that come right out of understanding epistemic feelings as affective experiences is that we can apply the wide range of knowledge that we have about the latter to understand the former. Before I will apply the conceptual resources introduced in chapter 4 to epistemic feelings in section 6.5, let me lose a few words on the dominant perspective on epistemic feelings. This framework emphasizes that epistemic feelings are

the result of monitoring and control mechanisms over cognition (see section 5.3.3) and that processing fluency is their principal determinant (see section 5.3.1.5 and 6.3.1). Now I want to make a few points about how one can fruitfully integrate these perspectives with the idea that epistemic feelings are affective experiences.

Regarding the metacognitive aspect: if epistemic feelings are affective experiences, then, I contend, a more general (and perfectly compatible) perspective becomes available. This perspective thinks the monitoring and control function of epistemic feelings back from them being affective experiences. The monitoring and control of cognition is effected because epistemic feelings are *evaluative* (Proust 2010). Now, feelings *in general* are valenced experiences that *evaluate* their particular objects which can be all kinds of things. Affective experiences evaluatively represent their objects as having feeling-specific properties, their formal objects. In order to fulfil this representational job, they engage in a division of labour with other representational states (their base) whose content supplies the particular objects of the feeling. In virtue of their evaluation, affective experiences call for action and modification of behaviour—of a physical and mental, an obvious and subtle kind (e.g. Frijda 2008). In other words, feelings are motivational. Ultimately, the monitoring and control of mental or cognitive processes is just a special case of the evaluative and motivational function of affect operative in all kinds of domains, be they part of the body, the world or the mind (cf. Proust 2015; Fulkerson 2019).

In fact, against the background of affect, it becomes clearer why epistemic feelings will often *not* appear to be about one’s own cognition—even when, on closer examination, they are. This is not very different from other feelings that display what has been variously called the “irreflexive or non-reflective consciousness” (Frijda 1986, 2009) or “implicit egocentricity” Pacherie 2002 of emotional feelings and the “narcissism” of bodily feelings (de Vignemont 2018; cf. Akins 1996).¹⁶⁸

Emotional experience is ‘objective’, in the sense that it grasps and asserts objects with given properties. Irreflexive emotional experience also, by its very nature, is ‘projective’: The properties are out there. These properties contain the relationship to the subject: Emotional experience is perception of horrible objects, insupportable people, oppressive events. They contain the relation implicitly: the “to me” or “for me” dissolves into the property (Frijda 1986, p. 188).

¹⁶⁸These notions also reverberate in Proust’s conception of affective experiences as “affordance sensings” (Proust 2015; cf. Gibson 1979, p. 129; see also Griffiths and Scarantino 2005).

[W]hat I called implicit egocentricity, is that although the pragmatic properties the target object is represented as having are relative to the capacities of the agent, the relations between agent and target are not explicitly represented. These relational properties are collapsed into monadic, causally indexical, properties of the object. (Pacherie 2002, p. 77)

Interoceptive feelings aim at protecting the organism. They have [...] a narcissistic function: they aim at securing what is best for the organism. [...] [W]hat matters is not to track one's heartbeats in order to be able to count them; what matters is to determine whether the rhythm is neither too fast nor too slow. [...] [T]he function of interoception is to regulate the physiological balance of the organism. Therefore, interoception is not about the state of the body *simpliciter*; it is about the state of the body given the organism's needs and interests. One might also say that it is about the state of the body *for the self*. (de Vignemont 2018, pp. 268 sq.)

These observations jointly point towards the following idea: The properties represented by affective experiences have a relational, subject-relative component. This component can come in the form of an “unspecific” self-reference (perhaps something like an “essential indexical”, Perry 1979) or a more specific reference to the subject's physical or mental constitution. This constitution may comprise the subject's background desires and beliefs, which contribute to the intelligibility of the affective experience in question (cf. Helm 2001). Naturally, it can also comprise the subject's physical and mental dispositions in the form of capacities and abilities: is a subject in a better position to fight or to flee upon encountering a threat? If, say, the subject is in good shape and a comparison of the constitution of subject and threat tips in favour of the subject, anger might ensue—but if not, fear is more likely to be the response.

Now, at times we might be or — perhaps upon reflection — become aware of the subject-relative components that figure into our affective experiences. Most of the time, however, these components remain implicit and are not explicitly represented in the experience itself. The properties that the affective experiences represent their particular objects as having appear objective — “out there” — in the sense of not being subject- or mind-dependent (cf. Fulkerson 2019, pp. 5 sq.). This is also how Intellectual Affectivism can account for Koksvik's feature of objectivity describing the feature of intuition experiences “to be about objective, mind-independent, facts” (Koksvik 2017, p. 7, see section

3.4.2).¹⁶⁹

So when e.g. my joy represents something as being joyous, it is not explicitly representing it as being joyous *for me* or as being joyous *because* I have such and such background desires and beliefs or have some (mental) resources to spare. In a similar vein, when my FOF represents a stimulus as being familiar, it does not explicitly represent it as being familiar *for me* or as being familiar *because* I processed it faster than expected, have encountered it before or am able to recognise it if I try harder.

These reflections explain why “the boundary between noetic feelings and other kinds of feelings is not very sharp. Some feelings seem to lie at the borderline between noetic feelings and feelings about the external world” (Dokic 2012, p. 303). What might be seen as being specific to epistemic feelings is not their monitoring and control function but that they are concerned with a specific set of particular and formal objects such as cognitive processes and epistemic properties that often appear to be “out there” in the physical or even abstract realm. Then, again, such specificity is common to *all* feelings—this is what makes them distinct in identity and kind from each other.

Relatedly, Fulkerson argues for the idea of genuine “Affective Perception” that is the result of dedicated affective systems within individual sensory modalities such as vision and audition (Fulkerson 2019; see also Prinz 2014). This phenomenon in part explains the appeal of perceptual theories of emotions:

Emotional perception thus forms an important bridge between our sensory and our emotional lives. Given this connection, it is not surprising that perceptual accounts of the emotions have been so popular (and makes it all the more surprising that there has been less interest in emotional theories of perception). (Fulkerson 2019, p. 13)

¹⁶⁹Here comes a speculative line of thought: One might wonder whether, in fact, subject-relative components are not implicit to almost all mental states, including perceptual experiences (Cohen 2009; Zahavi and Kriegel 2015). Against this background, one might try to reflect about the gain to represent properties that rely on subjective-relative components as explicitly subjective. If nearly all mental states have such implicit subjective components, then the question arises: why are some of them explicitly represented but most of them left implicit? From the point of view of the limited mental economy of consciousness, it seems crucial to use this low bandwidth to represent only information that is directly conducive to the subject’s ability to cope with the world. To be sure, sometimes subject-relative information falls into this privileged category. This is why *sometimes* this information *is* consciously represented. Still, to assume that this information *always* or even often falls into this category seems to go too far. It is, thus, unsurprising that, at first glance, many of our mental states do not appear subject- or mind-dependent. Reflections such as these were among my reasons not to include Koksvik’s feature of objectivity in chapter 2 (see also footnote 204).

I'd like to add that what is meant by a dedicated affective system here is not that sensory modalities comprise some self-sufficient modality-specific mini version of an affective system. It rather consists in specialised modality-specific *pathways* that serve as (fast-track) bridges to the suite of “general” affective systems that produce e.g. valence (Carruthers 2017b). This is why various modality-specific and -unspecific affective states have still things in common in virtue of being affective.

I find it plausible to assume that there are also dedicated affective systems within individual *cognitive* modalities such as memory and reasoning. I find it furthermore plausible that these systems will produce cognitive states and affective states in tandem and that these states will appear related in various ways. Among these results, we will sometimes find epistemic feelings that take the cognitive states as their bases. At times the epistemic feeling and the cognitive state will still be recognizable as individual states. And at times these two states will be so closely tied so as to appear as a kind of genuine “Affective Cognition” (cf. Goldie 2000; Mitchell 2019).

Another point concerns the flexible inferential mechanisms that bring about epistemic feelings. Again, the deployment of such mechanisms does not seem to be specific to epistemic feelings. In general, we know that flexible inferential mechanisms are at work in producing affective experiences (e.g. Scherer et al. 2001; Russell 2003; Carruthers 2011; Barrett and Simmons 2015; Pober 2018). These processes are highly context- and subject-specific and lead to, well, context- and subject-sensitive affective experiences. As Fulkerson nicely puts it:

We are often presented with very different emotional qualities. The pleasantness of a warm blanket differs from the pleasantness of a cup of cocoa. According to the affective-motivational account, there is no single affective system involved. One involves a particular combination of softness and warmth that the other lacks, and the emotional elements are tied to the unique contributions of each modality. The mechanisms by which we represent these qualities and the ways in which the signals are presented to us seem to introduce novel qualia, what Mohan Matthen [...] calls ‘qualia insertion’—generating a new conscious signal for when a particular taste or smell is bad. This seems intuitively right: I dislike equally the tastes of olives and of liver, but the particular affective qualities assigned to olives differ markedly from those assigned to liver. And these motivational and affective reactions are highly personal, and sensitive to small variations in the context. (Fulkerson

2019, p. 12)

In fact, this heavy inferential work is not unique to affective experiences at all. It is also at work in perception and cognition (e.g. Jacobs 2002; Matthen 2010). However, there is an important difference between perception and affect, on the one hand, and some forms of cognition such as inferential thought or belief, on the other: In the former case the inferential work is *not apparent* from a conscious point of view. Rather the experiences appear as immediate and as immediately meaningful. Things do not look like scattered light gradients and edges, they look like bears and faces. Things do not (typically) just feel negative or positive, they feel dangerous or familiar (cf. Teroni 2018).¹⁷⁰ This way, affective experiences, as well as perceptual experiences, appear to us somewhat more “simple” than e.g. some forms of cognition. However, this appearance is *not* an indicator of simplicity. It is an indicator of heavy evolutionary pressures towards tuning and optimising the respective processes towards speed and efficacy. Some things are just too important to be left to sluggish, resource-hungry serial processing.

Now there is also something new we can say about processing fluency. Note that there is a puzzle about processing fluency effects: The fluency of seemingly quite different processes such as perceptual or conceptual ones leads to seemingly quite similar effects on various property judgments. This is acknowledged but seldom *explained* in the literature:

[A]lthough processing fluency takes many forms, we argue [...] that fluency exerts the same influence on judgments independently of how it is generated (Alter and Oppenheimer 2009, p. 2).

[C]ognitive processes such as recognition and recall are constantly monitored and continuously evaluated. This evaluation process results in an experience or a feeling when there is a discrepancy (i.e., variations in processing fluency) between the evaluation and expectancies concerning the process. However, the feeling resulting from this discrepancy is nonspecific, and the discrepancy triggers a search for an explanation. (Unkelbach 2007, p. 229; see also Whittlesea and Williams 2000; Whittlesea 2002)

For instance, Topolinski manipulated how subjects judged the semantic coherence of word triads by manipulating the triads’ semantic coherence *or* readability. However, processing fluency by itself is clearly process-specific. Moreover, it seems plausible that

¹⁷⁰The binding mechanisms at work in perceptual and affective experiences differ in that the former has a higher degree of integration. We do not sometimes experience a luminance outside of objects—but we do sometimes experience feelings whose meanings we cannot quite place (cf. Wiese 2018).

the equally process-specific monitoring mechanisms that operate on this fluency should have the fluency source information available to them. If so, then there seems to be no reason why the detection of fluency by a perceptual monitoring mechanism would lead to anything of consequence for semantic coherence judgments. Nevertheless, we observe that perceptual fluency leads to effects on semantic coherence judgments. In other words, process-specific fluency has effects that are in conflict with its being process-specific. The fluency appears to switch modalities as it were. How come?

I think we can resolve this puzzle with an affect-oriented perspective. Recall that the semantic coherence judgments can also be manipulated by directly inducing affect. The hypothesis is that what ultimately matters is not process-specific fluency but affect.¹⁷¹ We know about affect that it is, in contrast to fluency, transmodal and that it can be the result of many different things, one of them being process-specific fluency. As a consequence, misattribution of affect is a well-known phenomenon. There is something about affect that is at odds with a fully reliable tracking of its sources. In other words, there is a stage in the generation of affect where its information of origin is lost (cf. Gendler 2006). This likely has to do with its evaluative function (see also section 4.4.3):

Affect [...] doesn't get tied to the representations involved in the cognitive appraisals that produce it [...] On the contrary, affect from different sources tends to combine to form a single evaluation of whatever happens to be the object of attention, or to be the most relevant among objects of current attention. (Carruthers 2011, pp. 146 sq.)

Processing fluency leads to affect because it, as many other things, is valuable. After all, it is one expression of processing success. Focusing on evaluation more generally, what is crucial is not so much the specific value attached to one particular object but rather it is the value of one object in comparison to the value of other objects. Put differently, one function that affect plays is one of a “common currency”:

For by leaving affect unbound from the various properties that produce it, a simple computation of overall value is facilitated, irrespective of the size of the set of affectively-relevant properties. (ibid., pp. 149 sq.)

In other words, affect is “de-modalized” for reasons of value cross-comparison:

¹⁷¹In making this point I do not assume that this affect is necessarily conscious. It can come in the form of, say, the previously mentioned unconscious epistemic nudges. However, sometimes it will be conscious and come in the form of necessarily conscious epistemic feelings.

[V]alence signals [...] are thought to provide an evaluative “common currency” for use in affectively-based decision making [...] Valence produced by many different properties of a thing or event can be summed and subtracted to produce an overall evaluative response, and such responses can be compared to enable us to choose among options that would otherwise appear incommensurable. (Carruthers 2017b, p. 1)

As a consequence, the specific source of the affect is lost (and needs to be coarsely reconstructed). This explains how process-specific fluency leads to judgments seemingly unrelated to the source-process of the fluency: The process-specific fluency leads to affect which is source-unspecific for matters of value cross-comparison. As a result, when experiencing, say, a FOK we choose to invest resources into retrieving the queried information from memory over an external source. This is because the FOK assigns a higher value to the latter than to the former. Consider a more complex FOK-like case: You have some idea that you want to write or say. Before you can do so, however, you get distracted for a brief moment. Upon returning to your idea you realise that you cannot spontaneously recall what it was that you wanted to write or say. What will often happen is that you will continue to query your memory until your fresh idea eventually emerges (hopefully). What is interesting about such cases is that at the moment your idea slips your mind and you find a gap in place of the content, the feeling of this mental gap infuses the absent content with a remarkably high value. In fact, when you finally manage to retrieve your idea, its content often fails to live up to your expectations, i.e. the high expected value you have put into it when searching for it. I, at least, have such experiences quite often. I take them to illustrate how, on a phenomenal level, a FOK-like state motivates retrieval attempts not via imbuing retrieval with comparative value but the absent content itself.

In the same vein that valence might potentially help us to make sense of the phenomenal polar opposite pair of positive and negative intuitions (see section 4.5), it can also do so with the numerous phenomenal polar opposite pairs we find in the class of epistemic feelings. The fact that there are phenomenal opposite pairs such as the feeling of familiarity and unfamiliarity, the feeling of knowing and not knowing, the feeling of confidence and uncertainty, the feeling of understanding and incomprehension, curiosity and boredom etc. can all be explained by positing that positive and negative phenomenal valence (partly) grounds these “inverted phenomenal twins”.

In this section, we have considered how Affectivism about *epistemic feelings* can con-

tribute to an understanding of several aspects of them. It can give us a more general perspective on the way some epistemic feelings are metacognitive without being explicitly about other mental states or about the subject. Furthermore, it can help us to solve a puzzle about how process-specific fluency can have quite process-unspecific influences on our behaviour: by giving rise to *transmodal* affect. Against the background of affective experiences being valenced, it also becomes easy to see why the class of epistemic feelings harbours many phenomenal polar opposites. I take the datum that understanding epistemic feelings as affective experiences can well accommodate—and, in fact, further elucidate—many observations about epistemic feelings to lend (abductive) support to the very idea.

6.5. Epistemic Feelings as Affective Experiences

I think that the findings and considerations of the foregoing sections establish some plausibility for Affectivism about epistemic feelings. In this section, thus, I go along and analyse epistemic feelings with the resources previously established for affective experiences. Before I plunge right into it, however, I don't want to remain silent on worries that some might still harbour concerning the idea that epistemic feelings are affective experiences. In the next section, I will try to actively address them with a mix of theoretical and phenomenological points about epistemic feelings.

6.5.1. The Mild Affective Phenomenology of Epistemic Feelings

In comparison to affective experiences such as migraines, fears or orgasms epistemic feelings are perhaps not as obviously experienced as positive, negative or highly arousing. As we have seen based on empirical findings, however, they nevertheless exhibit a subtle positivity or negativity and a degree of arousal. Still, this is not unlike milder affective experiences in general (see also Colombetti 2011). In this regard, epistemic feelings resemble certain instances of aesthetic experiences that are plausibly affective in nature (e.g. Prinz 2014; Goffin 2019; for a treatment of aesthetic experiences *as* epistemic feelings see Reber, Schwarz, and Winkielman 2004; Dokic 2016; Renoult 2016).

However, as James' approach demonstrates, when it comes to affective experiences, the focus tends to lie on a few "paradigm" cases of affective experiences such as pain and fear. In what sense are pains and fears paradigms of affective experiences? Without

doubt, they exhibit the features unique to affective experiences — valence, arousal and motivationality — to an *extraordinarily high* degree. But in being “very loud” *as* affective experiences, they are actually quite special, *rare* occurrences. A much bigger part of our affective life is constituted by the little, subtle movements of our affective sensibilities. These affective experiences are not only all too often neglected in favour of their few “violent” conspecifics but also *easy to neglect because of their calm nature*. It is very likely that the bad reputation feelings have acquired when it comes to our rational activities are to a good part due to the few “affective hooligans” that tend to get the lion’s share of our attention. This way we are likely neglecting the many essential roles that mild, low-profile feelings play in our “higher” faculties. This is a regrettable state of affairs since it gives away a great opportunity: an opportunity to understand better what drives us—what it is that guides our thoughts when we take ourselves to “deliberate” and “inquire”.

Now, we might be able to “triangulate” this mild part of our affective life that is so often lost to introspection by considering this:¹⁷² phenomenally obvious “paradigmatic” affective experiences are relatively rare occurrences in comparison to, say, perceptual experiences and thoughts which are with us all the time—literally. However, we know something about affective experiences that appears somewhat at odds with this apparent scarcity. Importantly, we can see the feature in question instantiated in paradigmatic affective cases: Feelings are caused by perceptual experiences and thoughts and they take them as their bases. Now, we have perceptual experiences and thoughts *all the time*. If these are involved with affective experiences, does it mean that they get only involved with them under exceptional circumstances? Does it mean that outside of these exceptional circumstances we go about our business as some kind of “Kantian Angels” driven purely by thought and perceptual experiences—only to be sometimes thrown off our enlightened path by affective seizures?

I don’t think this is the case. A more natural explanation seems to me that our ever-present perceptual experiences and thoughts lead to affective experiences that are just as ever-present. The majority of them, however, are not present as phenomenal ruptures but as gentle guides of thought and action.¹⁷³ The reason why we tend to think about

¹⁷²Thanks to Marco Inchingolo for making me think about this point.

¹⁷³Note that their gentleness doesn’t have to make them *less* persuasive. Quite the opposite, actually: they might be more likely to persuade us *because* of their measured nature—appearing to us in the form of the proverbial voice of reason. Maybe putting it metaphorically helps: when contemplating a subject matter I, for my part, am much more likely to be persuaded by the measured, tactful words of a well-respected friend than by my shrill, irascible uncle who, in his great and unmatched wisdom, is — once again — trying to impose his unsolicited advice on me. To who would you rather lend

affective experiences in the former “violent” way might be because we only tend to study the tip of the affective iceberg that happens to be phenomenally “on fire”. Furthermore, the function of affective experiences is likely not to be violent—it is to make things salient in the way echoed by Brady (Brady 2009, pp. 422 sq., see section 4.3.2 and 4.5). That is, they direct our attention towards something else *than themselves*, towards something that matters. It is thus not surprising that we are only able to get a good look at them in exceptional circumstances—such as when they *are* violent or when there is—consciously—nothing else relevant to look at (see section 6.5.2).¹⁷⁴

Now, add to this our documented unreliability to introspect the nature of our experiences, especially affective experiences, and you get a sense for why becoming aware of mild affect — while beneficial for theoretical and personal reasons — is not at all an easy task in itself (Haybron 2008; Schwitzgebel 2008).¹⁷⁵ If epistemic feelings are *mild* affective experiences, it is rather unsurprising that their affective nature tends to elude us.¹⁷⁶ This consideration is nicely echoed by Ben Bramble when reflecting upon the phenomenology of pleasantness (‘the pleasant feeling’) and unpleasantness (‘the unpleasant feeling’) in general:

Consider what ‘the pleasant feeling’ would have to be like [...] It would have to be the sort of feeling that can occupy an experience, and so make it count

your ear? And whose advice would you rather feel the urge to *resist*? I think this illustrates on, the one hand, how often we are willing to accept the insinuations of our calm affective guides as eminently reasonable (likely not even noticing *that* they are affective) and, on the other hand, how we are regularly annoyed by the overdrive of their vociferous relatives.

¹⁷⁴As Mangan relatedly points out (he refers to epistemic feelings as “non-sensory experiences”):

Even if we indirectly recognize that non-sensory experiences are present in consciousness, it is still very difficult to grasp them as objects of direct attention. The attempt to do so, in at least the great bulk of cases, instantly changes the character of the non-sensory experience, and brings a sensory content squarely into attention. (Mangan 2001, p. 15)

He goes on to quote James on this observation:

It is very difficult, introspectively, to see the transitive parts for what they really are. If they are but flights to a conclusion, stopping them to look at them before the conclusion is reached is really annihilating them. Whilst if we wait till the conclusion *be* reached, it so exceeds them in vigor and stability that it quite eclipses and swallows them up in its glare. (James 1890, pp. 243 sq.; for similar points for intuitions see Chudnoff 2011b, pp. 642 sq.)

¹⁷⁵Affective experiences likely do not rely on this kind of reflexive consciousness to assume (the lion’s share of) their functional roles. As Carruthers observes:

Indeed, it would be quite odd if the first-order causal role distinctive of a given type of mental state [...] required a higher-order representation to be present specifying that a token of that type of state is currently active. (Carruthers 2017a, p. 73)

¹⁷⁶That’s why we need to go beyond introspection and do empirical and theoretical work.

as pleasant, by *permeating* it. Consider, for example, pleasant experiences of listening to Bach, eating a juicy peach, solving a puzzle, sunbathing, etc. Clearly, if ‘the pleasant feeling’ exists, it does not make these sort of experiences pleasant by being ‘tacked on to them’, so to speak, in any crude fashion. Instead, it must be the sort of feeling that can come in extremely low intensities, and very finely discriminable locations within one’s experiential field, so that it can come scattered throughout one’s experiential field. If [...] I enjoy listening to Bach, while you do not, then the difference between our experiences of Bach has got to be that mine is permeated by ‘the pleasant feeling’, while yours is not. In this way, ‘the pleasant feeling’ might ‘brighten’ a whole experience, or lend it a ‘warm glow’ [...] This explains, it seems to me, why [...] we should not expect to be able to gain a clear sense of ‘the pleasant feeling’, or the way in which all pleasant experiences feel alike. The reason is that [...] most instances of ‘the pleasant feeling’ are, taken by themselves, virtually imperceptible. They occur in extremely small quantities (or low intensities), and in very abstract or ethereal locations in one’s experiential field, locations that are not at all easy to direct one’s attention toward, or focus upon. [...] That said, one can perhaps gain some sense of ‘the pleasant feeling’ by thinking, say, of the experience of orgasm as ‘the pleasant feeling’ delivered in a very large quantity (or a high intensity), explosively. Similarly, one might gain some sense of ‘the unpleasant feeling’ by thinking of a painful experience like stubbing one’s toe as ‘the unpleasant feeling’ delivered in a large quantity, and a particular bodily location, explosively. (Bramble 2013, pp. 209 sqq.)¹⁷⁷

Apart from that, there are straightforward reasons for the milder affective nature of epistemic feelings. A central one is that bodily sensations, while being not affective by themselves, are components of the *paradigmatic* affective experiences such as pain and fear. As a consequence, much of the phenomenology that people ordinarily take to be affective is actually somatic in nature.¹⁷⁸ Thus, in deciding whether a given

¹⁷⁷Similar considerations can be found in Mangan’s discussion of James’ concept of the fringe of consciousness (Mangan 2001).

¹⁷⁸Prinz — following James’ famous subtraction thought experiment — makes a similar point to the effect that if one subtracts all somatic phenomenology, no affective phenomenology remains. His primary reason for doing so is introspective.

The hypothesis that valence markers have a distinctive feel is certainly bolstered by intuition. It seems that negative and positive emotions feel significantly different in virtue of their difference in valence. There is, however, an alternative possibility. Perhaps the felt differences between negative and positive emotions is a consequence of the embodied

experience is affective people will not (only) look for affective (i.e. valenced) but for somatic phenomenology. However, bodily sensations are not very present in epistemic feelings. One reason for that lies in the intentionality, another in the motivational of epistemic feelings: 1) In contrast to bodily and emotional feelings, epistemic feelings do not typically take bodily sensations but exteroceptive and cognitive states as bases. 2) In contrast to bodily and emotional feelings, epistemic feelings naturally motivate *mental* rather than bodily behaviours (Proust 2008). Though only the motivation of bodily behaviour goes together with the (global) behavioural preparation of the body (physical action readiness), adding a myriad of attendant bodily sensations (Frijda 1986).¹⁷⁹ These two factors explain why there is a lack of bodily phenomenology in epistemic feelings.¹⁸⁰ This lack of “mock-affective” bodily phenomenology will occlude the *affective* nature of epistemic feelings.

I have been fighting here a somewhat defensive battle concerning the power of introspection to shed light on the affective nature of epistemic feelings. Yes, epistemic feelings are usually mild affective experiences – typically the positive or negative affect integral to them does not, as Bramble has described it, come “in a very large quantity (or a high intensity), explosively”. This is, however, not to say that epistemic feelings cannot be reasonably intense, giving us some introspective evidence for their affective nature.¹⁸¹ To demonstrate this, I ask you to read the following passage (while trying to understand what it is about):

A newspaper is better than a magazine. A seashore is a better place than the street. At first it is better to run than to walk. You may have to try

appraisals they contain. [...] it could turn out that the feeling of the emotion is exhausted by the appraisal. Alleged commonalities between the feelings of different negative emotions could be an illusion. Perhaps there is no phenomenal thread linking disgust, betrayal, and grief. I favor this view. I do not think that valence markers have any phenomenology in their own right. That claim may seem untenable. [...] Is there any reason to deny that valence markers have intrinsic feels? My main motivation is introspective. Despite a strong intuition that there is a special feeling associated with unpleasantness, I cannot find anything that answers to this feeling when I introspect. When I mentally subtract away the feelings associated with the embodied appraisals of disgust, betrayal, and grief, I do not find any remainder. (Prinz 2004b, pp. 176 sqq.)

The empirical findings and reflections presented above suggest that he is mistaken (see also Mitchell 2018).

¹⁷⁹Part of Bramble’s and mine observation above is perhaps also explained by the assumption that a larger amount of affect is needed to mobilise us globally, with “body and mind” as it were. In most instances, it presumably takes less affect to motivate a simple “move of the mind”.

¹⁸⁰However, as the *tip-of-the-tongue* experience and the measures of facial muscle activity mentioned above attest, it is not that epistemic feelings are never tied to bodily sensations.

¹⁸¹As we have seen, when James brackets out some emotions from his analysis, he willingly grants this.

several times. It takes some skill but it is easy to learn. Even young children can enjoy it. Once successful, complications are minimal. Birds seldom get too close. Rain, however, soaks in very fast. Too many people doing the same thing can also cause problems. One needs lots of room. If there are no complications it can be very peaceful. A rock will serve as an anchor. If things break loose from it, however, you will not get a second chance. (Bransford and Johnson 1972, p. 722)

How do you feel? Probably confused, unable to understand—this *feeling of confusion* or *incomprehension* is another negative epistemic feeling (e.g. Silvia 2010; Lodge et al. 2018; Arguel et al. 2019). Now try to attend to what phenomenally happens when I give you the following hint: kite. You likely feel much better now; all of a sudden everything seems to fall into place. What you just experienced is a *feeling of understanding*, sometimes also called *aha-* or *Eureka-experience* (e.g. Trout 2002; Bowden et al. 2005; Dodd 2014). I hope that this example does away with the idea that epistemic feelings cannot be intense.¹⁸² I'll provide more examples of relatively intense epistemic feelings in section 7.2.

In this section, I hope to have addressed some worries about the idea that epistemic feelings are affective experiences. In the course of this endeavour, I also strived to illustrate and explain some points about the affective phenomenology of epistemic feelings: that they are usually only mildly valenced and arousing and that the behaviours they motivate are often of a covert mental or cognitive nature. The next section will concern the affective intentionality of epistemic feelings.

¹⁸²In case the above example didn't do the trick for you or you just like to engage another time with this phenomenal curiosity try to get this:

The procedure is actually quite simple. First you arrange things into different groups. Of course, one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities that is the next step, otherwise you are pretty well set. It is important not to overdo things. That is, it is better to do too few things at once than too many. In the short run this may not seem important, but complications can easily arise. A mistake can be expensive as well. At first the whole procedure will seem complicated. Soon, however, it will become just another facet of life. It is difficult to foresee any end to the necessity for this task in the immediate future, but then one never can tell. After the procedure is completed one arranges the materials into different groups again. Then they can be put into their appropriate places. Eventually they will be used once more and the whole cycle will then have to be repeated. However, that is part of life. (Bransford and Johnson 1972, p. 722)

Confused? Frustrated? Well then: try *doing the laundry*.

6.5.2. The Affective Intentionality of Epistemic Feelings

Let's turn to the intentionality of epistemic feelings: as all affective experiences, epistemic feelings have intentional or particular objects, formal objects, and bases that provide the particular objects. I have said earlier that the bases of feelings can be diverse and have diverse kinds of contents. For instance, the content of a base (and thus of the feeling) might (appear to) be more or less conscious.¹⁸³ As Arango-Muñoz puts it for the FOK and TOT:

The subject is conscious of the valence of his or her feeling as positive or negative, but he or she is not yet able to determine the intentional object of the feeling. The FOK and the TOT point to the fact that the subject knows something without determining what is known. (Arango-Muñoz 2014, p. 201)

The FOK does have an intentional object, say, the answer-pointer or -cue “capital of Australia” or the one-word-answer itself “Canberra”, and it does have a base, say, the perceptual experience (hearing, reading) through which the answer-pointing question “What is the capital of Australia” was received. In fact, the base can be extended to include states that were triggered by hearing or reading the question and that more proximally led to the occurrence of the FOK. Such can include memory states that were prompted by retrieval attempts and that embody parts of the answer (e.g. starting letter, length of the word, semantic information) or other cues predictive of answer-retrieval success (e.g. process-properties of retrieval, familiarity with parts of the question) (Reder and Ritter 1992; Koriat 1995; Koriat, Levy-Sadot, et al. 2003, see section 5.3.1). Parts of this extended base might be conscious or consciously accessible (e.g. starting letter, semantic information) — just remember the felt “definite gap” James’ was talking about — while other parts might not (e.g. process-properties of retrieval). So although the base has content (question terms, parts of the answer) and part of it is conscious or consciously accessible, this content is not yet enough to enable the subject to fully identify the intentional object of her feeling.¹⁸⁴

¹⁸³Also the base itself can be unconscious, e.g. an unconscious perceptual or cognitive state.

¹⁸⁴An alternative to this construal is that the FOK's intentional object is a fully conscious part of the read or heard question such as “the capital of Australia” which the feeling represents as being “known” (for a related proposal of question-directed states see Friedman 2013; Carruthers 2018). The FOK might further be redescribed as the feeling that *I can answer this question (about the capital of Australia)* or, metarepresentationally, *I know the capital of Australia*, taking the latter, fully determinate propositional content as its intentional object.

Note that this is just another way of appealing to the familiar notion of *content-gradeability*. Some feelings, probably partly as a function of their base, are more *determinate* or *vivid*, i.e. clear and distinct in representing their particular objects than others.¹⁸⁵ We can think of the way the particular object is represented to be a function of the aspectual shape under which the base of the FOK represents its intentional object (see section 4.4.4). Plausibly, the feeling will motivate the subject to invest further resources to determine its object (or to determine the answer to the question that furnishes its intentional object). Importantly, since the feeling is an indicator of expected retrieval success, it will do so in a specific way: by relying on the subject's own memory instead of e.g. consulting external sources (Michaelian 2012).

Note further that while the particular object of the FOK might not be fully conscious, the feeling itself clearly is.¹⁸⁶ Indeed, it is the absence of the determinate intentional object or content¹⁸⁷ that makes the FOK salient in consciousness. Discussions of the transparency of (perceptual) experience often stress that it is the content of experience (and its features) rather than the experience (and its non-content features) that occupies centre stage in consciousness (e.g. Tye 1995). In the case of the FOK this appears to be different: there is no determinate content to bind consciousness. What stands out instead is the experience that points towards the promise of a (determinable) content: We become aware of knowing *without* knowing what we know.¹⁸⁸ The subject is aware of knowing in virtue of experiencing positive affect and being conscious of the fact that she *knows* something, i.e. of the formal object of the FOK.

¹⁸⁵I have couched content-gradeability in terms of varying amount of information available in consciousness. However, the dimension of content determinacy might be construed to be orthogonal to consciousness of content. Perhaps we can be fully conscious of *underdetermined* contents. To see this possibility, imagine the word "rose". Now: what colour did the letters have? Were they upper or lower-case? Were they big or small, close together or far apart? It is well possible that you cannot answer all these questions faithfully. And this is simply because the contents of your fully conscious imagining of the word "rose" were underdetermined. What this exercise shows is that (fully conscious) underdetermined content might be possible (but see Bourget 2017). Maybe James had something similar in mind when he wrote: "What a thought *is*, and what it may be developed into, or explained to stand for, and be equivalent to, are two things, not one." (James 1890, p. 279).

¹⁸⁶At least if the mental state in question is conscious, i.e. a FOK. If one accepts the notion of unconscious valence, one can imagine that in some contexts a role similar to the one of a conscious FOK is played by a functionally analogous but unconscious FOK-like nudge. This epistemic nudge would have unconscious valence and thus similar behavioural consequences (see also section 4.3.1 and 6.3.2). Unfortunately, much experimental work on FOKs does not rule out FOK-like nudges in favour of conscious FOKs.

¹⁸⁷Note that I use "content" here to refer only to the part of the content constituted by the intentional or particular object, not to other aspects of it, such as the formal object (or, depending on the reading, its aspectual shape).

¹⁸⁸In the sense that when I have a FOK when asked "What is the capital of Australia" I am not aware that the capital of Australia is *Canberra*. I might be aware that I know *the capital of Australia*.

The formal objects of epistemic feelings tend to be epistemic properties broadly construed: familiarity, uncertainty, understanding, knowing, rightness etc. It appears that these properties are not traditionally conceived as evaluative or affective. However, due to the subtlety of the affective experiences concerned with them, this is not particularly surprising. It is likely that we will have to rethink the way we routinely access these properties, at least when it comes to the phenomenal versions of them. How plausible is it that we have feelings that represent the mentioned epistemic properties? On second thought, don't philosophers quite often talk about the *value* of truth, knowledge and understanding? We seem to have hit upon a straightforward rationale for it: we are *literally* getting quite “emotional” about them. Think of it that way: Are these properties of relatively high survival value to our species (cf. Kozuch 2018)? And is the importance of epistemic properties — in contrast to e.g. specific colours — relatively invariant across contexts? If they are, then it is plausible that we have evolved a suite of mental states that efficiently detect these properties in our external and internal milieus (cf. Sperber, Clément, et al. 2010). It is natural to assume that this suite of states includes specific affective experiences: epistemic feelings that evaluate things as familiar, certain or right.

For such an evaluation it is not only crucial that epistemic feelings swiftly *detect* their feeling-specific properties but also that they behaviourally orient us in relation to them. This is where their motivational force comes in. As noted earlier, the motivational force of epistemic feelings stands naturally in close relation to mental and epistemic behaviours and actions (Proust 2008; Dokic 2012, p. 311). That affective experiences, in general, have a prominent role in non-overt mental behaviour has been often overlooked. However, it is something that is increasingly recognised (Frijda 2008, p. 72, see also section 4.5). While unexpected FOFs motivate sustained recognition attempts and expected FOFs engender a host of habitual stimulus-specific behaviours, FOKs privilege memory over web search (see also the FOK-like example from section 6.4). Bennett Schwartz and Janet Metcalfe nicely summarise these characteristics for the TOT:

TOTs implore us to continue our search to retrieve unrecalled words. The TOTs are like an itch that can only go away when the target word is retrieved. It is the “itch” of the TOTs that causes us to find them problematic, but it is the drive that they give us to continue our search that is likely their cognitive function [...] [O]ur experiential feelings are flags or markers that inform us that a particular task is possible, that a particular item is memorable, or that a particular word is retrievable. These markers can then alter our behavior

through control processes that direct our behavior. Just as the feeling of pain causes us to withdraw our hands from a sharp object, the TOT feeling nags us, compels us, or beckons us to engage in explicit retrieval processes. (Schwartz and Metcalfe 2014, pp. 16, 26)

Epistemic feelings emerge as experiences with a valence, arousal and motivational dimension. In other words: epistemic feelings are affective experiences. And that they are despite the fact that their inherent valence and arousal are mild and the behaviours they motivate often of a covert mental or cognitive nature. That is, albeit epistemic feelings are not *that obviously* affective, they *are* nevertheless affective experiences. As all affective experiences, epistemic feelings have particular objects, formal objects and bases that provide their particular objects under base-specific aspectual shapes. These aspectual shapes, in turn, might themselves vary in their quality such as their determinacy. The feeling-specific formal objects of epistemic feelings are epistemic properties broadly construed: familiarity, uncertainty, understanding, knowing, rightness etc. As feelings in general, epistemic feelings “encode”, “localize” or “embed” value or disvalue in the form of their determinate feeling-specific property within the contents of their bases and motivate us to behave accordingly. These feelings, in other words, epistemically evaluate their particular objects.

6.6. Chapter Recap

To sum up for this chapter: Based on empirical findings reviewed in section 6.2 and 6.3 as well as a mix of theoretical and phenomenological considerations interspersed throughout section 6.4 and 6.5.1 I tried to establish that epistemic feelings are affective experiences and analyse them along these lines (section 6.5).

In section 6.2 we have seen that epistemic feelings covary with interoceptive changes, variations in SCR and facial muscle activity, all well-established bodily proxies for the affective properties of valence and arousal. Furthermore, positive epistemic feelings lead to increased liking and can serve as positive affective primes—behavioural proxies for the presence of valence. Additionally, performance capitalizing on positive epistemic feelings are enhanced in a positive mood. On the negative side, performance that relies on positive epistemic feelings is impaired under conditions of negative mood and depression. Further behavioural evidence comes from reviews of the influences of emotional and epistemic feelings on judgments. They bring to light that both kinds of feelings

influence judgments in a similar way and that similar factors moderate the reliance on them when making judgments. Coming from the affective to the epistemic side: Induced affect co-varies with judgments of seemingly non-affective properties such as truth, grammaticality, logical validity and semantic coherence.

In section 6.3 I first made the case for the covariation between epistemic feelings and affective properties being not just a correlation but a constitution relationship (section 6.3.1). For that, I presented studies that observe false positives of epistemic properties on the basis of incidentally induced affect. That is, inducing non-diagnostic affect leads subjects to incorrectly judge that an epistemic property is present. This speaks in favour of an affective constitution of epistemic feelings. Secondly, in section 6.3.2 I made the case that the affect in question is *conscious*. It thus not only causally biases epistemic behaviour but phenomenally constitutes epistemic feelings that provide conscious guidance for the subject's epistemic behaviour. To make this idea plausible I recounted studies where the following happens: the experimenters make the subject believe that the affect they experience in a given epistemic task is *not* diagnostic for the presence of an epistemic property. As a consequence, the subject loses her ability to accurately detect the epistemic property. This does not only indicate that epistemic properties are at times detected on the basis of affect but also that the affect in question is conscious.

Section 6.4 considered how Affectivism about *epistemic feelings* can contribute to an understanding of several aspects of them. It can give us a more general perspective on the way some epistemic feelings are metacognitive without being explicitly about other mental states or about the subject. Furthermore, it can help us to solve a puzzle about how process-specific fluency can have quite process-unspecific influences on our behaviour: by giving rise to *transmodal* affect. Against the background of affective experiences being valenced, it also becomes easy to see why the class of epistemic feelings harbours many phenomenal polar opposites.

Section 6.5 started out by illustrating the relatively mild affective phenomenology of epistemic feelings (section 6.5.1): that they are usually (but not always) mildly valenced and arousing and that the behaviours they motivate are often of a covert mental or cognitive nature. In doing so it also strived to remedy some worries about the idea that epistemic feelings are affective experiences. Finally, section 6.5.2 was concerned with the affective intentionality of epistemic feelings. As all affective experiences, epistemic feelings have particular objects, formal objects and bases that provide their particular objects under base-specific aspectual shapes. These aspectual shapes, in turn, might themselves vary

in their quality such as their determinacy. The feeling-specific formal objects of epistemic feelings are epistemic properties broadly construed: familiarity, uncertainty, understanding, knowing, rightness etc. As feelings in general, epistemic feelings “encode”, “localize” or “embed” value or disvalue in the form of their determinate feeling-specific property within the contents of their bases and motivate us to behave accordingly. These feelings, in other words, epistemically evaluate their particular objects.

Intuition Experiences are Epistemic Feelings

7.1. Introduction: Looking Back and Ahead

We have covered some distance in the course of this thesis. It is high time to pause for a moment of reflection and to look back at what lies behind us. Chapter 2 gave intuition experiences a recognisable face in the form of a feature profile. Chapter 3 tried to demonstrate that intuition experiences do not quite resemble doxastic states. At the same time, there seemed to be a fair amount of resemblance with perceptual experiences. However, on closer look, this resemblance was bound to remain just that: resemblance. Intuition experiences ain't perceptual experiences. Comparing the former with the latter might help us to acknowledge the features of intuitions but mere resemblance with perceptual experiences is ultimately unable to *explain* the features of intuitions. Chapter 4 wandered off the beaten path of philosophical intuition theorising into the fascinating world of affective experiences. After taking a long, hard look at affective experiences, this expedition led us to discover a set of relatively familiar-looking qualities: juxtaposing the features of affective experiences with those of intuitions in section 4.5 revealed, again, a resemblance—quite a good one, in fact. This motivated the hypothesis that explains this good fit straightforwardly and fruitfully: intuitions and affective experiences resemble each other because intuitions just *are* affective experiences.

This is the seed of Intellectual Affectivism. Like the perceptual analogy, Affectivism

can acknowledge the features of intuitions. However, in contrast to a mere perceptual analogy, it can also *explain* the features of intuitions: Intuitions have the features they have not because they are similar to affective experiences but because they *are* affective experiences—and we know quite a bit about why affective experiences have the features they have. However, affective experiences are a diverse class of mental states with distinct subclasses. Thus, just claiming that intuitions are affective experiences and leave it at that appears a bit coarse-grained a classification. So we set out to identify a subclass of feelings that would make the most comfortable and plausible home for intuitions in our ontology of mind. We found this subclass in epistemic feelings, acquainted ourselves with them and reinforced their status as affective experiences (chapter 5 and 6).

And here we are now. We have seen Affectivism grow from an initial seed into a somewhat sturdy plant. Now, the present chapter marks the conclusion of Intellectual Affectivism by identifying intuitions more precisely with not just affective experiences and not just epistemic feelings but with *specific* epistemic feelings. The central claim I am going to advance here is that positive intuitions are feelings of rightness (FORs) and negative intuitions are feelings of wrongness (FOWs), two feelings within the subclass of epistemic feelings. Actually, the idea here is yet more precise: intuitions are specific varieties of FOR/Ws, namely feelings of truth (FOTs) and feeling of falsity (FOFs, from now on). Here is how I am going to proceed: In section 7.2 I will give you a first feel for FOR/Ws. Then, as I have done previously for other epistemic feelings, I will take you to a FOR/Ws science exhibition (section 7.3). You will see that the science that is explicitly concerned with FOR/Ws is as yet in its infancy (section 7.3.1) and that it has been so far conducted with a strong metacognitive focus, somewhat neglecting FOR/Ws that are not as obviously concerned with cognitive performance (section 7.3.2).

However, looking closer, we will discover that there are findings on feelings of logical validity that are plausibly construed as a variety of FOR/Ws. Finally, we arrive at what is commonly known as the truth effect on which there exists a rich empirical literature (section 7.3.4). The truth effect describes the heightened tendency to judge a statement as (more likely) true as a consequence of manipulating various variables. I will go on to argue that what we observe in the truth effect are actually the workings of two epistemic feelings and that these feeling are varieties of FOR/Ws: feelings of truth and feelings of falsity (section 7.3.5). After this diagnosis, I will provide an analysis of FOR/Ws based on the resources established for affective experiences and epistemic feelings (section 7.4). In the course of this analysis I will delineate FOT/Fs among FOR/Ws as those FOR/Ws that take propositions as their particular objects and represent them as right

or wrong, amounting to an evaluation of the propositions as true or false. This analysis will conclude in the realization that the feature profiles of FOT/Fs and positive and negative intuition experiences are identical. Thus, intuitions are identical to FOT/Fs and Intellectual Affectivism is successful. The next and last chapter 8 will consist of a deeper exploration of the mechanisms behind intuitions/FOT/Fs and their relation to actual truth and falsity.

7.2. Getting a Feel for Rightness and Wrongness

In the course of the present thesis, I have mentioned feelings of rightness (FORs) and feelings of wrongness (FOWs) (Mangan 2001; Thompson, Prowse Turner, et al. 2011; Thompson and Morsanyi 2012; Gangemi et al. 2015; Fernandez Cruz et al. 2016). What are they? Remember, for instance, the last time you were arranging furniture in your room until it “looked” or felt right. Recently, the internet has spawned a genre of video clips that capitalizes on the FORs of the audience. These clips show events and actions that typically involve the meticulous manipulation of physical objects such as peeling wood. In fact, “Oddly Satisfying” videos have become prominent enough to be featured in *WIRED* and *The New York Times* (Faramarzi 2018; Matchar 2019). Their appeal is admittedly better demonstrated than described. I recommend the same-named subreddit and YouTube channel.

I will now provide you with two long (but entertaining) introspective reports of sustained and highly intense (feelings resembling) FORs that are the result of ecstatic seizures and laughing gas (nitrous oxide gas) intoxication. These reports allow for an instructive peak into the phenomenal nature of FORs:

“Every detail of my perceptions was every bit as accurate as they ever are.”
“In dramatic contrast to the ordinary way of experiencing one’s surroundings, during my seizures, all of these boundaries would suddenly be erased. Although all my judgments of shape, size, color, texture, and so on would remain totally unchanged, the evaluation of my environment would undergo a sudden transformation. Everything would be joined together into one whole, as if every single thing in my surroundings were deliberately placed by an artist with the goal of composing a photograph. This would result in a sense of vividness which derived, not from any dramatic hallucination or visual “trick”, but from the fact that each object in my visual field was emphasized,

so to speak, by everything else. When these boundaries are erased, a second phenomenon begins - **all the ordinary facts about the environment seem suddenly to become infused with certainty and a sense of inevitability.** “One often has (what is sometimes called) an “aha!” moment when we can suddenly explain several puzzling facts simultaneously with the same answer. The sense that I had when I was experiencing some of these seizures was not unlike a continuous series of profound “aha!” moments. Although nothing around me seemed to have changed in any concrete way, every observation of my surrounding environment seemed to “make sense” in this way. It is the sense, as I mentioned before, that one might have when admiring an expertly composed painting or photograph – each detail seems to be the way it is for a reason, even if that reason is difficult to articulate or seems only to float at the very edge of one’s consciousness. “The great consilience, coherence, and vividness of everything in the world seemed to demand an explanation of its organization. And the explanation was immediate and completely convincing – some mind or purposive agency was at the root of the world’s organization. This was a conviction that was both incredibly vague and totally compelling. Speaking as someone who is not normally given over to such beliefs – I am an atheist who happens also to be a professor specializing in logic – these beliefs are wildly out of character for me. **But, looking back on these experiences, I am struck by two features of these beliefs, which I now believe may be related. First, they were absolutely immune to any rational doubt; indeed, they seemed not even to be possible candidates for any such doubt. Second, they seemed – perhaps ironically, given my earlier statement – to have many of the same features as the most justified, and rationally derived beliefs that a person could possibly hold.** For instead of merely being justified by one or several other considerations or observations, these seemed to be irrefutably supported by literally everything in the world. Literally everything that I could experience seemed to cohere with, and lend support to, the belief that there was some sort of agency behind it all. The level of apparent organization possessed by everything in the world simply demanded such an explanation.” (Picard 2013, 2496, original emphasis)

I have made some observations on the effects of nitrous-oxide-gas-intoxication which have made me understand better than ever before both the strength

and the weakness of Hegel’s philosophy. [...] With me, as with every other individual of whom I have heard, the keynote of the experience is the tremendously exciting sense of an intense metaphysical illumination. Truth lies open to the view in depth beneath depth of almost blinding evidence. The mind sees all the logical relations of being with an apparent subtlety and instantaneity to which its normal consciousness offers no parallel; only as sobriety returns, the feeling of insight fades, and one is left staring vacantly at a few disjointed words and phrases, as one stares at a cadaverous-looking snow peak from which the sunset glow has just fled, or at the black cinder left by an extinguished brand. [...] I have sheet after sheet of phrases dictated or written during the intoxication, which to the sober reader seem meaningless drivel, but which at the moment of transcribing were fused in the fire of infinite rationality. [...] What’s mistake but a kind of take? What’s nausea but a kind of -ausea? [...] Emphasis, emphasis, there must be some emphasis in order for there to be a phasis. [...] There are no differences but differences of degree between different degrees of difference and no difference. [...] The last phrase has the true hegelian ring, being in fact a regular *sich als sich auf sich selbst beziehende Negativität*. And true Hegelians will *überhaupt* be able to read between the lines and feel, at any rate, what *possible* ecstasies of cognitive emotion might have bathed these tattered fragments of thought when they were alive. (James 1882, pp. 206 sq.)

One thing that these examples show is that the intensity of a feeling does not have to correlate with its accuracy. At the same time, the extremely magnified feelings in these reports are the result of *exceptional* mental conditions. It stands to reason, however, that they are “just” extraordinary representatives of feelings that are usually ordinary and regular ingredients of our phenomenal life—but in *much* lower dosages. This is what makes these testimonies so remarkable: we encounter FORs with their volume turned up very high and a positivity that comes “in a very large quantity (or a high intensity), explosively”.

When it comes to FOWs my favourite illustration consists in making you look at upward flowing water. Here is a video of “upward flowing water”: <https://youtu.be/NiOafQZwn0g>. Looking at it, you supposedly experience an unpleasant FOW about what you *see*, i.e. the feeling takes a perceptual experience as its base. It might also work to look at a static impossible shape like a Penrose triangle:



Depiction of a Penrose triangle

Alternatively, you can *imagine* upward flowing water (or any upward “falling” object) — or can you *imagine* impossible shapes? — with similar results. Then your FOW would take an imagining as its base. Of course, there seems nothing really wrong about upward flowing water in a substantial sense. There is nothing really wrong about cars driving on the left side of the street, as you can see here: <https://youtu.be/Y8DqAIyuR0Qs>. Still, many people coming from countries with right-hand traffic feel that there is something wrong and not just left about it. Perhaps you also remember some occasion when something looked, sounded, felt (in the tactile sense), smelled or tasted “weird” to you. I submit that in such situations we experience FOWs where we lack a clear way to refer to or describe the kind of wrongness. Thus, we call it weirdness.

The difficulty to arrive at FORs and a high susceptibility to experience FOWs in relation to mundane tasks, such as washing one’s hands or locking the door, is a characteristic symptom of obsessive-compulsive disorder (OCD), a psychopathology marked by the display of debilitating obsessive-compulsive behaviour.¹⁸⁹ In the OCD-literature FOWs are known by other names: “not just right experiences” (NJREs) or “incompleteness experiences” (INCs). They are found to “occur in any sensory modality” (Ben-Sasson et al. 2017, p. 2) and are assessed via surveys where subjects respond to items such as “I must do things in a certain way or I will not feel right”, “When hanging a picture on

¹⁸⁹ Another complementary possibility is that there is an absence of an expected FOR that would typically terminate the activity. This appears similar to the absence of an expected FOF towards loved ones in Capgras disorder.

the wall, I have had the sensation that it did not look just right” and “When talking to people, I have had the sensation that my words did not sound just right” (Belloch et al. 2016). Although FOWs are pronounced in OCD patients, they are fairly typical experiences in everyday life (cf. Cogle and Lee 2014). Presumably, a majority of us experiences FOWs on seeing a crooked picture.¹⁹⁰

7.3. The Science of FOR/Ws

In this section I will make first some qualifying remarks about the state of FOR/W research: work that is explicitly on the FOR/W has been to the present-day relatively scant. Furthermore, various things can be “right” or “wrong” in various ways—I will clarify what this means for FOR/Ws. Then I introduce the distinction between performance- and object-directed FOR/Ws which will inform my take on explicit FOR research that has a performance focus. After that, I will discuss experimental work which I construe as dealing with FOR/Ws that are not performance-directed. It will emerge that different kinds of rightness evoke liking. One of them is at the core of the present thesis: truth. I will introduce research on the truth effect, a tendency to judge a statement as more likely true as a result of various experimental manipulations, and connect it with processing fluency and affect.

Evaluation:

rightness)

7.3.1. The Paucity of FOR/W Research and Ways of Being Right or Wrong

When it comes to FOR/Ws it is important to emphasise a couple of things. One is the paucity of FOR/W-specific research: “The work on FOR [...] is in its infancy, so there are a great many open questions” (Thompson 2014, p. 59). Thus, when discussing FOR/Ws, additionally to what we specifically know about them, it is useful to draw on what we know about better explored epistemic feelings such as the FOK and FOF as well as on what we know about affective experiences in general. As a consequence, many of the points about the nature of FOR/Ws are as yet speculative and open to

¹⁹⁰For more: there is an “Oddly Unsatisfying” analogue to Oddly Satisfying videos on the web.

empirical investigation. That they are so open, however, I consider an advantage of the approach.

Second, various things can be “right” or “wrong” in various ways. It seems that a crooked picture is wrong in another way than an inference that affirms the consequent or the computation $3 + 2 = 7$. In these cases the FOW seems to variously bear on (a)symmetry, logical (in)validity or arithmetic (in)correctness. Thus, FOR/Ws can indicate various kinds of rightness and wrongness. There are two ways in which this can play out. On the one hand, FOR/Ws might not be a unitary kind but actually a whole class of different feelings determinable in terms of specific kinds of rightness and wrongness. The other possibility is that there might be only one kind of FOR/W and that the specific rightness or wrongness it indicates is specified by its particular object that is (part of) its content.

I favour the latter unitary kind view to the former multiple kinds view. Consider the implications of the two views. According to the multiple kinds view, the FOW we experience towards a crooked picture relates to the FOW that we experience towards the computation $3 + 2 = 7$ in roughly the same way to how fear relates to sadness. Both are negative affective experiences but they are of a different kind and have different formal objects. According to the unitary kind view, the two mentioned FOWs relate to each other in roughly the same way to how one’s fear of a bear and one’s fear of having lost one’s keys relate to each other. The kind of negative affective experience would be the same and have the same kind of formal object. Different things can be fearsome and merit fear in different ways.¹⁹¹ Nevertheless, it appears that the affective experience in the bear scenario and the one in the lost keys scenario are both properly called fear. In other words, we do not posit that one is of a different kind from the other. Now, the same might be true for FOR/Ws: Although different things can be right or wrong in different ways, the affective experiences that alarm us to their rightness or wrongness are of the same kind. It seems plausible to me that the FOW directed at a crooked picture and the FOW directed at the computation $3 + 2 = 7$ are very similar, more similar, in fact, than the bear and the lost keys fears that differ in their particular objects. The formal objects of both FOWs seems to me properly described as wrongness, despite the fact that the particular kinds of wrongness differ due to the difference in the particular objects that are wrong.

Note that there is also plenty of room for the context in determining the *content and*

¹⁹¹Don’t ask me how losing one’s keys is supposed to be fearsome...

kind of the feeling that we experience. Taking in an inference that affirms the consequent might lead to a FOW that in this context will indicate logical invalidity. Nonetheless, depending on the context (including the subject's constitution), it might just as well lead to a FOR, confusion, boredom, enthusiasm or no feeling at all. Furthermore, the multiple kinds and unitary kind view disagree about the variability of FOR/Ws in *kind* across particular objects. *Both* views, however, allow for variability along dimensions such as phenomenology, intentionality or distal and proximal aetiology.¹⁹² It is, therefore, an open question on what level(s) a feeling of rightness (of kind/with content) X resembles or differs from a feeling of rightness (of kind/with content) Y.

7.3.2. A First Look on FOR/W Research Through a Distinction

At this point, I want to make one intentionality/content/particular-object-based distinction between *performance-directed* FOR/Ws and *object-directed* FOR/Ws. Most if not all work that is explicitly concerned with “feeling of rightness” or “feeling of wrongness” focuses on the former. Performance-directed FOR/Ws indicate that one has successfully or unsuccessfully performed some operation, activity or action, say computed the result of $2 + 3$, recalled the capital of Australia or drawn a conclusion. In other words, they evaluate a self-generated performance.¹⁹³

As Mangan notes “feelings of rightness and wrongness are able to indicate the success or failure of what are otherwise distinct cognitive activities” (Mangan 2001, p. 17).¹⁹⁴ Mangan mentions here specifically *cognitive* activities, giving FOR/Ws a clearly metacognitive spin (see also section 5.3.3 and 6.4). With this metacognitive focus, he is in good

¹⁹²The unitary kind view is, however, more restricted on the variability it allows. This is because it additionally requires that despite variability along some dimensions, some essential aspects remain in place as to justify speaking of the same kind.

¹⁹³One might furthermore distinguish between performance-directed and activity-directed FOR/Ws. While the former is directed onto a successfully or unsuccessfully *terminated* activity, the latter is directed onto an *ongoing* activity which in turn is successfully or unsuccessfully unfolding. It seems that the former can *inter alia* be based on the *outcome* (and its properties) such as the computed result or drawn conclusion while the latter can only be based on non-outcome feedback in the form of cues such as the amount of time an activity is taking to run to completion (i.e. fluency). James' “feeling of [...] a right or wrong direction in the thought” (James 1890, p. 261) is, I submit, best conceptualised as an activity-directed FOR/W where we feel that our thought is on the right (or wrong) track without being aware of where it is headed exactly.

¹⁹⁴In fact, Mangan appears to conceptualise the FOK as a kind of activity-directed FOR that is directed at retrieval processes (Mangan 2000). Mangan's position seems to be that most epistemic feelings are essentially FOR/Ws with different contents/particular objects. I disagree: I think that e.g. FOKs, FOFs and FORs are different in kind from each other. In fact, a FOK, FOF and a FOR can have the same particular object, say, Diego Velázquez' iconic painting *Las Meninas* while being still different kinds of experiences and having distinct formal objects (see section 7.4).

company (see e.g. also Thompson, Prowse Turner, et al. 2011; Fernandez Cruz et al. 2016). In fact, most if not all of the research explicitly concerned with FOR/Ws is not only specific to performance-directed FOR/Ws but to performance-directed FOR/Ws where the evaluated self-generated performance is *cognitive* in nature. Taking one’s hint from performance-involving Oddly Satisfying videos and OCD-research, there seems to be no principled reason, however, to restrict performance-based FOR/Ws (and other FOR/Ws, really) to the cognitive domain. That is, performance-based FOR/Ws do not have to be metacognitive in the sense of being about first-order *cognitive* processes or states—at least when it comes to what is apparent on the personal level.¹⁹⁵ Washing one’s hands, recreating the form of a specific bodyweight exercise or releasing an arrow when shooting a bow can feel right or wrong in a similar fashion.¹⁹⁶

Performance-directed FOR/Ws are closely related to feelings of confidence or certainty and might be (in some contexts) identical to them. To illustrate this point, it is instructive to take a look at the work of Valerie Thompson, the primary source of research that is explicitly on the FOR. Thompson and colleagues, for instance, presented subjects with a conditional statement such as (Thompson, Prowse Turner, et al. 2011, p. 113):

- If a car runs out of gas, then it will stall.

They were then asked to assess the logical validity of one of the following inferences, responding with “yes” in case it is valid and otherwise with “no”:

¹⁹⁵That is, I want to leave it open that there might be some sense in which (some of) these feelings are metacognitive when one looks at the subpersonal level (see also section 5.3.3 and 6.4). Take for instance the FOW you experience looking at upward flowing water. This particular instance of a FOW might actually be caused by a specific cognitive process or one or several of its properties, say, its disfluency. Nevertheless, on a personal level, the feeling itself represents *what you see* as wrong, not the seeing. It is the upward flowing water that appears wrong, not the processes through which you take in the upward flowing water. The disfluency can be understood as a metacognitive cue, i.e. as a property that is informative about a cognitive process. This cue, in turn, is causally implicated in the generation of the FOW. For some this might render the FOW metacognitive. Nonetheless, it should be clear that such a metacognitive nature is not apparent to the subject having the FOW. To her the FOW appears not about herself but about the world; it is as much about the wrongness of the upward flowing water as Linda’s fear is about the dangerousness of the approaching bear (and not, say, about her inability to fight the bear).

¹⁹⁶The FOR/Ws here would not be metacognitive if one understands “cognitive” in a restricted sense as relating to thought and its kin. They would be metacognitive if one has a liberal take on what is “cognitive” as related to information processing in general. Information processing, of course, also prominently figures in the performance of e.g. motor activity. In fact, some of the central elements of some theories of metacognition are directly inspired by the monitoring and control architectures that implement motor activity (Proust 2015). Note that if metacognition is defined with a very broad concept of “cognitive” in mind, it might make the resulting concept of metacognition unsuitable in respect to some explanatory projects such as explaining metarepresentational capacities (cf. Carruthers 2017a, see also section 5.3.3 and 6.4).

1. The car has run out of gas. Therefore it will stall.
2. The car has not stalled. Therefore it did not run out of gas.
3. The car has stalled. Therefore it ran out of gas.
4. The car has not run out of gas. Therefore it will not stall.

Crucially, the researchers were interested in measuring the subjects’ “feeling of rightness”. How did they go about it? After each validity judgment they required a subject to respond to the statement “at the time I provided my answer I felt”. For their response the subject could choose between seven options on a scale ranging from “guessing” (1 on a Likert-scale) to “Certain I’m right” (7 on a Likert-scale). This report was supposed to tap a FOR and its magnitude. The task is performance-directed and metacognitive by design, measuring a FOR directed at one’s cognitive performance (for an analogous take on the FOW see Gangemi et al. 2015; Fernandez Cruz et al. 2016). In fact, what is measured here can be straightforwardly characterised as a feeling of confidence or certainty about one’s response.¹⁹⁷ Or, perhaps more concisely, one could term the FORs at hand as “feelings of success” (FOSs). This would be analogous to the practice adopted in FOW research where performance-directed FOWs are aptly termed “feelings of error” (FOEs) (see e.g. Gangemi et al. 2015).

Be it as it may, in contrast to (cognitive) performance-directed FOR/Ws, *object-directed* FOR/Ws can be directed at *any* object. Insofar, (cognitive) performance-directed FOR/Ws emerge as a subset of object-directed FOR/Ws, namely those that take self-generated (cognitive) actions or activities as their intentional objects. As the example of a FOR directed at a particular furniture arrangement and a FOW directed at upward flowing water demonstrate, FOR/Ws plausibly do not have to concern anything the subject did. Instead, they can concern apparently external or subject-independent things, without any explicit self-reference. In fact, the apparent self-reference in performance-directed FOR/Ws is likely due to their specific particular objects: self-generated performances are particular objects with built-in self-reference. On the face of it, considering the personal level, there is nothing intrinsically self-referential, performance-directed or metacognitive about FOR/Ws. We can have feelings that mark all kinds of things as right or wrong—not only one’s own cognitive performances. I insist on this point because

¹⁹⁷Perhaps one could draw a distinction between *prospective* and *post-evaluative* feelings of confidence or certainty (Proust 2008). While the former *predicts* the success of an activity similar to the FOK predicting retrieval success, the latter signals the successful execution of an activity. One can then call the former feelings of confidence proper and the latter feelings of rightness. In the literature this distinction is often not made and both are called feelings of confidence.

my focus here will lie on object-directed FOR/Ws that are not (primarily) directed at performances or activities.

This is not to say that the relationship between performance-directed FOR/Ws and FOR/Ws directed at other objects is always straightforward. Consider again the task of Thompson and colleagues. What I am interested in is not so much the basis for the subject's judgment about her performance which, as the researchers seem to assume, might indeed be a performance-directed FOR. What I am interested in is the basis for the subject's judgment of logical validity since this will (sometimes) be the object-directed FOR concerning the conclusion of the inference. If both FORs (of different kind or content) exist, they are likely not the same. However, it might well be that (sometimes) the subject does *not* base her judgment of her performance on a performance-directed FOR but on her FOR that directly concerns the conclusion of the inference and its rightness.¹⁹⁸ That is, there is just one conclusion-directed FOR that does double duty: it guides the assessment of logical validity *and* the assessment of one's performance.¹⁹⁹ The FOR is used to answer the question "Am I right?" because it provides an answer to the question "Is this right?"²⁰⁰

¹⁹⁸Talking of object-directed FORs that are not performance/activity-directed and that concern inferences might be a bit confusing since inferences are best understood as transitional activities rather than as objects. Talking of FORs that concern the conclusion of an inference seem to be a better alternative here. These conclusion-directed FORs are still about the inference in virtue of being about the *result* of the inference: the conclusion. Inference-directed FORs, on the other hand, are best understood as performance- or activity-directed FORs that concern the act of drawing an inference. Note that the specific rightness of a conclusion and thus of a conclusion-directed FOR is intertwined with the specific premises that are given in a context. This implies that one and the same content might feel right as a conclusion against the background of given premises, and at the same time feel wrong on other grounds. For instance, $2 + 2 = 5$ feels wrong taken for itself. Though as a conclusion drawn from premises (1) "If the pope is a woman, then $2 + 2 = 5$ " and (2) "The pope is a woman", it might feel right. This shows that inferential validity can sometimes be a highly context-specific or localised kind of rightness that might be at odds with other kinds of rightness when e.g. conclusions might be valid but not sound. In fact, people often tend to confuse different kinds of rightness such as (context-specific) logical validity and plausibility (see e.g. Klauer et al. 2000, see also section 8.3.3).

¹⁹⁹Alternatively, since the two assessments take place at different moments, it is one's *recollection* of the conclusion-directed FOR that is the basis for the assessment of one's performance: one recalls to what extent the conclusion appeared right to one and judges one's confidence accordingly.

²⁰⁰This is similar to an idea of Evans (see also section 138):

I get myself in position to answer the question whether I believe that p by putting into operation whatever procedure I have for answering the question whether p . (Evans 1982, p. 225)

Gordon later gave this procedure a general formulation and termed it the "ascent routine":

Because this procedure answers a metacognitive question by answering a question at the next lower semantic level, I will call it an *ascent routine*. (Gordon 1995, p. 60)

To illustrate this double duty idea I ask you to have another look at the above-mentioned inferences from Thompson’s study and put yourself into the shoes of one of the participants. Likely you will have discerned that 1. is an instance of *modus ponens* (MP) and 2. of *modus tollens* (MT). Therefore, they are valid inferences. On the other hand, 3. is an instance of affirming the consequent (AC) and 4. of denying the antecedent (DA) and as such invalid inferences. The important aspect is how you arrived at these assessments. When reading the conclusions you might have had the feeling that some of them are right or, on the contrary, that some of them are wrong. These feelings presumably—perhaps together with some other determinants such as explicit analytic double checking²⁰¹—have engendered your judgement that 1. and 2. are correct while 3. and 4. are incorrect.²⁰² Of course, it is quite possible that your feelings led you, say, in the case of AC astray and you thought it correct. After all, AC is a widespread fallacy partly *because* it often exerts some pull to judge it correct by triggering fallacious FORs. Now, asked for your degree of confidence in your answer, what do you do? An obvious way is to use or recall the quality of one’s feeling.

This becomes all the more plausible if one imagines a slightly modified task where the subject does not judge a *given* conclusion but, provided with a conditional and an antecedent, has to come up with a conclusion herself (see also Thompson, Prowse Turner, et al. 2011, experiment 3). In some cases, one’s answer will be accompanied by a FOR that marks it as right. Is this now a performance-directed or a conclusion-directed FOR? It seems that it can serve as a basis for evaluating both: performance and conclusion. Or put differently: it can serve to evaluate the answering (performance-directed) as well as its result, the answer (object-directed). Which evaluative dimension

²⁰¹Such checks might have e.g. proceeded via constructing truth-tables or by rehearsing the premises and what follows and does not follow from them. One might, for example, get clear about the fact that in the case of AC it was never asserted that running out of gas is the *only* condition in which the car stalls. It is well possible that, confronted with such statements, one does not have a FOR/W right away but that they occur in the course or as a result of further reflective engagement with the elements of the inference.

²⁰²Compare this to an example from Boghossian 2003 discussed in detail by Chudnoff 2014b in the context of intuition:

- (1) If today is the 20th, then Martha Argerich is playing today in Carnegie Hall. [...]
- (2) Today is the 20th. [...]
- (3) Martha Argerich is playing today in Carnegie Hall. (Boghossian 2003, p. 225)

In its basic form, this *MP* example is very similar to 1. above. The claim is that we (might) have the positive intuition that (3). This is presumably so because (3) follows from (1) and (2), i.e. that (3) is true given (1) and (2). However, we draw the conclusion (3) not only because it demonstrably follows logically but because we have the intuition that (3) follows, given (1) and (2). This intuition—perhaps together with more explicit checks—leads us to infer or judge that (3). We might also say that (3) feels right and that is (partly) why we infer or judge accordingly.

will appear more salient will likely depend on the specific situation such as the specific task, instructions and one's goals. Note that, intuitively, it appears that one cannot assess one's performance independently of whether one's answer seems right or wrong to one.²⁰³ Thus, the conclusion-directed FOR seems more basic. There might be a *phenomenal* basis to this impression: When I contemplate an issue and conjure up related thoughts, those among them that appear right to me seem to be right independently of me, at least in the first place.²⁰⁴ However, on the subpersonal level, one's apparently "objective" FORs might ultimately be based on features of one's performance or on other subject-specific factors (more on this in section 7.3.4, 7.3.3 and 8.1).

These reflections establish *some* relevance of Thompson's work for my topic despite its metacognitive performance-focus. I will thus incorporate some of its insights. Thompson establishes that rightness/confidence ratings are *positively correlated* with 1) answer fluency, i.e. the speed with which a response is given, 2) probability of conclusion acceptance, 3) logical validity, 4) "believability", i.e. the real-world plausibility of a conclusion or the degree to which it is congruent with typical beliefs and *negatively correlated* with 5) conflict, i.e. if a conclusion is logically valid but not believable (e.g. "If a plant has roots, then it is a tree. This plant has roots. Therefore, it is a tree.") or vice versa, 6) rethinking time, i.e. in conditions where subjects were given the possibility to rethink and possibly change their initial "FOR-response" a lower FOR-rating was associated with longer rethinking time and 7) higher answer change probability. Thompson takes 6) and 7) to indicate analytic engagement, suggesting that a high/low FOR leads to a low/high amount of analytic engagement or scrutinizing behaviour (Thompson, Prowse Turner, et al. 2011; Thompson and Morsanyi 2012; Thompson, Turner, et al. 2013; Thompson and Johnson 2014a).

To sum up: in this section, I introduced and illustrated the distinction between performance- and object-directed FOR/Ws. Then I outlined how the bulk of research explicitly concerned with FOR/Ws (mainly work of Thompson and colleagues) focuses on the former. Finally, I summarised the findings of this research. In the next section, we will encounter

²⁰³As FOKs and prospective feelings of confidence or certainty show, one might nevertheless *predict* whether one's activity is likely to be successful.

²⁰⁴In fact, I often need to re-analyse certain thoughts in order to identify them as the product of my specific (theoretical) commitments or wishful thinking. Insofar it is hazardous to assume that the products of e.g. wishful thinking are phenomenally forthcoming about their wishful or subject-dependent nature *to the subject*. Wishful thoughts might, at first glance and without further checks, seem just as objectively right as non-wishful thoughts (cf. Koksvik 2011, p. 200). In fact, typically the wishful nature of one's thinking needs to be pointed out to one by someone else (see also footnotes 70 and 169).

work that is similar to Thompson’s but differs in the crucial respect that it looks at what is plausibly construed as object-directed FOR/Ws that are not about performance.

7.3.3. Liking Rightness

Using similar stimuli as Thompson, Morsanyi and Handley 2012 and Trippas et al. 2016 found that the believability and logical validity of a conclusion leads to a higher liking of it. This is evidence for a link between affect and properties such as believability and logical validity. The researchers provide an explanation inspired by Topolinski’s account of semantic coherence (see section 6.3, especially section 6.3.2): a believable and valid conclusion is processed in a way that gives rise to positive affect which can influence a variety of judgments, among them judgments of liking. Specifically, the researchers conjecture that as in the case of semantic coherence, believability and logical validity lead to *fluent* processing that triggers the affect used for the judgment. I take this as evidence that people have an affective sensitivity to various varieties of rightness of a conclusion. There is, however, a caveat in comparison to the findings of Topolinski. While the data suggests that believability and logical validity trigger positive affect, it is unfortunately inconclusive on whether this positive affect is typically also used as a basis for judgments of believability and logical validity. I say inconclusive rather than silent because Morsanyi et al. generated findings where subjects indeed had to judge the logical validity of the conclusion rather than provide a liking rating. These findings, in turn, seem to suggest that:

[P]articipants generated their liking ratings in a different way from their validity judgments and that different factors affected participants’ responses when they performed the two tasks. (Morsanyi and Handley 2012, p. 610)

However, the liking rating and the validity judgment task differed quite dramatically in their instructions. The former suggested a spontaneous, “intuitive” response while the latter suggested a deliberate, reasoned response.²⁰⁵ As the authors themselves argue,

²⁰⁵Specifically, the instructions for the liking rating condition were:

In this experiment you will have to read series of three statements. Each series will be presented in three parts. The first statement will be presented for 2 seconds, then the second statement will be presented for 2 seconds, and finally the last statement will be presented for another 2 seconds. After this a picture of a face will appear on the screen to indicate that you have to give a response, based on HOW MUCH YOU LIKE THE LAST STATEMENT. WHEN YOU DO THE TASK, READ ALL STATEMENTS VERY CAREFULLY. WHEN YOU MAKE THE LIKING JUDGMENT FOCUS ON YOUR FEELINGS ABOUT THE STATEMENT. DON’T THINK ABOUT WHY YOU LIKE OR DISLIKE THE STATE-

this disparity might have altered the degree to which subjects relied on their feeling to judge the target criterion. In fact, Topolinski and Strack 2008 have shown for semantic coherence judgments that while just reading a word triad automatically activates a common associate, performing an intentional, effortful search for a common associate does not. Engaging in the latter intentional task, in fact, impairs the coherence detection performance. Morsanyi et al. suggest that something similar might have been at work in producing their findings (Morsanyi and Handley 2012, p. 610).²⁰⁶

Unfortunately, their data does not allow for an assessment of whether *not* relying on one's feelings impairs the validity detection performance in a similar fashion to Topolinski's findings. Interestingly, Thompson and colleagues found repeatedly that, when given the opportunity, rethinking time and changing of the initial "intuitive" response (a measure of analytic thinking) does not lead to better performance, i.e. higher accuracy in producing correct responses (Thompson, Prowse Turner, et al. 2011; Thompson and Johnson 2014b).²⁰⁷ Despite the inconclusiveness of Morsanyi and colleagues' data, I speculate that it is highly plausible that (fluency-induced) positive affect *is* typically used for judgments of logical validity and believability. Although Thompson and colleagues did not test for the presence of affect, this is to some degree suggested by their findings: feeling of rightness/confidence ratings were not only positively correlated with answer

MENT, JUST GO WITH YOUR INTUITION AND GUT FEELINGS. (Morsanyi and Handley 2012, p. 608)

The instructions for the validity judgment condition were:

In this experiment you will have to read syllogistic reasoning problems. Each problem will be presented in three parts. The first premise will be presented for 2 seconds, then the second premise will be presented for 2 seconds, and finally the conclusion will be presented for 2 seconds. After this a picture of a face will appear on the screen to indicate that you have to give a response. Your task will be to decide whether the conclusion necessarily follows from the premises. You have to reason on the basis of the premises given to you, even if they don't mean anything or they are not true in real life. You have to TAKE THE POINT OF VIEW OF A PERFECTLY LOGICAL PERSON AND ONLY CONSIDER THE INFORMATION THAT IS GIVEN TO YOU, WHETHER IT MAKES SENSE OR NOT. THINK ABOUT YOUR ANSWER VERY CAREFULLY. DON'T RUSH. ALTHOUGH THE PROBLEMS WILL BE PRESENTED VERY QUICKLY, WHEN YOU MAKE YOUR DECISION ABOUT THE CONCLUSION, YOU CAN TAKE AS MUCH TIME AS YOU WANT.(ibid., p. 608)

²⁰⁶Lending some further support for this explanation, Halberstadt and Catty 2008 show that analytical thinking also hampers familiarity-related judgments. In fact, the authors trace this finding back to a long lineage of work. This literature suggests that for a host of often complex judgments, relying on analytical thinking instead of subjective experience leads to worse outcomes according to subjective and objective criteria (e.g. Wilson and Schooler 1991; Wilson, Lisle, et al. 1993; Dijksterhuis 2006; Dijksterhuis and Nordgren 2006; Dijksterhuis and Van Olden 2006).

²⁰⁷See also footnote 206.

fluency but also with logical validity and believability.

My speculation receives strong support from Topolinski and colleagues' findings clearly showing that judgments of semantic coherence *are* based on (fluency-induced) positive affect. Semantic coherence between words and “coherence” between different propositional elements of information in the form of logical validity and believability bear some resemblance as properties to each other.²⁰⁸ And, in fact, semantic coherence, logical validity as well as believability are found to lead to positive affect. Now it would be surprising if semantic coherence is sometimes judged on the basis of positive affect and logical validity and believability are not. It is more plausible, I think, that positive affect can be the basis of the judgment in all three cases. However, we will need further empirical work to see whether this is actually the case.

7.3.4. The Truth Effect

Another line of research under the header “truth effect” deals with a kind of rightness at the very core of the present thesis: truth. The truth effect describes the heightened tendency to judge a statement as (more likely) true as a consequence of manipulating various variables. Importantly, most of this work tests the truth effect in relation to *ambiguous* statements such as “the zipper was invented in Norway”²⁰⁹, i.e. the judging subject does not have definite information (e.g. knowledge) of the truth value of the statements. What are the variables that exert a truth effect? The most prominent and reliable one is repetition. Subjects rate statements they have been previously exposed to as more likely true (Hasher et al. 1977; Arkes, Boehm, et al. 1991; for a review see Dechêne et al. 2010).

This effect is well documented and robust across many experimental variations: it occurs for statements that are 1) read (Schwartz 1982) or heard (Hasher et al. 1977; Begg, Anas, et al. 1992); 2) known to be repeated (Bacon 1979; Nadarevic and Aßfalg 2017); 3) a matter of general knowledge (Hasher et al. 1977), trivia (Bacon 1979) or opinion (Arkes, Hackett, et al. 1989); 4) repeated minutes, weeks or months apart (Brown and Nix 1996;

²⁰⁸In fact, one can argue that the detection of semantic coherence is a precursor for the detection of coherence between the elements of a proposition (meaningfulness) which is, in turn, a precursor for the detection of coherence between propositions that are specifically given (logical validity) or between a specifically given proposition and internally stored (and currently activated) propositions (believability) (cf. DeLong et al. 2005; Öllinger and von Müller 2017). I will come back to this point in chapter 6.

²⁰⁹This example is taken from Dechêne et al. 2010, p. 238. According to the authors the zipper was actually invented in Switzerland and thus the statement in the main text is false.

Nadarevic and Erdfelder 2014); 5) are asked to be truth-judged repeatedly (Hasher et al. 1977) or only upon final presentation (Schwartz 1982; Nadarevic and Erdfelder 2014, experiment 2); 6) repeated verbatim or paraphrased (Silva, Garcia-Marques, and Reber 2017); 7) mixed (Dechêne et al. 2009) or not mixed (Garcia-Marques, Silva, Mello, and Hansen 2019) with unrepeated statements. Additionally, similarly to the truth effect in relation to paraphrases, people tend to judge statements that contradict a previously presented statement as more likely false (“falsity effect”) than unrepeated statements (Bacon 1979), implying that a truth effect has occurred for the previously presented statement. This even holds if the previously presented sentence and the “repeated” contradiction only differ with respect to a single (crucial) word such as “Crocodiles sleep with the eyes open” and “Crocodiles sleep with the eyes closed” (Garcia-Marques, Silva, Reber, et al. 2015; Silva, Garcia-Marques, and Reber 2017).²¹⁰

In a similar fashion, subjects tend to judge statements as more likely true or false if previously exposed to related (but differing) statements which support or undermine the presented statement (Ozubko and Fugelsang 2011). Relatedly, subjects judge statements more likely true if they were previously only exposed to parts of them or to conceptually related elements, e.g. the topic of the presented statement (Begg, Armour, et al. 1985, experiment 1-2). Interestingly, the repetition-induced truth effect appears neither affected by a subject’s explicit motivation to arrive at an accurate truth judgment (Garcia-Marques, Silva, and Mello 2016), nor by individual differences in epistemically relevant traits such as a subject’s general intelligence (“cognitive ability”), preference for certainty (“need for cognitive closure”) or reliance on “intuitive” versus analytical thinking (“cognitive style”) (De keersmaecker et al. 2019). Furthermore, as mentioned earlier, it has been long assumed that the repetition-induced truth effect only occurs for ambiguous statements. Against this background, Dechêne and colleagues note:

Overall, the truth effect appears to be very robust. The only constraint seems to be that the statements have to be ambiguous, that is, participants have to be uncertain about their truth status because otherwise the statements’ truthfulness will be judged on the basis of their knowledge (Dechêne et al. 2010, p. 239).

²¹⁰Surprisingly, the same authors found that while the falsity effect with contradictory statements could be observed a few minutes after initial statement presentation, it waned after a week. In fact, after one week, the falsity effect transformed into a truth effect: the contradictory statements were perceived as more likely true than new statements. In another context this phenomenon has been called the “sleeper effect”. “According to this effect, when people receive a communication associated with a discounting cue, such as a noncredible source, they are less persuaded immediately after exposure than they are later in time” (Kumkale and Albarracín 2004, p. 143).

More recent findings, however, cast doubt on this assumption. In fact, some weaker evidence was already around for a while. For instance, Begg, Anas, et al. 1992, Brown and Nix 1996 and, more recently, Henkel and Mattson 2011 presented statements that were paired with sources declared reliable or unreliable by the experimenters. Yet this did not significantly affect the truth effect for repeated statements that were paired with unreliable sources. Memorizing and recollecting sources, however, tend to be difficult and error-prone cognitive tasks (Mitchell and Johnson 2000). This is especially true if the encoded information lacks relevance to the subject: Source monitoring “involves a processing cost which is likely to be kept to a bare minimum when the information communicated is of no possible relevance to oneself” (Sperber, Clément, et al. 2010, p. 363).

However, more recently and forcefully, Fazio, Brashier, et al. 2015 demonstrated that participants give higher truth ratings to repeated statements, even though they seem to possess actual knowledge that contradicts the statements. In the same vein, Unkelbach and Greifeneder 2018 found that even providing 100% reliable advice on the actual truth of a repeated statement *during the judgment* modulated but did not eliminate the truth effect. At first glance, this appears alarming. On second thought, however, it is perhaps not all too surprising: One can and often should be willing to question what one or others appear to know. After all, “knowledge” is fallible and should be open to new evidence.²¹¹

Now, whatever happens during the presentation of a repeated statement, might constitute new evidence that may or may not conflict with the evidence coming from e.g. one’s own beliefs or from testimony.²¹² An integration of this conflicting evidence with other information one possesses would produce exactly the observed results: a repetition-induced truth effect tempered but not eliminated by other informational sources. What we see is that our judgment is influenced by various sources of evidence and not by either one or another (cf. *ibid.*). The finding that knowledge does not eliminate the truth effect indicates that if we take ourselves to possess knowledge about something, we do rely on

²¹¹In general, from a subjective point of view, it is difficult if not impossible to be aware of *knowing anything* since, ultimately, reality has to comply. This is the basic lesson of epistemic scepticism. Subjectively and objectively well justified beliefs can justify the subject in *believing* and *claiming* that she knows. Ultimately, however, these beliefs might turn out to be false and thus fall short of *being* knowledge.

²¹²In the case of the repetition-induced truth effect, whatever is brought about by repetition, might not constitute evidence if one assumes that repetition does not correlate with truth. This latter assumption might be ecologically questionable, however (Unkelbach 2007, p. 229; Reber and Unkelbach 2010; Herzog and Hertwig 2013). More on this later.

it—but we do not rely on it exclusively. Knowledge is *one* but not the only variable. There is a basis for truth judgment over and above knowledge and this basis is used for truth judgments in cases with *and* without knowledge.²¹³

That knowledge does influence our truth judgments is obvious and is shown by the mentioned studies. This is also evident in findings demonstrating that plausibility and implausibility, equivalent to what Thompson and Morsanyi called “believability” above, mediates the repetition-based truth effect (Pennycook et al. 2018; Fazio, Rand, et al. 2019). In fact, in their first experiment Pennycook et al. 2018 observed *no* effects of repetition on truth ratings for plausible or implausible statements, i.e. statements whose truth value subjects can actually assess based on their knowledge.

There is a catch, however: the statements for which no truth effect was observed were *extremely* plausible (e.g. “there are more than fifty stars in the universe”) or implausible (e.g. “the earth is a perfect square”). In contrast, when subjects were tested on real news headlines and highly (but not extremely) implausible fake news headlines²¹⁴, the truth effect emerged for both kinds of headlines. More strikingly yet, the truth effect even showed for fake news headlines that were 1) accompanied by the warning “Disputed by 3rd Party Fact-Checkers” during each presentation and 2) discordant with the subject’s political orientation. In fact, based on theoretical and empirical grounds, Fazio, Rand, et al. 2019 argue that repetition *does equally* increase perceived truth for ambiguous as well as for unambiguous statements, i.e. also for extremely plausible and implausible ones. What happens is that for ambiguous statements the increase in *perceived truth* exerts the largest *observable* truth effect while in extremely unambiguous cases one fails to *observe* the effect of this increase in perceived truth on the judgment, i.e. the behavioural truth effect:

²¹³ Assuming that the mentioned basis is one single state, an alternative idea might be that it is primarily this state itself (and not the judgment based on it) that integrates various sources of information, including one’s knowledge. I am sympathetic to this idea and will return to it in chapter 8.

²¹⁴ Both types of headlines were sourced from real online sources. Examples of real news headlines:

- “Majority of Americans Say Trump Can Keep Businesses, Poll Shows” (Pro-Republican, bloomberg.com)
- “The Small Businesses Near Trump Tower Are Experiencing a Miniature Recession” (Pro-Democrat, slate.com)

Examples of fake news headlines:

- “Donald Trump Sent His Own Plane to Transport 200 Stranded Marines” (Pro-Republican, uconservative.com)
- “Trump on Revamping the Military: We’re Bringing Back the Draft” (Pro-Democrat, realnews-rightnow.com)

[T]he increase is masked by the extreme disbelief. That is, the statements are so disbelieved initially that even with an increase in belief due to repetition they are still rated as definitely false. [...] [R]epetition increases belief in all statements equally, regardless of their plausibility. However, there is an important difference between this internal mechanism (equal increase across plausibility) and the observable effect. The observable effect of repetition on truth ratings is greatest for items near the midpoint of perceived truth, and small or nonexistent for items at the extremes. While repetition effects are difficult to observe for very high and very low levels of perceived truth, our results suggest that repetition increases participants' internal representation of truth equally for all statements. (Fazio, Rand, et al. 2019, pp. 2, 6)

So again: knowledge matters, but not only. Repetition affects “truth judgments independent of prior knowledge” but people “can, and often do, judge the truth of a statement based on their prior knowledge” (ibid., p. 5).

To sum up: the truth effect describes the heightened tendency to judge a statement as (more likely) true as a consequence of manipulating various variables. This effect is usually studied in relation to ambiguous statements such as “the zipper was invented in Norway” and the most prominent condition under which it occurs is repetition: subjects judge a statement as (more likely) true if they have encountered it before. This repetition-based truth effect is robust across many conditions. In fact, it even occurs if the statements are *not* ambiguous. That is, the truth effect can even be observed when subjects have either knowledge or perfectly reliable testimony on the truth status of a statement.

At first glance, this appears alarming. On second thought, however, it is perhaps not all too surprising: One can and often should be willing to question what one or others appear to know. After all, “knowledge” is fallible and should be open for new evidence. Now, whatever happens during the presentation of a repeated statement, might constitute new evidence that may or may not conflict with the evidence sourced from e.g. one's own beliefs or from testimony. What we observe is that our judgment is influenced by various sources of evidence and not by either one or another. The finding that knowledge does not eliminate the truth effect indicates that if we take ourselves to possess knowledge about something, we do rely on it—but we do not rely on it exclusively. Knowledge is *one* but not the only variable. There is a basis for truth judgment over and above knowledge and this basis is used for truth judgments in cases with *and* without knowledge.

7.3.5. Truth, Fluency and Affect

Now you might wonder: *why am I telling you all this?* It has to do with a familiar construct which is broadly accepted to provide a general theory of the truth effect, not only the one based on repetition. You might have noticed that I was withholding from you other truth-effect-inducing variables than repetition. Here they are: statements that rhyme (e.g. “birds of a feather flock together”) are judged more likely true than those that do not (e.g. “birds of a feather flock conjointly”) (McGlone and Tofighbakhsh 2000). Statements formulated in a concrete language (e.g. “the Naab flows into the Danube”) are found more likely true than those formulated choosing more abstract words (e.g. “the Naab is a confluent of the Danube”) (Hansen and Wänke 2010). Statements presented together with a picture of their grammatical subject (e.g. “macadamia nuts are in the same evolutionary family as peaches” with a photo of macadamia nuts) are judged more likely to be true (Newman et al. 2012, experiment 3; see also Kelley and Lindsay 1993). Finally, higher visual contrast (e.g. colour contrast between font and background) or readability of a statement makes it more likely to be judged true (Reber and Schwarz 1999; Scholl et al. 2014). What is more, statements uttered in accented speech are rated as *less* true (Lev-Ari and Keysar 2010).

What do repetition and these factors have in common? Or: what do they facilitate? The remote associate here is processing fluency. It is not repetition, rhyming, concreteness etc. *per se* that leads to the truth effect. Rather, all these factors make the various cognitive processes (perception, comprehension, encoding etc.) that deal with the statements more fluent (Unkelbach and Greifeneder 2013b). In fact, Unkelbach and Stahl 2009 show that apart from repetition, the factual truth of statements leads to processing fluency as well and that this “fluency due to truth” is a primary determinant of truth judgments. As a consequence, there is an “experienced equivalence of repetition and truth” and “people cannot distinguish between fluency due to truth and fluency due to repetition” (Unkelbach, Bayer, et al. 2011b, p. 598).

Furthermore, the truth effect is more pronounced if the processing fluency is *discrepant* (Hansen, Dechêne, et al. 2008; Dechêne et al. 2009; Garcia-Marques, Silva, Mello, and Hansen 2019, see also section 5.3.2.1). Processing fluency can be discrepant in several complementary ways: It can be discrepant relatively to 1) the rest of the currently ongoing processing background (i.e. a specific process can be fast and therefore fluent relative to *other* ongoing processes) 2) an expectation concerning the speed of the specific

processing formed outside of the present context or 3) formed in the present context (Hansen and Wänke 2013; Unkelbach and Greifeneder 2013b).

Now, assuming that processing fluency plays a major role, what happens in the truth effect studies? Put differently: What exactly makes people judge fluently processed statements as more likely true? I cannot see how it could be the processing fluency *itself* since processing fluency just means that some process is executed relatively fast—there seems to be no direct causal connection from that to a judgment (except that the judgment is made faster). Thus, what we have to look at are the *consequences* of fluency.²¹⁵ As I see it, there are (at least) four possibilities.

1. Processing fluency leads to a positive affective experience that influences the conscious truth judgment.
2. Processing fluency leads to an unconscious positive affective nudge that influences the conscious truth judgment.
3. Processing fluency leads to a non-affective experience of fluency that influences the conscious truth judgment.
4. Processing fluency leads to a non-affective nudge that influences the conscious truth judgment.

On the face of it, these possibilities are not mutually exclusive but complementary. Sometimes our truth judgments might be influenced by affective experiences, sometimes by affective nudges, sometimes by experiences of fluency and sometimes by non-affective nudges. In fact, it is conceivable that at times these things co-occur and each on their own influences truth judgments. However, based on the reasoning of chapter 4, it is questionable whether there are such things as non-affective experiences of fluency. There are good reasons to think that in cases where fluency causes conscious states, these states are affective experiences. This leaves still the affective and non-affective nudges on the table. As I said earlier, that there are (fluency-based) affective epistemic nudges is plausible. What about non-affective nudges? Here something similar can be said as for experiences of fluency. Note that we are talking about *results of* processing fluency. We have seen in chapter 4 that fluency leads to affective states, whether conscious or unconscious and that inducing subtle positive affect leads to essentially the same behavioural consequences in the form of judgments as inducing fluency. We cannot

²¹⁵Alternatively we could look at the fluent process itself or at the consequences of the fluent process. I will do that in chapter 8.

exclude that there are also non-affective states caused by fluency. It is, however, unclear what role there is to play for them if fluency and affect are already in the cast.

Based on these reflections I submit that, in fact, it is not processing fluency *per se* that is responsible for the truth effect. Rather, it is the *result* of processing fluency: affect. Based on the present data it is hard to say whether this affect is always conscious—there might be unconscious valence (or no valence at all) that establishes an observable behavioural bias (without a phenomenal state). At the same time, I think it is safe to say that, as Topolinski and colleagues have shown, sometimes this affect will be felt, and this will then be a potent phenomenal basis for the subject’s truth judgment. In other words, I suggest that what happens in cases when we experience an increase in “perceived truth” is the occurrence of an epistemic feeling: a feeling of rightness towards a statement or proposition, or more specifically a “feeling of truth” (Unkelbach, Bayer, et al. 2011b; Newman et al. 2012).

Do we have any empirical evidence for this proposal? Unfortunately, in contrast to e.g. Topolinski’s work on semantic coherence, this proposal has only been little explored. To my knowledge, there are only two studies that have explicitly investigated the role of affect in the context of the truth effect (Unkelbach, Bayer, et al. 2011b; Koch and Forgas 2012). And these yield inconclusive results.

Due to the assumption that it is fluency and not affect that is behind the truth effect, Koch and Forgas 2012 were interested in whether mood modulates the truth effect *via* modulating the influence of fluency on truth judgments. And in fact, they find that the truth effect is present in a neutral and a positive mood but is eliminated in negative mood. This they take as evidence that “positive mood promotes, and negative mood eliminates people’s reliance on processing fluency as an indicator of truth” (ibid., p. 485).²¹⁶ From an affect-oriented perspective, such a finding seems to be expected. I already mentioned in chapter 4 that negative mood, as well as depression, dampen epistemic feelings and influence one’s performance in tasks capitalizing on them (e.g. Sweklej et al. 2014). What we seem to observe in the present experiments is that feelings of truth (that are in this context uninformative about truth) are compromised by negative mood. Consequently, the truth effect vanishes. This is reminiscent of the drop in semantic coherence detection performance to chance level. On the other hand, a positive mood usually improves performances that capitalise on epistemic feelings

²¹⁶With regard to their findings I think “promotes” in the case of positive mood is a slightly misleading word choice. In their data, people in a neutral mood were actually numerically more likely to judge fluent statements as true.

(Balas et al. 2012). However, the present experiments had no accuracy conditions so that successful performance cannot be meaningfully defined and assessed.

Be that as it may, the primary question of the experimenters was not whether affect influences truth judgments. In fact, in relation to this question, the data is quite convoluted. This is due to the three different experimental manipulations of which each might have an affective pull, something the experimenters have not considered. They manipulated the mood of the participants (positive, neutral, negative) and presented statements that had a valence (positive, neutral, negative) and were perceptually fluent or disfluent (via visual contrast). The experimenters found no main effects for mood or fluency but only interaction between the two, namely that mood affected whether a fluency-induced truth effect could be observed (for positive and neutral mood) or not (for negative mood) (Koch and Forgas 2012, p. 483; but see Garcia-Marques, Mackie, et al. 2004, experiment 3). The sentence valence was manipulated in order to check whether statements that had a mood-congruent valence would be more likely to be judged true. This was found for the combination of positive statements and positive mood but not for negative mood.

On the other hand, we do not have a good idea of the more general effect of sentence valence since the authors did not investigate whether positive/negative statements in general are more/less likely to be judged true. Eventually, I think that any effect that affect could have had would have been concealed by different combinations of affect-relevant variables. For instance, a subject in a positive mood, confronted with a negative but fluent statement would have some positive affect from mood and fluency and some negative affect from sentence valence. It is possible that the mood manipulation had the strongest affective force, the sentence valence the second strongest, and the fluency manipulation the third strongest, in sum effectively increasing the likelihood of the sentence to be judged true. Though, ultimately, there is no way to prise the respective influences of these factors apart.

Taking additionally into account 1) that the authors asked about the mood of the participants before asking for truth judgments and 2) that the sentences were obviously valenced, makes it conceivable that the experiential effects of mood and sentence valence would've been discounted by the subjects when making truth judgments, i.e. rather than being taken to represent truth the affect would have been correctly attributed to one's cued mood and the apparent positive/negative content of the statements. Perhaps this partly explains why subjects in a positive mood were not more liberal with their

truth judgments than subjects in a neutral mood. In sum, however, the interpretation of these findings is difficult with regard to the question whether affect, due to fluency or something else, influences truth judgments.

Unkelbach, Bayer, et al. 2011b conducted the only direct test of the hypothesis that affect and not (only) fluency might be responsible for the truth effect. In their first experiment, they showed their subjects matched statements with positive or negative valence, e.g. “Each year over 100 people succeed (die) attempting to climb the 6163 m of Mount Chunasla” or “The divorce rate in Grenada is lower (higher) than in the rest of Spain”.²¹⁷ They found that positive/negative statements did not tend to be rated more/less likely true than negative/positive ones *across subjects*, i.e. positive statements in general did not tend to be judged more likely true than negative ones. Yet, *within subjects* they found a positive/negative correlation between statement positivity/negativity and truth ratings. In other words, when comparing the truth ratings of positive and negative statements of individual subjects, positive statements tended to be judged more likely true than negative ones. However, utilizing the same positive and negative statements together with neutral ones in a truth-effect-setup in experiment 2 and 3 the authors failed to find an effect of statement positivity/negativity while observing the typical repetition-based truth effect:

The present experiments are motivated by the observation that positivity could substitute processing fluency as the explanatory construct for the truth effect. [...] [W]e did not find systematic evidence for positivity influences. There was no amplification of the truth effect for positive statements [...] Accordingly, we conclude that the repetition-based truth effect is caused by processing fluency, over and above accompanying positive experiences. (ibid., p. 601)

I think it would be premature to dismiss the hypothesis that the truth effect is at least sometimes mediated by affect. As the authors themselves write: “A caveat is that this conclusion relies on null results; maybe we did not operationalise positivity adequately or the experimental context was not suited to capture positivity influences” (ibid., p. 601). Indeed, if we compare the rather obvious affect manipulation in the form of positive/negative statement content to the more subtle ones as e.g. utilised by Topolinski, it might appear as not optimal.

²¹⁷Note, there were no neutral statements in experiment 1.

For instance, the valence of the sentences was apparent to the subjects so that the possibly triggered affect might have been (correctly) attributed to the content of the statements rather than to them being true/false. As noted by Jacoby and Whitehouse for perceptual fluency in producing feelings of familiarity: “For the attribution process, the presence of alternatives to familiarity as a plausible source of effects is important” (Jacoby and Whitehouse 1989, p. 127). This is all the more likely for affect (that is a consequence of fluency): if (more plausible) alternatives than the truth of the statements are available in order to explain the affect, then the affect will be attributed to them. This seems to be the case for the statements used by Unkelbach and colleagues: instead of being attributable to the truth of the statements, the affect is correctly attributed to the affective valence of the content of the statements. In choosing the experimental setup it is important to keep in mind that the meaning of affect is highly context-sensitive. I suggest that a subliminal affective priming paradigm might have done a better job.

Apart from the issues with the specific studies here, there are numerous findings that point towards a role of affect in truth judgments. The prominent mere exposure effect describes the tendency of repeatedly presented (initially neutral) stimuli to be liked more (Zajonc 1968; Bornstein 1989). Importantly, the mere exposure effect has been explained by reference to repetition-induced processing fluency (e.g. Bornstein and D’Agostino 1994). Now, the more general mere exposure effect and the underlying mechanisms are highly reminiscent of the more confined statement-directed truth effect and its underlying mechanisms: The truth effect is most prominently induced by repetition and ultimately explained by processing fluency. For instance, a mere-exposure-like effect on liking judgments and a truth effect on truth judgments can also be brought about by increasing non-repetition-based conceptual fluency (Winkielman, Schwarz, et al. 2003, pp. 204-206; Silva, Garcia-Marques, and Reber 2017).

In contrast to the truth effect literature, work on mere exposure has made an additional theoretical step by proposing that it is not ultimately the repetition-induced fluency that is responsible for the mere exposure effect but the result of fluency: affect (Harmon-Jones and Allen 2001; Winkielman, Schwarz, et al. 2003; Fang et al. 2007).

[W]e suggest that the mere-exposure effect is driven by the impact of stimulus repetition on processing fluency [...] [W]e propose that the positive hedonic marking of the fluency signal is the crucial ingredient, consistent with the accumulating evidence that high fluency elicits positive affect. (Winkielman, Schwarz, et al. 2003, p. 204)

I suggest that, after we have clarified the affective nature of epistemic feelings in chapter 4, something similar applies to the truth effect. In fact, the mere exposure effect and the repetition-induced truth effect are likely two sides of the same coin: fluency-induced positive affect that in one context expresses itself in the (likely primitive) form of liking judgments and in another in the form of truth judgments. I conjecture that one could potentially even drop the fluency from the picture and replace it with whatever manipulation one chooses that reliably triggers subtle, transient affect.

I think this proposal receives additional support if we think back to the results of Morsanyi and colleagues further above. Recall that not only logical validity led to conclusions being liked more but so did *believability*. Believability is, in fact, just another word for coherence with one's background beliefs. In other words: believability is equivalent to knowledge-based perceived truth, discussed in the truth effect literature.

Now, Thompson and colleagues showed that believability modulates logical validity judgments whereas Morsanyi and colleagues found that believable conclusions are liked more. Similar to Unkelbach and Stahl 2009, I propose that believability expresses itself in a similar way as repetition, visual clarity or conceptual relatedness: in the form of processing fluency (between the sources of which subjects cannot distinguish). This, in turn, leads to positive affect: As a consequence, believable, repeated, visually clear etc. statements are liked more or judged true.

Additionally, I want to repeat a point I made above: Semantic coherence between words and coherence between different propositional elements of information in the form of logical validity, believability *and truth* bear some resemblance as properties to each other. It is plausible that all of them lead to (fluency-induced) affect and that this affect can be the basis of judgments about corresponding properties. Then again, we will need further empirical work to see whether this is actually the case. As I see it, the case is still open.

Note that the present proposal does not only imply that one should observe the truth effect if one (subtly) manipulates affect. It also suggests that the truth effect should be modulated by whatever (subtly) causes positive affect, processing fluency or something else (e.g. coherence). I will come back to some potential underlying mechanisms that can trigger affect or be fluent in the next chapter. For now, however, I want to analyse FOR/Ws further with the resources provided in the previous chapters.

To sum up: in this section, I introduced the dominant theory about the common cause

of the truth effect across manipulations of different variables such as previous exposure, readability and rhyming. This theory states that the proximal cause of the truth effect is a familiar construct: processing fluency. Now, we have previously discussed processing fluency as the cause of epistemic feelings and its relation to affect quite extensively (section 5.3.1.5, 5.3.2.1, 6.3 and 6.4). These previous reflections have brought to light that what actually matters for epistemic feelings is not so much processing fluency *per se* but rather its result: affect. Following this line of thought, I argued that, in fact, it is not processing fluency *per se* that is responsible for the truth effect. Rather, it is the *result* of processing fluency: affect. Based on this I tried to make the case that what we observe in the truth effect are actually the workings of two epistemic feelings which are proposition-directed varieties of FOR/Ws: feelings of truth and feelings of falsity.

7.4. FOR/Ws as Affective Experiences and Intuitions as FOT/Fs

In the following, I want to provide a general analysis of the FOR/W drawing on the resources introduced in the previous chapters. In the course of this analysis, I will delineate FOT/Fs among FOR/Ws as those FOR/Ws that take propositions as their particular objects and represent them as right or wrong, amounting to an evaluation of the propositions as true or false. It will turn out that the features of FOT/Fs are identical with the features of intuitions. Thus, I will conclude that intuitions are identical to FOT/Fs.

I hope to have established that FOR/W can be directed at all kinds of objects. As we have seen they can concern particular conclusions and their rightness or wrongness (i.e. logically valid, believable) or they can concern a self-generated response to a specific problem (e.g. Thompson, Prowse Turner, et al. 2011; Thompson and Morsanyi 2012; Thompson, Turner, et al. 2013; Morsanyi and Handley 2012; Trippas et al. 2016).

We experience a FOR/W towards a certain particular object delivered by the base of the feeling. The formal objects of these feelings is determinable rightness and wrongness respectively. These feelings will usually attract us to or repel us from their intentional object and motivate approach or avoidance behaviour with respect to it. The determinate kind of rightness or wrongness that such a feeling represents depends on its particular object. If the feeling is directed onto a particular, then this particular will appear (in)adequate (or apt or fitting). In this case, the more specific kind of rightness

or wrongness that is the formal object of the respective feelings is best described as kinds of adequacy and inadequacy. If the intentional object is a conclusion, then the rightness in question will be logical validity (or plausibility). If the intentional object is a sentence, then the rightness in question will be (specific kinds of) linguistic well-formedness.²¹⁸ Alternatively, the feeling of rightness or wrongness can be directed onto a proposition. In this case, the determinate rightness or wrongness that is the formal object of the respective feelings is best described as truth and falsity.²¹⁹

I want to call the subclass of FOR/Ws that have propositions as their intentional objects *feelings of truth* (FOTs) and *feelings of falsity* (FOFs, from now on). As affective experiences in general, these feelings are motivational. More specifically: they motivate us to assent to or to withdraw assent from a certain proposition, probably motivating some additional inquiring behaviour. That FOT/Fs have this kind of *pushiness* is one of the things that the previously reviewed empirical work has abundantly demonstrated (see section 7.3).

FOT/Fs can potentially take any mental states with propositional content as their bases. When it comes to philosophical practice, a frequent kind of base would presumably be thought or, more specifically: propositional, non-assertive thought which goes by various names such as “entertaining”, “grasping”, “supposition”, “conceiving” or “propositional imagining” (see e.g. Williamson 2007). There is some debate on the continuity between the kind of state that some like to call propositional imagining and imagining proper (Kind 2016, pp. 3 sq.; Arcangeli 2017, pp. 2 sq.). For my purposes, it does not matter whether propositional imaginings are continuous with imaginings proper or not. I thus

²¹⁸Against this background, the *Grammaticality* intuition in chapter 2 might be best understood as a FOW directed at *the sentence* “the boy the man the girl saw chased fled”, not a FOR directed at the proposition “The sentence ‘the boy the man the girl saw chased fled’ is ungrammatical”. This FOW signals linguistic inaptness or, more specifically, ungrammaticality (see also footnote 33).

²¹⁹Actually, it is the whole context that determines the kind of rightness and wrongness. The context includes the particular object but comprises yet other factors such as task requirements and instructions (external factors) and the subject’s expectations, intentions and goals (internal factors). Giving it an active spin, philosophers like to speak of imaginative or cognitive projects or endeavours in this connection (e.g. Currie and Ravenscroft 2002; Chudnoff 2013c, pp. 49, 63 sq.; García-Carpintero 2013; Jackson 2016, pp. 43 sq.; Stock 2017). Now the context that a subject finds herself in and the project that a subject currently pursues will play not only a pivotal role in determining the specific rightness or wrongness of an occurring FOR/W but whether a FOR/W occurs at all. For the sake of simplicity, I will omit this complication and speak as if the particular object determines the kind of rightness. Additionally, I acknowledge that some particular objects can be assessed with respect to various kinds of rightness. Thus, the kind of particular object taken by itself will often be underdetermined in respect to the kind of rightness that applies to them. For instance, a conclusion can be assessed as being logically valid or as plausible. Similarly, propositions such as “Women are incapable of rational thought” or “Torturing kittens for fun is OK” are not only assessable as factually false but also as morally wrong (see footnote 33).

invite the reader to replace what I call propositional imagining with whatever is the most palatable label for his or her commitments. In fact, if you're generally unhappy with the imagining label (and not only with the *propositional* imagining label) pick whatever occurrent propositional attitude you like. As already said: FOT/Fs can take *any* mental states with propositional content as their bases.

Note further that what matters for me here about propositional imaginings is that they are non-assertive and have propositional content.²²⁰ In the literature, the contrast between propositional imagining and imagining proper is often drawn by saying that the former is belief-like or “cognitive” while the latter is perception-like or “sensory”. However, assuming that perception can have propositional content (Wringe 2015) and that all imaginings are non-assertive, this does not necessarily map onto the distinction between belief-like and perception-like imagining since the latter is non-assertive and can be propositional as well. The contrastive case for propositional imaginings becomes clearer in the next distinction: Propositional imaginings are often distinguished from objectual imaginings that present us with individual objects rather than propositions (Yablo 1993).

I will be concerned primarily with propositional imaginings here.²²¹ ²²² Bayne and Pacherie further helpfully distinguish between three different forms of propositional imaginings: simple, counterfactual and indicative propositional imaginings:

In one sense, to imagine a proposition is simply to entertain it, independently of any attitude to the truth-value of the proposition in question. One might imagine that the population of Nepal has doubled in the last 25 years without having any attitude towards the truth-value of this thought. Call

²²⁰In fact, the non-assertive bit might ultimately turn out to be unimportant. Arguably, there is nothing that prevents us from having FOT/Fs towards propositions that are the contents of *assertive* propositional attitudes. Relatedly, it seems that some already endorsed propositions might sometimes take quite a bit of time before we not only judge but really *feel* them true/false.

²²¹However, if the unitary view about FOR/Ws is roughly on the right track, there is no difference in kind between FOR/Ws directed onto objectual contents and FOR/Ws directed onto propositional contents, i.e. FOT/Fs (see section 7.3.1).

²²²We can construct parallel accounts of propositional and objectual intuitions as FOR/Ws with either propositional or objectual imaginings as their bases. Furthermore, we can speculate about the possible interplays and relations of the two. Chudnoff, for instance, contends that intuitions mark not only a certain propositional content as true or false but that they also present us with the truth-maker for the proposition they assert. Such a double phenomenology can be accommodated by the co-occurrence of a FOT/F with a base that has propositional content and a FOR/W with a base that has objectual content and whose objectual content directly bears on the propositional content's truth. Perhaps also a *single* FOR can take both, a state with objectual content *and* a state with propositional content as its base; or it can take one single state which has propositional *and* objectual content as its base..

this *simple imagination*. One can also entertain in imagination a scenario in which a proposition holds while believing of this scenario that it is not actual. This is the sense of imagination at work in our understanding of fiction. Call this *counterfactual imagination*. There is yet a third form of propositional imagining, according to which to imagine P is to have some inclination—however slight—to think that P is the case. Call this *indicative imagination*. It is this sense of imagination that is at work in such claims as, ‘I can’t imagine that the Mafia killed JFK’. Someone who says this is not meaning to deny that she can *entertain* the proposition <the Mafia killed JFK>, rather, she means to communicate the thought that it is extremely unlikely that the Mafia killed JFK. Indicative imagination and belief seem to be on a continuum, insofar as both are attitudes to the way the world actually is; simple and counterfactual imagination, by contrast, are not on a continuum with belief. (Bayne and Pacherie 2005, p. 167)

I will be concerned with simple propositional imaginings here.²²³ Consider Stock’s minimal characterization which I think captures the essence of simple propositional imagining:

Propositional imagining “involves, minimally, taking a certain attitude of ‘thinking that’ to a given content that one reads [or hears etc.], without a commitment to its truth, any automatic integration with one’s belief set, or any automatic relevance to one’s behaviour. It may not involve a substantial phenomenological aspect. It can be largely passive and involve little deliberate activity on part of the reader [or listener etc.] other than reading

²²³But a brief comment on the other two kinds of propositional imaginings is in order. In the case of counterfactual imaginings we entertain a proposition which we believe to be false. Nevertheless, how is it that we are aware that we engage in counterfactual rather than simple imagining? One possible answer is: we aren’t. It is not uncommon that a (dispositional) belief and a currently entertained propositional content fail to connect for some reason (cf. Fazio, Brashier, et al. 2015). In this case, on a personal level, we are actually having a simple rather than a counterfactual imagining. Another possibility is that, as an expression of our belief, we judge the proposition we’re entertaining to be false but, for the sake of our imaginative endeavour, imagine it to be true. A third possibility is that we are made aware of the counterfactuality of our imagining by experiencing a FOF towards the proposition upon simply entertaining it. In fact, this FOF sometimes might be the basis for the judgment in the second case. In a similar vein, I believe that what Bayne and Pacherie call indicative imagining might in some cases be, in fact, a combination of two things: a simple imagining that, say, *the Mafia killed JFK* and a low-intensity FOF towards it. The corresponding utterance that *I can’t imagine that the Mafia killed JFK* is then the expression of a judgment based on the combination of a simple propositional imagining and a low intensity FOF. Alternatively, if one is susceptible to the Mafia-killed-JFK-conspiracy-theory, one could have a (low-intensity) FOT upon simply entertaining *the Mafia killed JFK* and express it by saying that *I can imagine that the Mafia killed JFK*.

[or hearing etc.] and processing of lines of text [or heard utterances etc].”
(Stock 2017, p. 27)

Philosophers in their inquiry think about and imagine all kinds of things. The contents of their imaginings will sometimes be more perception-like and sometimes more belief-like. In imagining certain propositions mental imagery is often recruited. This mental imagery can take the form of sentences, pictures, sounds etc. The contents can be concrete or abstract. This accounts for the fact that, not unlike perceptual experiences and intuitions, FOT/Fs can be about concrete things. Still, they can also be about very abstract subject matters, unlike perceptual experiences and more like beliefs, judgments and intuitions. And this they can by having (propositional) imaginings as their bases whose contents can range from the concrete to the abstract, linguistic, modal, moral, logical etc.

Imaginings by themselves are not assertive and not committal (Kind 2016, p. 3). The FOT/Fs that co-occur with them and take them as bases can be said to “flip the assertiveness switch”: In virtue of the presence of the feeling, the imagining gets the gloss of being about something true or false. Another noteworthy feature of imaginings is that they can be deliberate as well as spontaneous. At times “we can choose whether and how to engage in the imaginative episode” (Spaulding 2016, p. 210) and at times imaginings “can be spontaneous, nondeliberate passive experience rather than something one *does*” (Walton 1994, p. 48). I can ask you to imagine the first three prime numbers and you will be able to voluntarily comply with the request. However, sometimes imaginings come unbidden. This accounts for FOT/Fs in different contexts. For example, we might account for FOT/Fs had in the context of a presently circumscribed inquiry such as thinking through some thought experiment. But also for FOT/Fs that occur suddenly and spontaneously, apparently unrelated to the present situation, that is in a context where we would speak of sudden insight.

Furthermore, imaginings have the quality of being content-gradable. Imagining that there is a red house in front of you or that kittens are being tortured can be more or less vivid, more or less clear and distinct, varying with the context and purpose of your imagining as well as your capabilities to imagine. Note that, additionally to the “native” content-gradeability of imaginings, the presence of a feeling will enhance the vividness of the imagined content.²²⁴

²²⁴My guess is that those who prefer to speak of intuitions as inclinations (e.g. Sosa 2007c) have cases in mind where the aspect of content-gradeability and perhaps other phenomenal features are less salient while those that prefer to speak of intuitions as experiences (e.g. Bengson 2015) heed the instances

The fact that the feeling will come to “inherit” properties of imaginings such as content-gradeability is a result of *base property mirroring*, introduced in section 4.4.4: it describes the idea that an affective experience will come to mirror certain dynamical, intentional and phenomenal properties of its base. Now, such a kind of influence is not a one-way street: the feeling modifies the content delivered by the imagining in characteristic ways as well. It does not only “flip the assertiveness switch” but also makes the imaginative content more vivid, thereby contributing to a feeling-specific increase in content-gradeability (see section 4.5).

Since I have mostly talked about propositional imaginings, one might wonder here: are propositional imaginings “natively” content-gradeable? Let us differentiate a bit: As mentioned earlier, propositional imaginings can be perception-like and they can be belief-like. In the former case, there is no reason to deny native content-gradeability. What about belief-like propositional imaginings? This will, of course, depend on one’s view. One might argue that belief-like imaginings are actually perception-like in that they are also based on sensory imagery—but not of a vivid visual kind but of, say, an auditory or linguistic kind (Jackendoff 1987). This would perhaps secure content-gradeability again.

One might want to reject this view of belief-like imaginings. Perhaps, then, I would have to accept that for this subset of imaginings I cannot secure content-gradeability for FOT/Fs via base property mirroring. In this case, the potential vividness effect of these feelings would presumably peter out since it would lack a base property that could be its target. However, I conjecture that some people would still be willing to talk of the content of the imagining becoming “clear and distinct”. This impression would, however, not be based on the vividness effect of FOT/Fs but on the independently supplied *pushiness*. FOT/Fs do not only make content appear more vivid, they also motivate us to accept or reject it. I find it plausible that, sometimes, when we speak of some content seeming “clear and distinct”, what we actually mean is that we feel “attracted” to it. “Clarity” can thus at times be couched in terms of how strongly attracted one feels towards a certain proposition—one “firmly grasps” the proposition in this sense. In the end, I do not want to commit too firmly to content gradeability as an essential feature of FOT/Fs. What matters is that we manage to explain that we sometimes have FOT/Fs that *have* content-gradeability. Against the background of the present account, there can be FOT/Fs with content gradeability, namely those that take a base with content gradeability (and can thus be targeted by the affective vividness effect) and there can be

where content-gradeability and perhaps other phenomenal features are more salient.

FOT/Fs that have no content gradeability but only pushieness (gradeability) because they latch onto a propositional base that has no content-gradeability.

Finally, the idea of unconscious imaginings is getting traction lately (e.g. Sullivan-Bissett 2018). As demonstrated by phenomena such as affective priming and blindfright, feelings can certainly take unconscious perceptions as their bases (Scarantino 2010). So if there are unconscious imaginings, I can see no reason why conscious feelings cannot take such unconscious imaginings as their bases. Unconscious imaginings provide us with a straightforward way to explain cases where the contents of FOT/Fs are not only somehow indeterminate but are plainly unavailable to consciousness. This does not mean that the FOT/Fs are not intentional, neither does it mean that they are unconscious (see end section 4.4.3 and beginning of section 4.4.4). Something is right or something is wrong and—but we don't know what. Nevertheless, we are usually able to point into the general direction: something about *this argument*, something about what *he said*. What is unconscious in these cases is the imagining and its content. In virtue of taking this imagining as a base, the content of the FOT/Fs will appear inexplicit (see also Bengson's feature of inexplicitness in section 3.4.2 and FOK inexplicitness in section 6.5.2).

So, as for other (epistemic) feelings, the contents of FOT/Fs, i.e. *what* is felt as true or false, can be conscious to various degrees. In cases where the amount of consciously available content is insufficient to form a proposition that can be consciously assented to, the feeling might motivate attempts to articulate the content further. This is one of the ways in which FOT/Fs motivate inquiry (cf. Vaidya 2010, p. 399).

In any case, the subject will often be unaware of *why* exactly a (conscious) content feels true or false to her. Feelings in general are representing something as being a certain way but they are by themselves opaque about the reasons and causes for their occurrence (cf. Haidt 2001; Hauser et al. 2007). It is often this opacity paired with the subject's inability to integrate their occurrence with the current contents of consciousness that makes epistemic feelings stand out phenomenally. On the other hand, subtle epistemic feelings (and their influence) tend to go little noticed where congruent reasons and causes can easily be identified.

Finally, note that FOT/Fs have a phenomenal valence, propositions as their particular objects and truth and falsity as their formal objects. As such, they are psychological expressions of truth and falsity (more on the psychological mechanisms in chapter 8). However, that does not mean that in them truth and falsity *appears* psychological or subjective to the subject—they *appear* objective: as truth or falsity (see section 6.4).

Just as in a visual experience of redness the redness does not appear to us as psychological or subjective, or in our disgusts disgustingness (or contagion) does not appear as psychological or subjective (at least at first glance).

FOT/Fs, then, exhibit the following features: they feel positive or negative in relation to the truth or falsity of a proposition. This is just another way of saying that they have *phenomenal epistemic valence* and *assertiveness*.²²⁵ Taking a proposition in a particular context and throwing phenomenal valence into the mix gives us these *phenomenal* properties. Truth and falsity seem to emerge as formal objects that are phenomenally represented by the feeling—in the form of *phenomenal* truth or falsity. This appears like a kind of phenomenal intentionality on the part of FOT/Fs. How plausible is this? I think this depends on the answer to the following question: When Rita is afraid of the ferocious bear, does the bear appear as fearsome to her because of the phenomenal qualities of her fear? In other words: Does the bear appear fearsome because her fear phenomenally represents the bear as fearsome?²²⁶ It seems obvious that it does. This phenomenal datum seems to explain, among other things, why Rita takes to her heels. Phenomenal valence seems to ground the fearsomeness (or danger) as a phenomenally represented property. Without it, fear would not be frightening in the same way as valence-less pains are not painful (see section 4.2 and 4.3). Now, for FOT/Fs phenomenal valence plays the same role for truth and falsity as it does for the formal objects of other affective experiences: It grounds them phenomenally. Without it, FOT/Fs would presumably not be phenomenally assertive and epistemically valenced. FOT/Fs phenomenally “encode”, “localize” or “embed” truth and falsity (their determinate feeling-specific property) within the contents of their bases (e.g. imaginings) and motivate us to behave accordingly. They, in other words, phenomenally represent or evaluate their propositions as true or false.

Here is the same point from an angle that thinks the idea back from the way FOT/Fs motivate us to behave (cf. Weiss 2016, pp. 44-46). Fears, for instance, motivate flight and avoidance behaviour, among other things. What do FOT/Fs do? In our discussion of the truth effect we have seen that FOT/Fs lead us to judge statements as true or false. In other words, FOT/Fs seem to have specific functional roles: we use them to infer the truth or falsity of a statement. Against the background of inferential or

²²⁵Considering other epistemic feelings it might be more accurate to speak of a *specific* phenomenal epistemic valence, namely phenomenal valence in relation to truth or falsity. FOFs and FOKs can be said to have other phenomenal epistemic valences in relation to familiarity or knowing.

²²⁶Or more technically: Her fear phenomenally represents the bear as bad in the determinate form of being fearsome (cf. Teroni 2018).

conceptual role semantics, this seems to secure the idea that FOT/Fs are intentional and that truth and falsity somehow figures in their intentionality (e.g. Block 1986). As I said previously, different kinds of intentionality do not have to be mutually exclusive (see section 2.2.1). Now I think that in the case of affective experiences and FOT/Fs there is a natural suggestion that, in turn, explains how they play (part of) the functional roles they play. They play these functional roles *because* they phenomenally represent propositional contents as true or false. In other words, they assume (part of) their functional roles in virtue of their phenomenal intentionality (e.g. Horgan and Tienson 2002).

Now somebody might object: “Wait, are you saying that truth feels positive and falsity feels negative? I don’t think they do.” Note that I am talking here of cases where truth and falsity feature as *phenomenal* qualities—in other words, when they feel *somehow*. Now, how do they feel if not positive or negative? Do they feel like—well—truth and falsity? The problem here is that we lack anything towards which we could point and say “look, *this* is how truth and falsity is supposed to feel”. This, I think, is because truth and falsity is *only felt* in FOT/Fs. There is thus no way of pointing towards some other state and say “look, *this* is how truth and falsity feels—and it is not positive or negative”.²²⁷ As I tried to make clear: *sometimes* affective phenomenal properties are quite subtle when it comes to their affectivity, especially if one tends to think affectivity narrowly and only present in feelings whose positivity or negativity comes “in a very large quantity (or a high intensity), explosively” (see section 6.5.1). I think we are fortunate that not all affective experiences are like that. Most epistemic feelings are not like that. Most FOT/Fs are not like that. I think we should acknowledge that the majority of our affective life is constituted by the little, subtle movements of our affective sensibilities. Bipolar disorder and related mental conditions demonstrate that it is all but desirable to have an affective life that is dominated by affective ebb and flow.

Note this is not to say that we have no other ways to arrive at truth or judge truth except via FOT/Fs. It is only to say that we have no other ways to experience truth or falsity *as a phenomenal property*. Just as we have established good tools to detect danger or contamination independently of phenomenology, we have found plenty of ways

²²⁷FOT/Fs are not unique in this regard. The same points apply to e.g. visual experiencing something as red or affectively experiencing something as fearsome. In the former case you might be able to point towards some kind of “visualness” or “colouredness” about the experience of red while in the latter case you might be able to point towards, well, the negativity of the experience, or rather towards the negativity or badness of what you are afraid of (section 4.4.3). I think that FOT/Fs are more like the latter case.

to detect truth and falsity independently of phenomenology (e.g. through testimony). It is undeniable that truth as a property has received a tremendous amount of attention—and for good reasons. We are a species that heavily relies on social coordination fueled by the exchange of large amounts of information. Our well-being thus delicately depends on the quality of this information: it better be true (see Sperber, Clément, et al. 2010 and section 8.2). All this attention might have concealed, however, that truth is a formal object *among many*. By that I do not mean that it is a non-phenomenal, non-affective formal object to belief (Williams 1970; Deonna and Teroni 2015). What I mean is that it is a phenomenal, affective formal object—a *concern* or *core relational theme*—to FOT/Fs. In other words, there is a psychological phenomenal seed to truth (and falsity) (see chapter 8).

To summarise, FOT/Fs have the following features: they are phenomenally conscious and intentional states that take propositions as their intentional objects. They represent the content of these propositions assertively, i.e. as true or false and they motivate or push—sometimes stronger, sometimes weaker—the subject to assent or dissent to what they represent as true or false. In doing so, however, they—as affective experiences in general—fall short of fully committing the subject to their contents. Furthermore, FOT/Fs are, depending on their base, sometimes content-gradable. One of the characteristic features of FOT/Fs is that they exhibit a phenomenal epistemic valence, i.e. they can directly represent their contents either as true (FOTs) or as false (FOFs). This phenomenal epistemic valence, as well as their assertiveness, is grounded in the general phenomenal valence of affective experiences. This valence also explains why FOTs and FOFs are phenomenal polar opposites and that FOTs feel genuinely different from FOFs even if their contents are the same. Finally, FOT/Fs—as all affective experiences—are non-voluntary, that is, they are not under voluntary control but happen to one. Call this (part of) the feature profile of FOT/Fs. Here is a condensed version: FOT/Fs are 1) intentional, 2) assertive, 3) motivational, 4) non-committal, 5) gradable in 5.1) (sometimes) content and 5.2) pushiness, 6) phenomenally epistemically valenced and 7) non-voluntary

Now, you might recall the idea of the present thesis: that intuition experiences are epistemic feelings. What was the feature profile of intuitions again? Here is the short version: positive/negative intuitions are 1) intentional, 2) assertive, 3) motivational, 4) non-committal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced and 7) non-voluntary. It appears that, apart perhaps from a small qualification on content-gradability, the feature profile of intuitions is identical to the

one of FOT/Fs. I think this is no coincidence (or wishful thinking on my part). In fact, I think we have established a firm basis to back the central claim of this thesis—I suggest that there is a straightforward explanation for this observation: the two feature profiles are identical because intuition experiences are identical to FOT/Fs. Put differently, another way of saying that some proposition *seems* or *appears* true or false is to say that it *feels* true or false (or right or wrong). This concludes Intellectual Affectivism. The following picture emerges: Suppose we have the positive or negative intuition that *P* and this leads us to judge accordingly that *P* is true or false. According to Intellectual Affectivism there is a propositional state (e.g. an imagining) that *P* co-occurring with a FOT/F directed onto the content of the propositional state. This feeling, in turn, pushes the subject to assent or dissent to *P*. The subject follows her feeling and judges *P* in accordance with it.

Crucially, in the course of this thesis, we have not *only* recovered the features of intuitions from Intellectual Affectivism but, at the same time, we have significantly advanced our understanding of them. This is no coincidence either: it is because we know quite a bit about affective experiences, epistemic feelings, FOR/Ws and FOT/Fs. We have seen that assertiveness and phenomenal epistemic valence are (partly) grounded in the core affective dimension of feelings: phenomenal valence. The pushiness of intuitions turns out to be explained by the general motivational nature of affective experiences. We have seen that intuitions are occasionally content-gradable because content-gradable mental states such as imaginings are involved as their bases. We have also seen that the affective dimensions of intuitions interact with content-gradable (imaginative) contents to make them more vivid, clear and distinct. We have made sense of the occasional inexplicitness of intuitions in reference to unconscious bases.

In contrast to what the perceptual analogy seems to suggest, Affectivism does not *overstate* the degree to which we are at the mercy of intuitions' non-voluntariness and apparent belief-independence. We can lose illusory intuitions with time, effort and some know-how. At the same time, Affectivism does not *understate* the persuasive power of intuitions. From a behavioural point of view, “cognitive” illusions are not as easy to dismiss or override as perceptual illusions (see section 4.5). In all this, one of the most important aspects about intuitions that Affectivism emphasizes is this: intuitions as FOT/Fs have many brothers, sisters and cousins among FOR/Ws, epistemic feelings and affective experiences. They are members of a happy, at times chaotic but ultimately mutually supportive family of mental states. In other words, intuitions are not on an solitary ontological post of their own but part of a rich, well-established psychological

family: affective experiences. Thus, Affectivism emerges as a good theory of intuition experiences:

Sufficiency of Affectivism (SA):

SA-P1 Affectivism is true if it can accommodate the feature profile of intuitions.

SA-P2 Affectivism can acknowledge the feature profile of intuitions.

SA-P3 Affectivism can explain the feature profile of intuitions.

SA-C1 Therefore, Affectivism can accommodate the feature profile of intuitions.

SA-C2 Thus, (by inference to the best explanation) Affectivism is true.

Another advantage of Affectivism in the form presented is that we are able to go beyond appearances and speculation when it comes to understanding intuitions. That is, we can start fathom their depths and try to get a grasp on their *actual* causal determinants. In fact, we already started doing so in section 7.3. In this vein, I will take a closer look on what we specifically know about FOT/Fs in order to further elucidate one of the perhaps most interesting aspects of intuitions: FOT/Fs and their relation to actual truth and falsity. This is the topic of the next and final main chapter.

The Mechanisms of Truth: Of Fluency, Coherence and Comprehension

8.1. The Mechanisms behind Truth and Falsity

We have performed an analysis of FOR/Ws and FOT/Fs on the phenomenological level, leading to the conclusion that intuition experiences are to be identified with these specific epistemic feelings. One of the major advantages of this result is that intuition experiences become ripe for a deeper mechanistic exploration. We can thus go some way in defusing Boghossian’s worry about intuitions:

To be sure, the idea that we possess a quasi-perceptual faculty—going by the name of ‘rational intuition’ [...] has been historically influential. It would be fair to say, however, that no one has succeeded in saying what this faculty really is nor how it manages to yield the relevant knowledge. ‘Intuition’ seems like a name for the mystery we are addressing, rather than a solution to it. (Boghossian 2000, p. 230)

If the theory [...] is to serve as a genuine explanation [...] rather than simply acting as a placeholder for such an explanation, it must consist in more than a suggestive label; it must somehow lay bare, in appropriate detail, how some capacity that we have gets to work on the properties we are able to think about so as to disclose their natures. (Boghossian and Bonjour 2001, p. 635)

Now, a mechanistic exploration has the potential to answer an important question: Do FOT/Fs result from properties that are good guides to actual truth and falsity? The answer to this question would be straightforward if these properties would be, well, truth and falsity. On first glance, however, the story behind intuitions appears to be more complicated. The present chapter is an attempt to tell this story. I will first consider the role of processing fluency and its relation to other things, such as actual truth (section 8.2). Then I will go beyond fluency and present a plausible memory coherence mechanism that is regularly employed to understand and validate incoming information (section 8.3). I will argue that this mechanism regularly produces FOT/Fs independently of whether it is fluent or not. Against this background, we will see that FOT/Fs are often produced on a basis that appears like an excellent psychological implementation of truth. informational coherence. This will also allow us to link up intuitions with our capacity for understanding and with beliefs.

8.2. The Status of Fluency

We have gone some way in “laying bare” the mechanisms behind intuitions by talking about fluency and the truth effect, as well as about epistemic feelings and affective experiences more generally. Against this background, it seems plausible that in specific contexts, fluency of various origins will lead to FOT/Fs. Now one important question is about whether fluency or disfluency are good guides to truth and falsity. As Stefan Herzog and Ralph Hertwig point out:

[F]luency’s validity depends simultaneously on the answers to two questions. The first is the question of *ecological correspondence*: Can fluency, in principle, accurately reflect environmental criteria [...] and thus potentially enable valid inferences about our world? [...] [T]here is little research available to corroborate this belief. This dearth of evidence, we suspect, stems from researchers mainly focusing on the second question, which we call the *susceptibility to manipulation* question: Can fluency-based judgments and decisions, in principle, be influenced by obviously irrelevant factors [...] that sabotage the potential correspondence between fluency and external criteria? As the extensive literature on fluency effects shows, the answer is unambiguously positive [...] Ecological correspondence and susceptibility to manipulation together imply that fluency will only lead to valid judgments and decisions

to the extent that there is both an ecological correspondence between fluency and environmental criteria and an absence of sabotaging influences, which would otherwise dilute fluency's validity. (Herzog and Hertwig 2013, p. 192)

Now, it is important to note that it would be premature to dismiss fluency-based FOT/Fs as reliable in tracking truth simply because they do not seem to track the distal property of truth directly but are based on proximal (psychological) properties such as fluency.

This rather seems to be expected: propositions that radically differ in their contents can be true or false. And they seem to have no single thing in common that makes them so.²²⁸ Furthermore, whether an individual proposition is true or false cannot simply be read off its content. An Externalist lesson is that the truth or falsity of a proposition is more often than not determined by a complex net of facts that are *external* to the proposition. In fact, these facts are not only external to the proposition but also lie outside the very situation where a subject tries to determine a proposition's truth. In other words: truth tends to be a property to which there is no direct access in principle. It is thus unsurprising that one has to rely on factors that go beyond the content of the individual proposition to determine its truth or falsity.

On the face of it, mental representations in general do not track any but the simplest properties directly. Our representations of many "simple" distal properties such as colour and depth are typically based on a myriad of inferences over accessible proximal properties or cues, something forcefully demonstrated by visual illusions (Jacobs 2002; Howe and Purves 2005). So the question concerning the reliability of fluency-based FOT/Fs should not be whether they track truth and falsity directly but rather whether

²²⁸Prinz makes this point concerning the high variability among particular objects exhibiting the formal object of danger, leading to feelings of fear:

[D]angers don't share any morphological properties in common. They don't look alike. There is no obvious set of appearances uniting all and only dangers. (Prinz 2006a, p. 149)

In explaining this he goes on to draw a somewhat skewed but nevertheless instructive analogy with the visual experience of red:

Assume that red is the power to cause a certain experience in us. Two things follow from this. First, the property of being red is morphologically heterogeneous. The physical entities that have the power to cause red experiences in us are highly varied. They have nothing intrinsic in common in virtue of which they might be grouped together. Their unity lies in the effect they have on us. Second, our perceptual representations of red do not resemble what they represent. Red is represented as a specific phenomenal quality. There is nothing out there that is intrinsically red. (ibid., p. 149)

I think similar things can be said in relation to FOT/Fs and how they represent truth (see section 7.4).

processing fluency is a good proxy for truth and falsity.²²⁹

This will in part depend on what, in turn, causes fluency (e.g. repetition). So the *first two questions* are: 1) Is fluency a good proxy for truth? 2) What causes fluency and is this a good indicator of truth?

Let me briefly summarise some reflections that go some way in answering these questions. As we have seen, prior exposure to a stimulus results in facilitated processing of that stimulus, i.e. fluency. Fluency is thus a consequence of prior exposure and therefore a relatively reliable proxy for it. Furthermore, exposure frequency, in comparison to other fluency-modulating and obviously truth-orthogonal factors such as rhyming, is actually a cause of fluency that might ecologically (although not logically²³⁰) correspond with truth. This leads to the question: is exposure frequency (which is a diachronic cause of fluency) a good proxy for truth? What would our informational environment have to be like in order for exposure frequency to be correlated with truth?²³¹ Assuming that truth is determinable, it would have to be the case that one gets more frequently exposed to (the same) truths than to (the same) falsehoods.

Three considerations speak in favour of this being the case: First, we are a species that heavily relies on social coordination fueled by the exchange of large amounts of information. Our flourishing is thus intimately intertwined with the quality of information gathered and traded by individuals. As a consequence, telling the truth is typically socially rewarded and circulating falsehoods is usually socially penalized. This establishes a social pressure towards a higher proportion of truths in the informational pool.

Secondly and relatedly, if something like a Sperberian relevance expectation (that a communicated piece of information is believed to be relevant by the communicator) and a Gricean maxim of quality (that a piece of communicated information is believed to be true by the communicator) indeed guide a good portion of communicative behaviour,

²²⁹By “direct tracking” I mean to capture here the idea that FOT/Fs are somehow directly caused by actual truth and falsity, without intermediary variables, similar to the idea of direct causation introduced in section 5.3.1.2 and footnote 144.

²³⁰That repetition is not *logically* connected to truth was famously illustrated by Wittgenstein. According to him, taking a statement to be true because it is repeated is as “if someone were to buy several copies of the morning paper to assure himself that what it said was true” (Wittgenstein 2009, section 265).

²³¹One assumption that needs to be made is that the truth of a statement is actually *determinable*. This seems more straightforward for some subject matters—such as physical, mathematical or logical reality—than for others—such as social or metaphysical “reality”, including some philosophical, spiritual and religious speculations (cf. Reber and Unkelbach 2010, pp. 577 sqq.). Exposure frequency would, in any case, remain a good indicator for *social consensus* (within a sourced informational ecosystem) if not for truth (cf. Koriat 2012).

then we should expect that one will encounter more truths than falsehoods (Grice 1975; Sperber and Wilson 1987; Reber and Unkelbach 2010, p. 576).

Third, false propositions (ideas, theories etc.) are not constrained by reality and thus infinite while true propositions (ideas, theories etc.) are finite in virtue of being so constrained. Even if we make plausibility constraints on the false propositions, there is still a large number of plausible false propositions (e.g. “The present height of the Cheops-Pyramid is 123/124/130 meters”) while there is only *one* true proposition (e.g. “The present height of the Cheops-Pyramid is 138.75 meters”) from which the plausible falsehoods are often derivatives. As a consequence, it is much more likely to encounter repeatedly one and the same true proposition than to encounter repeatedly one and the same false proposition. Thus, one should encounter true propositions more frequently than false ones (Unkelbach 2007, p. 229; Reber and Unkelbach 2010).²³²

So to summarise the tentative emerging answer to question 1) and 2): in virtue of being a product of prior exposure, fluency emerges as a relatively good proxy for truth. As Unkelbach concludes:

Consequently, there is a good chance that people encounter the same true statements, concepts, or ideas more frequently or repeatedly than they do false statements, concepts, or ideas, and therefore, these true items are more fluently processed, creating the proposed positive correlation between processing fluency and truth. To the same extent, as false statements are repeated (such as with rumors, urban legends, or false propaganda), they may gain credibility as well; yet, these cases should be the exception rather than the rule, and they exploit the positive correlation between fluency and truth rather than create them. (Unkelbach 2007, p. 229)

Apart from these general reflections, the picture grows yet more complex by acknowledging that there are different kinds of fluency (e.g. perceptual fluency vs. conceptual fluency, absolute vs. relative fluency) and among them some might be better guides to truth than others. After all, fluency is a process property, i.e. a property a process might possess or lack. The fluency of, say, a specific conceptual or memory process (conceptual or mnemonic fluency) might be more closely tied to the truth of a proposition than the fluency of a specific perceptual process (perceptual fluency). So, a *third question*

²³²For more considerations and details about the ecological validity of fluency see Reber and Unkelbach 2010; Herzog and Hertwig 2013. For considerations on the informational value of *discrepant* fluency see Hansen and Wänke 2013.

arises: 3) What is it that is fluent and what does it do?²³³ For instance, comparing the relatively stronger repetition-based truth effect to the truth effect based on perceptual fluency, “it seems that repetition has a stronger connection to truth, which is also less malleable than in the case of perceptual fluency” (Silva, Garcia-Marques, and Mello 2016, pp. 830 sq.).

In effect, invoking processing fluency in the explanation of a phenomenon is fairly uninformative. The common way to measure processing fluency is by simply measuring response times: “response latency [...] can be used as a proxy for processing fluency; however, it is important to note that this index is imperfect, because many other factors not related to the experience of fluent processing can contribute to response latencies” (Unkelbach 2007, p. 221; see also Forster et al. 2013). So, based on the ambiguous observation that a subject responds faster (taken as a measure of processing fluency), an explanation in terms of processing fluency states that there is some process that is (relatively) fast and that is why, say, we judge a statement as true. Having said that, just as the observation that a subject responds faster tells us very little about *why* this is the case, fluency explanations tell us very little about *what* the fluent process in question is and does: “Processing fluency per se does not explain *why* people respond to fluently processed information with higher-rated truth” (Unkelbach and Rom 2017, p. 111). In other words, fluency explanations are quite thin explanations.

This foreshadows the, as far as I can tell, neglected *final question*: 4) What, other than fluency, causes affect and leads to FOT/Fs? In the remainder, I want to focus on this last question. We will see that it naturally bears on the other questions. Now, by asking this question I do not primarily ask about experimental ways one can devise to sneakily induce affect that is not based on fluency and to trick subjects into taking this affect to signal truth. What I mean is: are there processes and (mental) conditions other than fluency that naturally and routinely lead to affect in the form of FOT/Fs? And are *these* reliable guides to truth and falsity? Or even somehow substantially bound up with truth and falsity?

After all, work on epistemic feelings such as FOKs and TOTs have uncovered that these feelings are, in fact, *multiply determined* (see section 5.3): “FOKs are not governed by one underlying mechanism. Rather, a host of [...] mechanisms seem to work in concert to produce the single subjective state” (Thomas, Lee, et al. 2016, pp. 88 sq.). This

²³³This question can be seen as a more confined version of question 2, targeting not the diachronic cause (e.g. repetition) but the synchronic cause of fluency, i.e. the character of the very process that is fluent and on whose workings fluency supervenes.

obviously makes sense: if several proximal cues for the presence of a distal property, say information possession, are available, it is useful to use them all (Dretske 1981). By implication, if only one cue, say, perceptual fluency, is available, then *this* will be used, possibly resulting in a weaker (i.e. less intense) FOK than one based on multiple cues (e.g. perceptual fluency, conceptual fluency and availability of partial information) (cf. Schwartz 1994; Schwartz, Travis, et al. 2000). In a similar vein, looking at perceptual fluency Reber and colleagues emphasise that:

the relationship between objective and subjective perceptual fluency is more complicated: [...] subjective fluency is a feeling based on objective perceptual fluency at different stages of perceptual processing. Subjective perceptual fluency is a unified experience based on objective fluency from several sources. (Reber, Wurtz, et al. 2004, p. 48)

Something similar will plausibly also be the case for FOR/Ws and FOT/Fs. That is, they will be multiply determined by fluency and other properties that are hopefully useful as proxies for rightness/wrongness and truth/falsity. This is suggested by the observation that the repetition-based truth effect is typically stronger and more robust than the truth effect based on perceptual fluency. With this in view, Silva and colleagues instructively suggest:

[O]ne other characteristic of repetition may contribute to its strong impact on truth judgments. Repetition aggregates different levels of fluency. Besides the perceptual fluency that comes from reprocessing the wording and phrase structure of the statements, repetition also increases conceptual fluency due to the reprocessing of the semantic content and meaning of the stimuli. (Silva, Garcia-Marques, and Mello 2016, p. 831)

Relatedly, Unkelbach and Stahl 2009 have shown that there is fluency due to repetition and fluency due to truth, and that people seem not to be able to distinguish between the two. It seems that different process-specific fluencies combine to multiply determine the unitary affective state of FOT. In fact, it is not obvious that varieties of fluency are the only or even primary determinants here. This is why Unkelbach and Greifeneder forcefully argue that:

processing fluency influences judgments and evaluations jointly with other available information. [...] [A]ny given judgment is influenced by informational cues which are weighted according to their subjective validity and

then linearly integrated. Thus, both processing fluency and other information (e.g., knowledge) should jointly influence judgments and evaluations. [...] [A]ny cue may “rationally” influence judgments, as long as it has an ecological correlation with the to-be-judged criterion; and fluency might be such an ecologically valid cue [...] Thus, even if someone believes to *know* that Osorno is in Chile, presenting this piece of trivia in a difficult-to-process way may lead to lower truth ratings just because the statement *feels* false. (Unkelbach and Greifeneder 2018, pp. 78 sq.)

Now, I suggest combining the realization that individual epistemic feelings such as FOKs are multiply determined with the findings of Topolinski and colleagues, which indicate that semantic coherence, perceptual fluency and other affect-triggers integrate into *one* affective experience, a feeling of semantic coherence (Topolinski and Strack 2009b,a). I submit that the same might be the case for FOTs: it is not necessarily the judgment that integrates multiple sources of truth-bearing information (“experience of fluency”, apparent knowledge) but the FOT itself that is multiply determined by these sources (see also Winkielman, Ziembowicz, et al. 2015 and section 7.3.4 and footnote 213).

We have already seen that fluency, logical validity, believability *and* conflict have an influence on FOR- and liking ratings and thus probably on the quality of the underlying feeling, making it multiply determined.²³⁴ Now, it is conceivable that other properties such as truth (or other truth-related properties) have ways of leading to affect without having to go the detour over fluency (although they likely typically go both ways). This is perhaps the reason why people fail to distinguish between different kinds of fluency such as fluency due to repetition and fluency due to truth: both issue into a unitary affective state on the personal level that can be jointly fed from different sources such as the instantiation of different kinds of fluency and other properties such as truth.

This possibility becomes more plausible if one considers the following: if affect functions as an evaluation of certain states of affairs, it would be surprising that the only affectively valuable property in relation to truth would be fluency. On the face of it, one could well

²³⁴Of course, fluency and perhaps conflict, on the one hand, and logical validity and believability, on the other hand, are not on a par when it comes to their causal proximity and the underlying processes. The idea I am leading up to here is that the proximal causal properties behind the more distal properties such as logical validity and believability are not obviously processing fluencies. Among other things, this is hinted by the fact that e.g. Thompson and colleagues find variables such as logical validity, believability and conflict to explain part of the variance in FOR ratings *independently* of response time, i.e. answer fluency (e.g. Thomas, Bulevich, et al. 2011, see also end of section 7.3.2). This appears at odds with the idea that processing fluency is the sole proximal causal property in these cases. More on that below.

imagine other truth-relevant properties whose instantiation would be affectively valued. One such property is plausibly the coherence of (currently activated) information *itself* (not potentially resultant processing fluency). That this is the case is actually suggested by one of the most prominent theories in psychology: the theory of cognitive dissonance (Festinger 1962; Zanna and Cooper 1974; Elliot and Devine 1994; for a review see Harmon-Jones, Amodio, et al. 2009).

The original theory of cognitive dissonance predicted that when an individual holds two or more elements of knowledge that are relevant to each other but inconsistent with one another, a state of discomfort is created. This unpleasant state is referred to as “dissonance.” [...] [T]he unpleasant state of dissonance motivates individuals to engage in psychological work in an effort to reduce the inconsistency between cognitions. (Harmon-Jones, Amodio, et al. 2009, p. 121)

Conversely, I find it eminently plausible that the instantiation or detection of coherence would have (subtle) affective consequences, independently and additionally to the affective consequences of processing fluency.

In fact, it seems clear that coherence itself would potentially facilitate processing fluency as well (cf. Gawronski and Strack 2012; Winkielman, Huber, et al. 2012). As Winkielman and colleagues helpfully note:

It is useful to highlight a few things about the relation between features and fluency. First, both can be available simultaneously, with each contributing to the final reaction. For example, positivity from detecting a “smile” feature can combine with positivity from the ease of face recognition. Second, features and fluency can play off each other. For example, the same feature, such as symmetry, might create a positive reaction because of its cognitive implications (e.g., health), but also make the face easier to recognize. In other words, a feature might not only create an evaluative reaction directly, but also indirectly, via its influence on fluency. (Winkielman, Huber, et al. 2012, p. 90)

Insofar, fluency emerges as a potential, difficult to disentangle *confound* for studying the effects of many properties that themselves trigger affect *and modulate fluency* (cf.

Landwehr et al. 2017). These properties count among themselves coherence as well as the above-mentioned logical validity, believability and conflict.²³⁵

In fact, Topolinski and colleagues posit processing fluency to account for the link between semantic coherence and affect. They instructively mention two things, however:

The present model assumes fluency to be, procedurally, the first link in the intuitive chain and affect to be its consequence. [...] [O]ne could object that perhaps coherence first triggers positive affect, which in turn increases the fluency of processing coherent triads. (Topolinski and Strack 2009b, p. 58)

It is still not entirely clear why the semantic coherence of a triad leads to more fluent processing. That the automatic semantic activation of the common associate [...] facilitates processing is not a satisfying explanation because it does not identify the underlying process. (Topolinski and Strack 2009c, p. 1497)

Both these tensions can be resolved when we assume that semantic coherence can trigger positive affect by itself *as well as* fluency (which additionally triggers positive affect by itself). There is no need to exclude a direct coherence-affect link because there is a fluency-affect link and these links are complementary rather than mutually exclusive. In the same vein, the need to find a coherence-fluency link to account for affect becomes less pressing because there is a direct coherence-affect link.

Note further that we do not even have to go much further than semantic coherence to find one way in which *propositions* trigger fluency and affect in a fashion analogous to (coherent) word triads: they just need to include elements that have common associates. In fact, the disambiguation mechanisms underlying the detection of semantic coherence are plausibly part of machinery that is routinely and automatically employed to make sense of sentences in everyday life. This is why the rather artificial word triad task works.

²³⁵Alternatively, one might want to argue that in these domains it is in fact *only* fluency and disfluency that leads to affect and properties such as coherence lead to affect *only* via their link to fluency. In fact, *whatever* property is suggested to lead directly to affect, one might always be able to argue that this property also leads to fluency and that is how affect is triggered. I acknowledge that it is quite difficult to discriminate between this possibility and the proposed confound idea. Note that using the findings showing that one could trigger affect without fluency and that this affect is indistinguishable from fluency-triggered affect like e.g. in Topolinski's semantic coherence tasks cannot exclude the all fluency proposal. This is because the proponent of this view can discount these instances as anomalous and what we are after is something that is natural. At the same time, it seems to me that the all fluency proposal lacks independent motivation.

The overarching guiding question of this chapter is: Do FOT/Fs result from properties that are good guides to actual truth and falsity? Against this background, I considered four questions in this section: 1) Is fluency a good proxy for truth? 2) What causes fluency and is this a good indicator of truth? 3) What is it that is fluent and what does it do? 4) What, other than fluency, causes affect and leads to FOT/Fs? To summarise the tentative emerging answer to question 1) and 2): since fluency is a typical product of prior exposure, fluency emerges as a relatively good proxy for truth. Delving on question 3) brought to the fore that fluency explanations are in general quite thin. They do not tell us much about what it is that is fluent and what it does. But much will depend on this. The search for mechanisms to flesh out such explanations, potentially even independently of fluency, brought us to question 4). I pointed out that FOT/Fs like other epistemic feelings will plausibly be multiply determined—not only by fluency. Considering what we know so far, revealed another plausible candidate property that itself often leads to fluency: coherence. What is coherence? That is, how are we to understand it in terms of psychological processes? We will find out in the next section.

8.3. Getting to the Truth

8.3.1. Coherent References and Truth

When we realise that there might be direct links between affect and properties such as coherence, things start to get quite interesting. In fact, Unkelbach and Rom 2017 have recently proposed a new mechanistic explanation of the repetition-induced truth effect that goes beyond thin processing fluency explanations. Here is the basic idea: When we encounter a statement, say, “The world’s most poisonous snake is the Australian Inland Taipan” (or any stimulus, really), “corresponding references” (e.g. “snake”, “Australia”) in the form of memory traces are activated.²³⁶ These memory traces give meaning to the elements of the statement. Furthermore, if the statement one encounters is not yet in one’s memory repertoire, it will create a localised network of memory traces that represents the statement and connects its elements to each other. These elements can either come in the form of activated pre-existing corresponding references (e.g. “snake”) or ones that are newly created by the encounter (e.g. “Taipan”).

²³⁶For stimuli in the form of objects one could potentially develop similar points by reference to mental files (Recanati 2012, 2016).

Importantly, the elements of the localised network also connect to relevant background knowledge of the subject (if any) that, of course, organise in-memory trace networks as well. Among other things, these networks embed the pre-existing corresponding references (e.g. “Australia”, “snake”). In our case, activating the memory trace for “Australia” and “snake” would activate and connect to memory elements that in concert represent what one knows about Australia and snakes. This implies that the meaning-bearing corresponding references can be, based on one’s background knowledge, coherent or incoherent with each other. Taking into account one’s background knowledge about Australia and snakes (e.g. that Australia is home to many deadly animals), the corresponding references for our example statement provide coherent meaning for the statement. In contrast, this would not be the case for the statement “The world’s most poisonous snake is the Swedish Inland Taipan”. Despite the similarities in the kinds of corresponding references (e.g. country name, name of snake, qualifier about how poisonous), Sweden, unlike Australia, is not associated with poisonous animals. In other words, Sweden and highly poisonous snakes are incoherent elements; the integration of the memory traces activated by the statement results in incoherence. As a result, the statement involving Australia will be more likely judged true due to the triggered coherence while the statement involving Sweden will be more likely judged false due to the triggered incoherence.

Against this background, the determining relations between memory traces or corresponding references can be (in a graded fashion) either positive, negative or neutral relative to their coherence or relational consistency (see also Thagard 1989; Kunda and Thagard 1996; Betsch and Glöckner 2010; Öllinger and von Müller 2017). Note that the connections between elements are learned and can express relations of “covariation, rates of reinforcement and punishment, logical relations such as negation or causation, and so forth” (Betsch and Glöckner 2010, pp. 288 sq.). Furthermore, it appears that the relation holding between elements can be either 1) taken at face value from the statement (especially if one lacks sufficient background knowledge), 2) informed by one’s background knowledge or 3) a combination of both.

Another consequence of this picture is that localised networks representing a statement’s meaning can be more or less embedded or integrated, depending on how much relevant background knowledge the subject possesses and brings to bear in the moment of the statement encounter. Statements about whose subject matter one is knowledgeable will integrate with global networks, e.g. for our example if one is a snake enthusiast. Conversely, statements about whose subject matter one knows little will integrate with

more localised networks, e.g. if one is a snake layperson. This is the picture assuming that there are some existing memory traces that a statement activates, i.e. that there *are* corresponding references for a statement. In our case, that one knows some things about snakes and Australia. If, however, such references are missing, encountering a statement (say, “Mimas has more spin than Pallene”) will lead to the formation of new memory traces that organise into a new localised and wholly isolated network, an “island in one’s web of knowledge”, so to say.

The explanation of Unkelbach and Rom 2017 for the repetition-based truth effect goes now as follows: The first presentation of an unknown statement (e.g. one does not know that Taipans are the most poisonous snakes and live in Australia) activates corresponding references (e.g. “snake”, “Australia”) and their respective links, potentially, strengthening pre-existing references and links; a localised network that represents this specific statement is formed, i.e. corresponding references are created. As a result, when a statement is repeated, it will have more corresponding references that are coherently linked (as suggested by the formerly presented statement) than new statements:

The theory then explains the truth effect as follows: In the presentation phase, a new statement is processed and stored in memory; that is, if the statement’s elements have corresponding references (e.g., Australia, Taipan), they will be linked and form a localized information network. In the test phase, repeated statements then have on average more corresponding references which are coherently linked due to the prior presentation. New statements have on average fewer corresponding references that are coherently linked; thus, they appear relatively less true (ibid., p. 113).

Furthermore, Unkelbach and Rom note that this explanation is well compatible with a fluency explanation. In fact, it can be thought to put some flesh on the bones of the thin fluency explanation:

This process is fully compatible with a fluency account. The only necessary assumption is that many corresponding references that are coherently linked lead to fluent processing of a given statement. Importantly, integrating processing fluency as an output variable and not as the ultimate causal variable solves the theoretical problem of how people learn to associate fluent processing with truth. If one assumes that corresponding references that are coherently linked lead both to judged truth and to fluent processing, people would be constantly exposed to the correlation of truth and fluency, leading

to “illusory truth effects” when fluency and truth are factually orthogonal, or when fluency is manipulated independent of repetition (e.g., by color contrast or rhyming). (Unkelbach and Rom 2017, p. 113)

In effect, the authors provide an explanation for why fluency comes to be associated with truth. It is in virtue of two other interconnected factors that are closely associated with both fluency and truth: the presence of pre-existing corresponding references and their coherence. On the one hand, coherent references lead to fluency. On the other hand, they are closely related to truth (while, of course, not entailing it). In other words, “the fluency experience is constantly paired with experienced truth from coherent networks of corresponding references” (ibid., p. 122). Because of these associations, the two associates of coherent references come to be associated as well: fluency becomes associated with truth. If now fluency occurs but is not due to coherent references, fluency will still be taken for truth. Consequently, other fluency-related factors lead to a truth effect as well. The present account also seems to offer a straightforward explanation for how “repetition has a stronger connection to truth, which is also less malleable than in the case of perceptual fluency” (Silva, Garcia-Marques, and Mello 2016, pp. 830 sq.). Clearly, coherent references are more intricately linked with truth than perceptual fluency.

Unkelbach and Rom find a way to demonstrate that coherence is prior to fluency as a determinant of truth judgments: they show that statements that contradict but are otherwise identical to previously presented statements (e.g. “Crocodiles sleep with their eyes closed” vs. “Crocodiles sleep with their eyes open”) are judged promptly, i.e. fluently. However, they are not promptly judged to be true but to be *false* (see also Silva, Garcia-Marques, and Reber 2017, experiment 2).

As the old-contradicting statements have corresponding references with an incoherent element, they should lead to more, and relatively fast, “false” responses in comparison to new statements. [...] The prediction of relatively fast “false” responses in comparison to new statements sets the present theory apart from a fluency explanation, where relatively faster responses should map on relatively higher truth ratings [...] And because processing fluency is assumed to be a unitary experience [...], the truth effect should emerge independent of the specific source for processing fluency, that is, independent of whether fluency originates due to color-contrast, rhyming, previous reading, conceptual priming, or repetition (Unkelbach and Rom 2017, p. 113).

Thus, old-contradictions lead to incoherence and incoherence leads to “false” responses. At the same time, incoherence does not lead to disfluent “false” responses but to fluent “false” responses. In other words, incoherence does not lead to disfluency but to fluency, yet this does not lead to “true” responses but to “false” responses. That is, we observe an association between coherence and truth and a dissociation between fluency and truth.

This also brings to the fore that fluency, measured usually via response latencies or times, might not always be a sign of process success. Instead, it can indicate the prompt *failure* of a process or other negative outcomes. In other words, the same is true for the fluency-success link as for the fluency-truth link: it does not always hold. Perhaps, however, it is still possible to argue that while the outcome of the fluently *failing* process is ripe for a negative evaluation (it is a failure, after all), the fact that it fails *fluently* might still be of positive value. This illustrates the fact that while the link between fluency and some other properties might be fragile at times, the link between fluency and speed is an identity relation and—from a resource allocation point of view—speed is nice (all things being equal).

There is something else that is shown by these findings: corresponding references alone do not do the truth trick. The references must be coherent in order to be judged true and incoherent in order to be judged false.²³⁷ This account nicely illustrates that fluency does not have to be the determinant of the truth effect. There are properties that lead to the truth effect that seem to be more robustly linked with truth than fluency.

Note that Unkelbach and Rom essentially go the way that I suggested for semantic coherence at the end of section 8.2. Topolinski and colleagues were under the impression that it is necessary to find a connection between semantic coherence and the positive affect which they found to ground semantic coherence judgments. They assumed the connection to be fluency. I noted the possibility that additionally to a fluency-mediated link, there might be a direct link between semantic coherence and the affect on whose basis semantic coherence judgments are made. Unkelbach and Rom, now, identify a direct link between coherence and truth judgments, additionally to a fluency-truth link. They, however, do not reflect on the role affect might play in the picture and whether there might be sometimes an experiential state that mediates between coherence and truth judgments.

²³⁷The thinnest basis for the coherence/incoherence of corresponding references is to take the statement that created them at face value in the absence of background knowledge (cf. Gilbert 1991; Unkelbach and Rom 2017, p. 112). For instance, the statement “Mimas has more spin than Pallene” would create a localised and isolated network where e.g. the connection between spin and Mimas is positive and thus having a spin is coherent with being Mimas.

I suggest combining these two lines of reasoning: Coherence, additionally to causing fluency (and fluency causing positive affect), directly causes positive affect which is used as a basis for a truth judgment (see also Hookway 1993, pp. 162 sqq.). This way we can also tell another, potentially complementary story for why fluency leads to a truth effect as well: because it causes affect, whether it is due to coherence or not. Not only is there a (learned) connection between coherence, fluency and truth. There is also a connection between coherence and affect as well as fluency and affect. And having a link with affect that simulates the coherence-affect link is all it takes to potentially bring about a truth effect.

As a consequence, the repetition-induced truth effect is stronger than the one induced by perceptual fluency not only because it aggregates fluencies but also because it aggregates sources of affect. Perceptual fluency is only one source of affect while in the repetition-based truth effect there is perceptual fluency *and* coherence (and, perhaps, some additional conceptual fluency that might be in part caused by coherence) that cause positive affect. Affect, in turn, provides the common experiential currency with which various disparate, process-specific but valuable properties are priced and cashed out (see also section 6.4). Various fluencies and coherence are de-modalised to combine into a unitary affective state that puts a comparative value on a statement to be judged true. This is also likely why discrepant fluency counts for more than “just fluency”: a statement that is processed faster than other statements is valued higher relative to those other statements, and is consequently more likely to be judged true.

Of course, various properties do not always have to pull into the same affective direction:²³⁸ in the case of quickly judged statements that contradict but are almost identical to previously presented statements, it might be that negative affect from incoherence offsets positive affect from fluency, issuing into a net negative affective experience: a feeling of falsity (cf. Kelley and Jacoby 1998, p. 130).²³⁹ An affective experience is typically the outcome of a computation that integrates several sources of positive and negative value, valence or affect. This is also how fluencies of seemingly irrelevant origin can exert an effect on various judgments. Note, however, that the mere speed of a judgment (the observed fluency variable here) does not seem like an ideal indicator of whether

²³⁸Thompson’s findings that answer fluency is *positively* and conflict is *negatively* correlated with FOR-ratings might be taken to suggest this as well (see end of section 7.3.2).

²³⁹There are also findings that suggest that fluency taken for itself (in contrast to coherence or incoherence) might not be fixed in its affective value across contexts, so that there might not be positive affect that needs offsetting (Unkelbach 2006, 2007; Olds and Westerman 2012; Silva, Garcia-Marques, and Mello 2016; but see Coutanche and Thompson-Schill 2012).

the affective state guiding this judgment is positive or negative. You might judge very quickly on the basis of a negative affective state such as fear. Thus, the assumption here is only that a prompt judgment is indicative of processing speed (i.e. fluency) which *ceteris paribus* is positively valued but might not drive the overall computation of value (see also above).

Nevertheless, fluency (and other affect-triggers) might have a positive effect on truth judgments despite the presence of incoherent corresponding references (Fazio, Brashier, et al. 2015). Fazio and colleagues found that the repetition-induced truth effect was modulated but not eliminated by knowledge (i.e. corresponding references). This is in line with FOTs being multiply determined: it is not *only* coherence or incoherence (as it is not only knowledge) that counts.²⁴⁰ In some cases, fluency/disfluency might perhaps even offset incoherence/coherence. This might be fostered in cases when corresponding references might fail to be cued or activated in a specific context (e.g. because the knowledge is not well integrated, highly context-specific or of low relevance to the subject), when the quality of corresponding references is compromised (e.g. because the memory traces are dated and/or their encoding was not sufficiently specific)²⁴¹ and/or when memory is fragmented (see also Isberner and Richter 2014, pp. 249 sq.; Elga and Rayo 2015).

In this section we have first considered an alternative, mechanistic explanation of the repetition-based truth effect. This explanation relies on the idea that a previously encountered statement activates previously established corresponding references in memory. If these references are coherent, then this leads the subject to judge the statement as more likely true. On the other hand, if the references are incoherent, then this leads the subject to judge the statement as more likely false. I suggested a plausible way to understand how the subject becomes aware of her coherent/incoherent corresponding references: coherence or incoherence directly (and indirectly via fluency) triggers positive/negative affect. In other words, the coherence computations over corresponding references lead to FOT/Fs. In the next section I will take this explanation a step further and integrate these insights with how we decide whether something is true based on our actual knowledge about it, and not on an “educated guess”. For that I will draw on resources from the field of language comprehension.

²⁴⁰Except if one assumes that there are some kind of “monopolizing determinants” that block the (affective) integration of other variables (that are usually integrated).

²⁴¹This might be behind the findings of Garcia-Marques, Silva, Reber, et al. 2015 and Silva, Garcia-Marques, and Reber 2017 that an initial falsity effect (when corresponding references are freshly formed) turns into a truth effect after a week.

8.3.2. Epistemic Monitoring

There is an interesting aspect of Unkelbach and Rom’s proposal that the authors do not mention: While their aim is to identify mechanisms for the repetition-induced truth effect, it is plausible that the very same mechanisms are at play when truth is evaluated based on actual knowledge of a statement. Unkelbach and Rom focus on cases where one does not “know” the specific statement. These unknown statements target corresponding references whose coherence/incoherence allow for a more or less educated guess about the truth or falsity of the unknown statement. But what happens if one actually knows or seems to know that, say, “Freddie Mercury was born with the name Farrokh Bulsara”? Plausibly something very similar: corresponding references are activated and are integrated relatively to their coherence/incoherence.

A known statement will target corresponding references that are already interconnected in the way (not) suggested by the statement. In other words, we get a match (or mismatch) between a statement and information currently activated in memory. In fact, this is essentially what happens when a statement is repeated. If a statement was previously presented, it has established matching corresponding references for the repetition case. There is really something uncanny about the way repetition can be epistemically hijacked:

It was Napoleon, I believe, who said that there is only one figure in rhetoric of serious importance, namely, repetition. The thing affirmed comes by repetition to fix itself in the mind in such a way that it is accepted in the end as a demonstrated truth. (Le Bon 2001, p. 77)

However, repetition does not have such a profound impact on our belief system for no reason. It has this influence because it usually *is* a good proxy for truth—not least because coherent references are a good proxy for truth—and relying on it is, under certain assumptions about the informational environment, an adaptive feature of our cognitive system (Unkelbach 2007; Reber and Unkelbach 2010; Herzog and Hertwig 2013; Marsh et al. 2016, see section 8.2). Furthermore, memory systems themselves exhibit features that to a degree (can) hold the influence of repetition in check. There *is* a quantitative if not a qualitative difference between repetition matches and knowledge matches: in comparison to statements that found their way into memory through simple prior exposure, statements that are “known” will have corresponding references that are more firmly consolidated. Yet this could also be achieved with sufficient repetitions. However, statements that are “known” will (hopefully) be more interconnected and

embedded with the rest of the subject’s knowledge, therefore accumulating more and broader coherence. This will also have the result that if a subject knows a presented statement to be false, the localised network created by the encounter will connect to embedding networks that represent discordant knowledge. These embedding networks will modulate the coherence that the localised network can muster. One upshot of this is that fragmented belief systems are (epistemically) hazardous and that “defragmentation” (via e.g. bridging pieces of information in memory or presenting the “big picture” from the start) might be an effective way to curtail the influence of isolated localised memory networks (cf. Elga and Rayo 2015; Levy 2015; Mandelbaum 2016). Be it as it may, the matching between a presented statement and information currently activated in memory has a name in the language comprehension literature: “epistemic validation” or “epistemic monitoring” (e.g. Richter et al. 2009; Singer 2013; Isberner and Richter 2014; O’Brien and Cook 2016). Epistemic monitoring describes an automatic, effortless, “routine, nonstrategic validation process that detects knowledge violations” (Isberner and Richter 2014, p. 246). Epistemic monitoring is a step in the language comprehension process that itself consists of activation, integration *and validation* of information.

It denotes the idea that comprehenders use their prior knowledge and beliefs to monitor linguistic information for consistency with previous text information and with their own knowledge and beliefs (Richter 2015, p. 338).

Insofar, validation is a coherence computation under conditions of matching knowledge. It appears that the comprehension process can fail at different stages, potentially resulting in different phenomena: A failure at the activation and integration stage will likely result in (a feeling of) incomprehension and/or confusion while success will result in (a feeling of) understanding. On the other hand, something that comes closest to be properly described as a failure at the validation stage is presumably found for statements for which the subject lacks sufficient information, resulting in ambiguity that can lead to uncertainty, curiosity or doubt (see also Hookway 1998; Thagard 2004; Carruthers 2017b, 2018). Such kinds of statements are ordinarily employed in truth effect tasks. Note that detecting knowledge violations is not a failure of validation—rather it is its function. In sum, the coherence computation behind validation can result in ambiguity, have a positive result or detect knowledge violations of various kinds:

They may comprise information that is clearly false (based on semantic or world knowledge), information that is merely implausible, or information that is not false or implausible per se but inconsistent with antecedent text.

(Isberner and Richter 2014, p. 247)

In this framework, the knowledge basis against which incoming information is validated will depend on what information is currently activated to construct a meaningful representation based on a language element.

[T]his monitoring process relies on the activation of knowledge through memory-based processing [...] However, it is important to note that not all knowledge that is potentially relevant for evaluation will be activated during situation-model construction; rather, the activation will be a function of accessibility. It is reasonable to assume that the accessibility of knowledge will, among other things, be influenced by how recently it was previously activated, by its typicality given the reader's experience of the world [...], by how well connected it is with other stored knowledge [...], by how the currently processed text information is phrased (resulting in more or less surface overlap with knowledge in long-term memory [...], by how focused the information is in the text [...], and by the depth of processing required by the task [...] Thus, violations of knowledge that is not activated for situation-model construction may well go unnoticed. (ibid., pp. 249 sq.).

The computed coherence will thus be coherence with the subject's "knowledge" (i.e. information in memory) that is cued by a given language item (word, sentence, text, speech) and context, not coherence with the hypothetical totality of a subject's knowledge and much less with the "totality of knowledge simpliciter" (see also Elga and Rayo 2015).

Accordingly, we do not propose that comprehension entails a full analysis of the (potential) truth of information but rather a quick and incomplete analysis based on the knowledge that is activated for situation-model construction. If an inconsistency with the activated knowledge is detected, the incoming information is initially rejected to protect the situation model from contamination with false information (Isberner and Richter 2014, p. 250).

As a consequence, epistemic monitoring emerges as a process "which protects the system from false information and thereby—in general—promotes accurate and stable mental representations (ibid., p. 257). Interestingly, this validation mechanism does not seem to differentiate between various kinds of information and seem to be identical for statements that express moral convictions or opinions (Van Berkum et al. 2009; Gilead et al. 2019,

see also footnote 33 and 219). By and large, epistemic monitoring can be considered as a member in the “suite of cognitive mechanisms for epistemic vigilance, targeted at the risk of being misinformed by others” (Sperber, Clément, et al. 2010, p. 359).

In this section, I have motivated the idea that the coherence mechanism outlined in the previous section is actually also at work when we *do* have knowledge about the truth of a given sentence. To elaborate on this point, I introduced the notion of epistemic validation or monitoring found in language comprehension research. Epistemic monitoring describes an automatic, effortless, “routine, nonstrategic validation process that detects knowledge violations” (Isberner and Richter 2014, p. 246). Importantly, next to activation and integration, this validation process is the third of three steps in our routine language comprehension processes. Against the background of the coherence mechanisms introduced in the previous section, epistemic validation can be understood as a coherence computation under conditions of matching knowledge.

8.3.3. The Truth in Intuitions

You can probably see my proposal coming: Apart from being part of a more complex comprehension machinery which can run more or less fluently, I submit that it is epistemic monitoring that routinely produces feelings of truth (in case of a positive validation outcome) and feelings of falsity (in case of a negative validation outcome). In fact, against this background FOT/Fs appear to be caused by something that seems as close to a psychologically realistic version of coherentist truth as we can get (Young 2018). I doubt that we can ask more of a connection between a psychological state and (abstract) truths. The presence of validated coherent corresponding references appears as an excellent psychological operationalization of truth. In fact, due to its clear isomorphism to coherentist truth it might be called “psychological coherence truth”.²⁴²

I am not claiming, however, that epistemic monitoring is the only supplier of FOT/Fs. As we have seen, FOT/Fs can also be (jointly) caused by ambiguous but coherent memory references as well as by a host of fluencies. In fact, I conjecture that in the right

²⁴²One might now wonder: Aren’t frameworks like the one of Thagard 1989 or Unkelbach and Rom 2017 inspired by coherentist theories of truth? Isn’t the isomorphism a consequence of this? This might well be the case. Still, I don’t see a problem with it: a theory developed for one domain might turn out to be (better) suited for another domain. In fact, I would even go a step further and claim that the reason why we came up with the coherentist theory of truth in the first place is because it taps a psychological reality—psychological coherence. This psychological reality is one precursor for an explication and formalization of truth (cf. Hookway 2002). It is because of *this* that Coherentism is appealing and intuitive as a theory of truth.

circumstances FOT/Fs can be caused by whatever causes affect in a way that mimics the dynamics of affect resulting from coherence and fluency. Insofar, some FOT/Fs turn out to be indeed more reliable than others (cf. Weinberg 2007; Nado 2014).

We find that those intuition theorists who have proposed an aetiological constraint for intuitions in the form of being based on one’s understanding were not far off the mark (e.g. Bealer 1999; Ludwig 2007; Sosa 2007a,c; Brogaard 2013, p. 279). Similar considerations apply to those making content constraints (e.g. to analytic, *a priori* or modal contents) that seemingly reduce the reliance on empirical contingencies and favour the possibility of validation based on “sheer understanding” (Sosa 2007a, p. 52; see also Bonjour 1998; Boghossian 2000; Boghossian and Bonjour 2001).²⁴³ In fact, such contents appear to be not too far from the plainly true (e.g. “there are more than fifty stars in the universe”) or false statements (e.g. “the earth is a perfect square”) mentioned in section 7.3.4 (Pennycook et al. 2018; Fazio, Rand, et al. 2019). Now, it turns out that FOT/Fs are in fact sometimes plausibly the products of a validation step in the comprehension process, which is just a more technical term for understanding. As such, it is well possible that intuitions that *are* based on understanding (certain kinds of contents) are more reliable than others, e.g. those based on a validation of more complex subject matters, on incidental processing fluency or on other subtle sources of affect. However, it is not that it is easy (or even possible) to tell the difference on the level of phenomenology—FOT/Fs can but must not be caused by validation (over specific contents) (cf. Topolinski and Strack 2009b; Chudnoff 2014a).

The emerging picture does not only provide a connection between intuitions and understanding, it also provides a connection between intuitions and beliefs. Now you might think: Of course! Intuitions lead to beliefs. Wasn’t this the whole point with the philosophical intuition debate (see section 3.1)? Indeed. Nonetheless, what I am getting at here is something else. Namely, that sometimes beliefs lead to intuitions. Remember, this was emphasised by Cohen (see end of section 3.3.1):

Belief that p is a disposition [...] normally to feel it true that p and false

²⁴³In this light, consider how Paul Benacerraf describes the dialectic between Kant and Logicism on the notion of synthetic *a priori* truths:

[I]n reply to Kant[’s doctrine that mathematical propositions are synthetic a priori], logicians claimed that these propositions are a priori because they are *analytic*—because they are true (or false) merely “in virtue of” the meanings of the terms in which they are cast. Thus, to know their meanings is to know all that is required for a knowledge of their truth. No empirical investigation is needed (Benacerraf 1981, p. 18).

that not-p, whether or not one is willing to act, speak, or reason accordingly. [...] The standard way to discover whether you yourself believe that p is by introspecting whether you are normally disposed to feel that p when you consider the issue. (Cohen 1992, p. 4)

That is, dispositional beliefs are (among other things) dispositions to experience FOT/Fs. I mentioned that coherence computations operate over a piece of incoming information and information in one's memory.²⁴⁴ Sometimes the information retained in memory will (not) cohere with incoming information and sometimes it will (mis)match with the incoming information. Now the memory contents over which the coherence computation and validation operate will plausibly include the subject's beliefs (among other things).²⁴⁵ Having a (mis)matching belief will ordinarily contribute to a positive (negative) validation outcome, possibly giving rise to an intuition experience that phenomenally expresses one's dispositional (mis)matching belief.²⁴⁶ Even in the absence of (mis)matching dispositional beliefs, one will be able to speak of *dispositions to believe* (Audi 1994) resulting from closely related beliefs in the form of coherent corresponding references. As a consequence, a belief that *A*, a belief that *B* and a belief that *C* will figure into a coherence computation over an incoming proposition *D*, and give rise to the intuition that *D*. Ultimately, this connection between understanding, intuitions, and beliefs obtains because the system that enables understanding, regularly produces intuitions and embodies dispositional beliefs is one and the same: memory.

Note that against the sketched background FOR/Ws concerning logical validity and FOR/Ws concerning the truth of a proposition or a whole set of propositions emerge as different facets of the same thing. In epistemic monitoring recently and currently encoded contextual information as well as reactivated world knowledge play a role. The

²⁴⁴A yet more encompassing framework on epistemic monitoring is the reality- and source-monitoring framework which, in a similar fashion, also operates over incoming non-linguistic sensory information (Mitchell and Johnson 2000; Schneider 2013). Disturbances in reality monitoring lead to highly distressing phenomena such as derealisation and depersonalisation (dokic2012bDokic and Martin 2015a; Bouzerda-Wahlen et al. 2013). This opens up the possibility to unite the views and to develop a unifying account of how we *in general* retain a sense of what is real (or true) (Singer 2013, p. 364). Perhaps the assertiveness and alleged pushiness of perceptual experiences is actually due to an affective overlay state signalling what's real and what's not. Insofar, affective states might be the source of all phenomenal assertiveness and pushiness. This is just a speculation, of course.

²⁴⁵That is, if one is willing to accept some form of representationalist picture of belief and grant that memory (e.g. semantic memory) will feature prominently among the ingredients of this picture.

²⁴⁶However, there are other factors that plausibly figure into the mentioned computation, possibly trumping the contribution of one's (matching) beliefs. If one, for instance, assumes that beliefs are grounded in semantic memory, then one can imagine that in some coherence computations other things than semantic memory will play a role, for instance, episodic memory.

difference between assessing inferences, single propositions and sets of propositions (e.g. texts) is the relative role of these two different sources of information in driving the validation. The operative processes, however, seem to be essentially the same. In other words, the validation of a conclusion against the background of premises, the evaluation of a stand-alone proposition and the evaluation of a proposition embedded in a network of propositions (e.g. in a text) differ only in the degree to which contextual information (premises, rest of the given text) has to be taken into account and to which one can rely on one's activated background knowledge during validation.

As it stands, FOT/Fs can phenomenally “mimick foundationalism” (Hookway 1993), being based on psychologically proximal truth-cues such as fluency and/or (validated) coherent references. Now, how good a proxy these cues are for actual truth and falsity will depend on the context and the portion of the subject's world model recruited by the context.²⁴⁷ The world model, on the other hand, is a function of the subject's biography and the informational environments traversed in the course of this biography. Ultimately, answering questions about the quality of our truth-tracking devices will depend on an examination of the evolutionary pressures that shaped the cognitive system embodying our world models. As Herzog and Hertwig aptly note:

Our cognitive machinery has evolved in the service of enabling us to navigate an often dangerous and uncertain world. How successfully we deal with this world depends, among other factors, on the fit between the cognitive machinery and environmental structures (Herzog and Hertwig 2013, p. 190)

The goal of this chapter was to make good on one of the advantages of Intellectual Affectivism: It sought to go beyond appearances and speculation when it comes to understanding intuitions and tried to fathom their actual causal determinants. The resulting overarching question that guided this chapter was: Do FOT/Fs result from properties that are good guides to actual truth and falsity? The exploration in this chapter led us from processing (dis)fluency (section 8.2) to (in)coherent corresponding references (section 8.3.1) to epistemic monitoring (section 8.3.2) as possible sources of FOT/Fs. This exploration has not only revealed a connection between intuitions, on the one hand, and understanding and beliefs, on the other, but also provided a tentative answer to the guiding question: yes, FOT/Fs regularly result from properties that are good guides to actual truth and falsity (section 8.3.3). The deployed truth-cues turn out to be good proxies for actual truth and the feelings that phenomenally represent

²⁴⁷As well as, perhaps, one's theory of actual truth.

truth (and falsity) appear as rather reliable trackers of actual truth and falsity. In other words, *phenomenal* truth and falsity seem to covary with the *actual* properties.

Naturally, this verdict can not stand without qualification: Of course, it depends. FOT/Fs are not exclusively produced by the described mechanisms. In fact, I conjectured that, in the right circumstances, FOT/Fs can be caused by whatever causes affect in a way that resembles the dynamics of affect resulting from coherence and fluency. To that extent, some FOT/Fs turn out to be indeed more reliable than others. On that score, those intuition theorists that have suggested aetiology and content restrictions were actually onto something. Now we can really see how these can contribute to epistemic damage control in a non-arbitrary way. On the other hand, it does not seem straightforward at all to establish whether these restrictions are *actually* satisfied for a given FOT/F—not on the basis of phenomenology alone, at any rate (see also Egler 2019).²⁴⁸

Be it as it may, I think this is a good tentative result. Intuitions appear to be the outputs of fairly reliable mechanisms (and, thus, they can be thought to be derivatively reliable themselves). At the same time, intuitions are not perfect arbiters of truth—they can be misleading. Against the established background, however, they (and their mechanisms) are fallible in ways that can be precisely predicted and tested. This advances our understanding of the mechanisms behind intuitions and points into directions for further theoretical and empirical investigation: Now we can go *beyond* a crude wholesale of intuitions' epistemic worth and can actually begin to test its limits in systematic ways. I cannot see what else we can realistically hope for when it comes to a psychological state.

²⁴⁸It is true that the FOT/Fs will likely not differ *qualitatively* (cf. Sosa 2007a, pp. 61 sq.). However, there might be some *quantitative* differences along the various gradable dimensions of FOT/Fs.

Concluding Affectivism: The End and the Beginning

9.1. Recounting the Journey

In the end, the present thesis can be summarised quite succinctly: It asks the question “What kind of states are intuition experiences?” Now, intuition experiences have a specific feature profile (chapter 2): they are 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content, and in 5.2) pushiness, 6) phenomenally epistemically valenced, and 7) nonvoluntary. A good answer to the guiding question thus could accommodate this feature profile, i.e. it could acknowledge and explain these features. Extant intuition theories provide the following answers: Eliminativism claims the term “intuition” has no extension—intuitions do not exist. Doxasticism claims they are doxastic states. Perceptualism claims they are states similar to perceptual experiences. We found the existing answers lacking based on roughly the following three arguments (chapter 3).

Insufficiency of Eliminativism (IE):

- (IE-P1) Eliminativism is true if it can accommodate the feature profile of intuitions.
- (IE-P2) Eliminativism cannot (and does not want to) accommodate the feature profile of intuitions.

(IE-C1) Therefore, Eliminativism is false.

Insufficiency of Doxasticism (ID):

(ID-P1) Doxasticism is true if it can accommodate the feature profile of intuitions.

(ID-P2) (If Doxasticism could acknowledge the feature profile of intuitions then it could explain them).

(ID-P3) Doxasticism cannot acknowledge the feature profile of intuitions.

(ID-C1) Therefore, Doxasticism cannot accommodate the feature profile of intuitions.

(ID-C2) Thus, Doxasticism is false (or at least incomplete).

Insufficiency of (Non-Literal) Perceptualism (IP):

(IP-P1) Perceptualism is true if it can accommodate the feature profile of intuitions.

(IP-P2) Perceptualism can acknowledge the feature profile of intuitions.

(IP-P3) Perceptualism cannot explain the feature profile of intuitions.

(IP-C1) Therefore, Perceptualism cannot accommodate the feature profile of intuitions.

(IP-C2) Thus, Perceptualism is false (or at least incomplete).

The rest of the thesis developed a new intuition theory: Intellectual Affectivism. The answer it gives to the guiding question is the following: Affectivism, as developed here, claims that intuitions are affective experiences, or more specifically: they are specific instances of epistemic feelings, feelings of truth and feelings of falsity. In order to set the ground for Affectivism, the psychological kind of feelings or affective experiences was introduced and characterised (chapter 4): affective experiences, of which bodily feelings such as bodily pain or pleasure and emotional feelings such as fear or joy are paradigmatic subclasses, are valenced, arousing, motivational and richly intentional by engaging in a division of representational labour with other mental states. Then the class of epistemic feelings was introduced and characterised (chapter 5). The thesis proceeded to make a case for epistemic feelings being affective experiences (chapter 6). Having established that, it went on to identify and analyse specific epistemic feelings as promising candidates for an identification with intuition experiences (chapter 7): feelings

of rightness and feelings of wrongness. It turned out that a propositional variety of these feelings, feelings of truth and feelings of falsity, has the same feature profile as intuition experiences. Thus, the claim went, positive intuitions are to be identified with feelings of truth and negative intuitions are to be identified with feelings of falsity. In virtue of these feelings being affective experiences, Affectivism cannot only acknowledge the features of intuitions but also explain them. Intuitions have the features they have for essentially the same reasons as bodily and emotional feelings have them—because they are (specific) affective experiences. The thesis thus established roughly the following argument:

Sufficiency of Affectivism (SA):

- (SA-P1) Affectivism is true if it can accommodate the feature profile of intuitions.
- (SA-P2) Affectivism can acknowledge the feature profile of intuitions.
- (SA-P3) Affectivism can explain the feature profile of intuitions.
- (SA-C1) Therefore, Affectivism can accommodate the feature profile of intuitions.
- (SA-C2) Thus, (by inference to the best explanation) Affectivism is true.

Finally, the last main chapter made good on one of the advantages of Affectivism (chapter 8): it went on an exploratory expedition into the mechanistic depths of the relationship between feelings of truth and falsity and actual truth and falsity. The expedition returned with a relatively favourable result.²⁴⁹ Let me now turn briefly to some of the implications of Affectivism.

9.2. Back to the Future of Affectivism

Some find intuitions mysterious. But intuitions understood as epistemic feelings become — as we have seen — a fairly well-defined target for philosophical and empirical investigation. There is no mystery here, only underexplored territory. Moreover, the fact that the existence of epistemic feelings is empirically well established might be used to persuade those sceptical about the existence of intuition experiences.

In our ontology of mind, affective experiences are not only agreed to be *canonical* mental states, i.e. mental states agreed to have a firm (theoretical) place in our mental ontology

²⁴⁹In the unlikely case you prefer a longer summary, have a look at the appendix A.

or psychology, but also canonical *phenomenal* states, i.e. mental states with a distinctive phenomenology. Identified as a kind of affective experience, intuitions thus not only inherit the firm grounding of a canonical mental state but their phenomenology can also be accounted for along the lines of a canonical phenomenology, namely affective phenomenology. This way, one would, for instance, not need to subscribe to the existence of *sui generis* cognitive phenomenology.

Now I was emphasizing the fact that Perceptualism is ontologically unstable, laying itself open to reductive attacks. Doxasticism seemed to be ill-equipped to lead a successful offensive along this line. However, forms of Non-Doxasticism might be better positioned. Affectivism is such a non-doxastic position from which, I have hoped to show, a successful run on perceptualist intuitions is viable. However, there are (at least) two potential points of contention here:

First, the Perceptualist might counter:

You think Perceptualism is ontologically unstable because it does not settle the ontological status of intuitions. Then what about affective experiences—what *are* affective experiences?

This brings up a legitimate point. I have said quite a lot about affective experiences but I did not *explicitly* answer the question for affective experiences which I sought to answer for intuitions: “What kind of states are affective experiences?” I don’t think that we absolutely need an answer to this question in order for Affectivism to work its (explanatory) magic. This is because we know a lot about affective experiences — knowledge we can directly apply to intuitions because intuitions are affective experiences — independently of whether we know what kind of states affective experiences are.

Be it as it may, here is my preferred answer to the question: affective experiences are *sui generis* kinds of states—they are what they are: affective experiences. I have told you quite a bit about the *unique* phenomenal and intentional features of affective experiences such as their phenomenal valence, motivationality and their involvement with bases. These aspects, it seems to me, set them apart from our other natural psychological kinds such as perceptual experiences or doxastic states. Surely, they *are* experiences and so, in virtue of that, they do share features with e.g. perceptual experiences. However, that does not make them the same *kind* of experience. Affective experiences as well as perceptual experiences, if you will, are subclasses of a superordinate psychological class: phenomenal states or experiences.

Now the Perceptualist (and Doxasticist) might object:

So you're saying Perceptualism is problematic because it either posits intuitions as *sui generis* or identifies them as members of a class that itself is *sui generis* such as seemings or presentations. However, you yourself seem to say essentially the same. You say intuitions are affective experiences. But then you say that affective experiences are *sui generis*. How is this different from saying that intuitions are *sui generis* seemings?

I admit that this objection gets the analogy *in strategy* right. However, there is a difference in how sound the strategy is in relation to *sui generis* seemings and *sui generis* affective experiences. This is because—as I hoped to show—we know a lot about affective experiences while we know only little about seemings. Affective experiences are *sui generis*—but they are *sui generis* in the same way perceptual experiences and doxastic states are *sui generis*: they are well-established canonical mental states—canonical phenomenal states, even.

In contrast, seemings are highly contested in their very existence—it is not clear at all what explanatory gain — apart from some perceptual metaphors — one gets by identifying a mental state as a seeming (e.g. Byerly 2012). Seemings understood as experiences are modelled on perceptual experiences and then, for one purpose or another, one attaches intuition experiences to the class of seemings as well.

In comparison, affective experiences harbour numerous subclasses such as bodily feelings (with its own subclasses), emotional feelings (with its own subclasses) as well as epistemic feelings (with its own subclasses).²⁵⁰ The subclasses of affective experiences share many concrete and relatively well-understood properties among each other—this is what establishes affective experiences as a coherent class of mental states. In fact, the subclasses of affective experiences have so much in common that the boundaries between the different subclasses appear at times blurred and porous.²⁵¹

So we observe a dramatic asymmetry here: seemings as experiences are not familiar kinds of states, there seem to be few independent reasons to postulate them except for one's wish to account for intuitions. On the other hand, we have many good reasons to postulate affective experiences. And these reasons are *wholly independent* from

²⁵⁰The subclasses of epistemic feelings might be e.g. “proper” epistemic, noetic, metacognitive and metaperceptual feelings (see e.g. Dokic and Martin 2015b).

²⁵¹This is why it does not seem implausible to conceive of epistemic feelings as another subclass of *emotional* feelings, alongside “basic”, social, moral and aesthetic emotions. This is why some prefer to use the label “epistemic emotions” (see footnote 126).

issues with intuition experiences.²⁵² Thus, by identifying intuitions as *sui generis* affective experiences we embed them deeply and *stably* into the existing net of our mental ontology.

However, as we have seen when e.g. discussing the relationship between affective experiences and bodily sensations, it is not that there are no *reductive* accounts of affective experiences out there. I think I have said enough about bodily sensations in chapter 4. Yet there is more on the reductionist menu, there is also Judgmentalism and Perceptualism about affective experiences or particularly about the subclass of emotional feelings. Judgmentalism had its ups (Solomon 1993, 2003b; Nussbaum 2001). The Judgmentalist Golden Age, however, lies in the past (see e.g. Deigh 1994; D’Arms and Jacobson 2003; Scarantino 2010). To see part of the reason why you could again go back to chapter 3.3 and replace “intuition” with “emotion”. I think many of the arguments would still hold.

This leaves Perceptualism. You perhaps remember what Deonna said:

It has been argued that emotions are very much like perceptions or even that they are kinds of perceptions. [...] The appeal of such a view, if true, is clear enough. It is the promise of illuminating a more elusive phenomenon by means of one that is slightly more familiar. (Deonna 2006, p. 29)

Now, it is true that Perceptualism about emotions has been trending of late (e.g. Tappolet 2000, 2016; Döring 2004, 2007; Prinz 2004b, 2006b). It has also been thoroughly criticised, however (e.g. Salmela 2011; Brady 2013; Dokic and Lemaire 2013; Kurth et al. 2018). In any case, it is important to differentiate here. Most emotion theorists do not endorse Literal Perceptualism about emotions. They are happy to stick to a mere analogy. In other words, they are Comparative or Quasi-Perceptualists, using the analogy as a conceptual tool. This brings us back to the unattractive Swampman/Orfolog-scenario. A mere analogy is useful for some purposes such as *describing* the (phenomenal) features but is not very useful in explaining them.

As a matter of fact, Perceptualists about emotions do not *need* Literal Perceptualism in order to understand emotions. This is because they can rely on a wealth of *emotion-specific* knowledge that is independent of the perceptual analogy. Often, in fact, they

²⁵²Relatedly, seemings and presentations as classes seem to be short hands for bundles of states from different kinds rather than psychological natural kinds (cf. Byerly 2012). This bundling can be based on e.g. their assertiveness or on some epistemic qualities (cf. Huemer 2001).

do not even *want* to make an ontological claim on the basis of the observed similarities with perception. Perhaps most of them would rest content with my suggestion that these similarities are not best explained because emotions are perceptions but because *both*, emotions and perceptions, are experiences—thus sharing certain features. This seems compatible with what Perceptualists about emotions usually rather want to do: they simply make an epistemic claim based on the analogy. In other words, many Perceptualists about emotions are not Ontological but Epistemic Perceptualists, suggesting to treat emotions and perceptions epistemically somewhat on a par. Consider Oliver-Skuse’s conclusion:

These challenges encourage limiting the scope of the analogy. For my purposes the perceptual analogy is most interesting if it is taken to stress a partial analogy between perception and emotions the core of which is that they stand in the same rational relations to beliefs that perceptions do. (Oliver-Skuse 2016, p. 28)

So, I said *most* emotion theorists do not endorse Literal Perceptualism, implying that there are some that do. Prinz seems to do that most clearly at times, and perhaps—more cautiously—Sabine Döring. It should come as no surprise that these versions of Perceptualism are the ones most liable to critique. However, suppose they can be made to work. And suppose further that they can be extended from emotions to *all* affective experiences. Well, then, perhaps it is worth “biting the bullet” and buy into Literal Perceptualism about affective experiences. In this case, the good news is that Affectivism about intuitions would give us a straightforward road to Literal Perceptualism about intuitions, together with a plausible rationale for an analogous perceptual epistemology. The uncanny creatures of Swampman and Orfalo can pack their bags and we can celebrate with Sam and Don. Hooray, everybody’s happy.

But now, back to reality: I agree that comparative work is helpful. However, as intuitions and affective experiences exhibit unique features, it is hardly the way forward when it comes to what is special about intuitions and affective experiences. I think the most promising way is to conceive of affective experiences as *sui generis* and develop accounts specifically *of* affective experiences.²⁵³ Fortunately, we have no lack of such accounts (e.g. Stocker 1983; Goldie 2000; Helm 2001; Roberts 2003; Deonna and Teroni 2012, 2015; Brady 2013; Weiss 2016; Mitchell 2018, 2019).

²⁵³Of course, we would need some subclass-specific specifications of these theories when we apply them to subclasses of affective experiences such as bodily, emotional and epistemic feelings.

Second, after having been quiet for a while—politely listening to me ramble on and on about her first point—the Perceptualist might raise her second objection:

As I see it, Intellectual Affectivism seems to be based on descriptive observations about our (epistemic) psychology. I now wonder: Doesn't it lend an *epistemological* flank? What I mean is: You seem to have straightened the ontology of intuitions—but what does it mean in terms of epistemology? Where does this leave us regarding justification?

This is an important question and I can see that it will have many sit on the edge of their seats now. Actually, I bracketed the question out intentionally to see where an inquiry of the more basic question “What kind of states are intuitions?” would take us. Now, I obviously cannot provide a definite answer to the epistemic question in the conclusion of a thesis—and so it will have to be for another time. I want to offer some remarks, however. First, I think the question has a clear *normative* dimension. To decide what counts as justification and what does not is not a purely descriptive but a normative endeavour. In this endeavour, people have variously gone internalist or externalist ways. As I already said, however, my aim here was precisely *not* to take a normative stance on intuitions. My aim was—it seems to me—more basic: I wanted to supply a picture of how things actually stand about intuitions, psychologically. So, by all means, I want to leave *normative* epistemological questions open. I thus embrace the mentioned flank and am curious to see what passes through.²⁵⁴

On the plus side, however, I might be able to say a few things on what the present picture might imply *given* that one adopts Epistemic Internalism or Epistemic Externalism. After all, I have said quite a few *descriptive* things about how FOT/Fs as intuitions relate to truth and falsity. I think this is relevant insofar as *ought implies can*—the way things stand needs to inform what we can normatively ask for.

Now, Internalists tend to take phenomenology as the decisive datum for justification (e.g. Koksvik 2011; Chudnoff 2013c; Bengson 2015). Here it seems that affective experiences and, in particular, intuitions understood as epistemic feelings do *appear* to us as justifiers. This is essentially why Gerald Clore and Karen Gaspar propose their “feelings-as-evidence hypothesis, which says that cognitive and affective feelings [...] may be experienced as internal evidence for beliefs that rivals the power of external evidence

²⁵⁴As I have noted in footnote 129 this is perhaps not unlike Perceptualism concerning its ontological flank. It might well be that Perceptualists want to leave the ontological question open as well and only reject doxasticist proposals along these lines. Perhaps, then, they would welcome my non-doxastic ontological proposal .

from the environment” (Clare and Gasper 2000, p. 26). In fact, against the established background, it is hard to see what could be a better or even alternative *phenomenal* basis for judgments of hedonic, evaluative and epistemic properties (see section 7.4). So *if* appearance is all that is needed to *be* justifiers, then I see no reason to deny this status to intuitions.

One thing that potentially needs adjustment however: Accounts of intuitive justification that give a central stage to phenomenology have been so far developed based on a perceptual analogy. Yet, on many counts, the phenomenology of affective experiences seems quite different from the phenomenology of perceptual experiences. This is not to preclude that there are important epistemic analogies between intuitions and perceptual experiences despite the fact that, *pace* Perceptualism about emotions and intuitions, they might be metaphysically distinct. It is rather to say that what we might need for intuitions is not (or not only) an account of perceptual justification but an account of *affective* justification.

On the other hand, what can be said for Externalists? For justification to obtain, Externalists tend to ask for a reliable process on whose basis a belief is formed (e.g. Bealer 1996; Sosa 2007c). Now, I think that chapters 5 and 8 have much to recommend themselves to Externalists. It seems that, by and large, an Externalist can rest content with what he finds in terms of reliable processes when it comes to FOT/Fs—with some qualifications, of course, but, then again, everything needs qualifications. Brian McLaughlin’s comment on the passage of Clare and Gaspar above might be thought somewhat of a summary of what we have uncovered in these chapters:

But what is most important here is not that people experience the feelings as evidence, but that the feelings function for our belief systems as evidence; our belief-fixing mechanisms respond to them as evidence. [...] Indeed, [at times] the evidence seems to be treated by the belief system as weighty enough to trump visual evidence. [...] What is represented by [...] [a] feeling may or may not be the way it is felt as being. Thus, [...] feelings can be illusory. [...] Still, such feelings function as evidence for the belief system. Indeed, they are normally fairly reliable, though, of course, some are more reliable than others. (McLaughlin 2010, pp. 151-153)

We can thus say that there is a wealth of empirical work on feelings, epistemic feelings and specifically on FOT/Fs. These await an analysis against the background of externalist frameworks. As we have seen, so far there seems to be reason for qualified optimism.

Furthermore, conceiving of intuitions as a kind of affective experience resonates with virtue epistemology not only in its virtue reliabilist but also in its virtue responsibilist variety (e.g. Fairweather and Zagzebski 2001; Sosa 2007c; Battaly 2008; Candiotti 2017). The latter is more akin to virtue ethics which is concerned with virtuous agents marked by well-tempered *affective* dispositions. This is also insofar significant as Affectivism, in contrast to Perceptualism, casts a quite different light on the possibility to educate intuitions, i.e. to learn and to improve when it comes to them (see also Hogarth 2001, 2011). If one takes the perceptual analogy to be an analogy with perception as traditionally understood in philosophy — with an emphasis on modularity — then, there is relatively little room for perceptual and consequently intuitive learning.

Yet, this seems not to be what we observe with feelings nor with intuitions: Surely, as affective and cognitive “illusions” suggest, they are somewhat headstrong pupils at first, especially if compared to teacher’s pet belief. However, with patience, time and some emotional intelligence they can be made to embody what is properly called *virtue*. This is also why expert intuitions are possible. So, in contrast to what a perceptual analogy would seem to suggest, the system producing intuitions does not come to us as a rigid encapsulated system that, luckily, also happens to be natively reliable. It comes to us, luckily, as a system capable of moderately paced learning, especially in the case of feelings that are hooked up to a relatively plastic cognitive system such as memory.

Affectivism, as developed here, has also the potential to cast a new light on the somewhat confrontational dialectic between experimental philosophers, who are sceptical about intuitions on empirical grounds, and those in favour of philosophical armchair practice (e.g. Sosa 2007b; Weinberg 2007; Weinberg and Alexander 2014; Bengson 2013). This debate has been often couched in terms of a forced choice between the wholesale rejection or acceptance of the epistemic worth of (philosophical) intuitions. Now, it seems to me that Affectivism is in a good position to mediate between experimental and armchair philosophers in a way that might satisfy both. Instead of trying to dismiss the empirical findings of experimental philosophers we can *integrate* them in an empirically informed picture of intuitions—and we can do so without throwing in the intuitive towel. This is because we can go beyond wholesale rejection or acceptance now: we can conceive of intuitions in a way sympathetic to armchair philosophers *and* make nuanced and empirically testable predictions about them and their epistemic powers.

Affectivism harbours some potential for unifying hitherto disparate fields and phenomena. First, it brings together perspectives on intuitions from philosophy, psychology

and folk psychology. Second, it can offer analogous accounts for intuition experiences in other domains understood as FOR/Ws directed at other things than propositions. It can, for instance straightforwardly explain logical intuitions as conclusion-directed FOR/Ws. Similar things can be said for linguistic intuitions and FOR/Ws directed at e.g. sentences or other linguistic items. I think it is also worth exploring that FOR/Ws might be at the bottom of many aesthetic, social and moral intuitions. And even if this might go too far, Affectivism makes us realise that intuitions as affective experiences have many siblings. So even if FOR/Ws cannot do the explanatory job alone, then we could still call upon some other family members to help out. This opens the possibility of a unifying account of intellectual (logical, modal, metaphysical etc.), linguistic, aesthetic, social and moral intuitions under the broader banner of affective experiences. Finally, understanding intuitions as affective experiences might be able to bridge the “intuitive gap” and provide us with means to identify (mental states homologous to) intuitions in animals and young children,

Furthermore, identifying intuitions as a form of affective experience can shed light on hitherto neglected features of intuitions. It has, for instance, the potential to solve the puzzle of why we engage in epistemic activity in the first place. Why do we seek truth, understanding and meaning? Many philosophers consider epistemic goods such as truth as intrinsically valuable. Now we might have uncovered part of the story concerning the *psychological* and *phenomenal* basis of this assessment: The intrinsic phenomenal positivity or negativity of epistemic feelings *is* intrinsically rewarding and motivating (Carruthers 2017b). In other words, the psychological motivation behind the pursuit of truth might be essentially similar to the pursuit of happiness.

There are important practical and *political* implications as well: the affective nature of a big chunk of our “internal epistemic guides” lends itself to biases and exploitation by epistemic predators such as populists and propagandists. The analysis that these “just exploit people’s feelings” is on the right track but does not hit the mark. In truth, what is exploited are the very psychological mechanisms behind what we take to be true. Neglecting this affective dimension of our epistemic activity strips us of the means to adequately deal with socially highly costly phenomena such as fake news and conspiracy theories. Affectivism, I suggest, can take us the first steps towards filling this lacuna. We have seen that there are ways to educate intuitions.

A Brief History of Affectivism

We have arrived at the very end of this thesis—in the appendix. Thank you for making it this far. In the course of this thesis, I have tried to make the case for Intellectual Affectivism and go some first steps in exploring its merits. In case you need a refresh, this appendix is supposed to serve you as an extended summary of what happened.

A.1. Chapter 2

In chapter 2 I delineated the target state of which Affectivism is a theory from other things we call “intuition”. I went on to give “intuitions” in my sense—intuition experiences—a recognisable face: Intuition experiences are occurrent conscious mental states that are (partially) characterised by their specific (but not necessarily *sui generis*) phenomenology—in other words, there is something it is like to have an intuition experience. This phenomenology is part of the specific “feature profile” of intuition experiences. In the remainder of chapter 2 I used examples to zero in on this feature profile of intuitions. I pointed out that intuitions are not only phenomenally conscious but also intentional states, often taking a proposition or propositional content as their intentional object. This content they represent assertively, i.e. as true or false. And they motivate or push the subject to assent or dissent to what they represent as true or false. In doing so, however, they fall short of fully committing the subject to their contents, akin to perceptual experiences and unlike beliefs and judgments. Furthermore, intuition experiences are

gradable in two ways: On the one hand, similar to the content of perceptual experiences and imaginings, they can be more or less determinate in the way they represent their contents. On the other hand, they can push you to assent or dissent more or less strongly. A peculiar feature of intuition experiences is that they exhibit phenomenal epistemic valence, i.e. they can directly represent their contents either as true (positive intuitions) or as false (negative intuitions). Positive intuitions feel genuinely different from negative intuitions even if their contents are the same. They constitute phenomenal polar opposites. Finally, intuition experiences are non-voluntary, that is, intuitions are not under voluntary control but happen to one. In short: intuition experiences are 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content, and in 5.2) pushiness, 6) phenomenally epistemically valenced, and 7) nonvoluntary.

A.2. Chapter 3

After outlining this feature profile, I put it on the wish list for a good theory of intuitions and went shopping through the aisles of extant philosophical intuition theories in chapter 3. What did we find there on the subject matter of intuitions? We found that philosophers are mostly concerned with the epistemology rather than the ontology of intuitions. Concerning the latter we found Eliminativism lurking in a corner, Doxasticism — the majority view — occupying most of the shopping shelves and then, finally, we arrived at an assortment of fairly popular niche products: Perceptualism. Eliminativism about intuitions states that whatever it is that we call “intuitions” and whatever we think they are, they — like Swampmans and Orofolei — do not exist. This appeared to be a hard sell since at least some things that we call “intuitions” certainly do appear to exist. For instance, sometimes we call beliefs or judgments “intuitions” and most people — except, perhaps, for some *intuition-unspecific* Eliminativists — seem to agree that beliefs and judgments exist. Or aren't you *judging* right now that Affectivism is on the wrong or the right track? Then we saw that in order to demystify intuitions, Doxasticists use exactly this datum about beliefs and judgments to their advantage and claim that intuitions *are* doxastic states, i.e. beliefs, judgments or something in the ballpark. In the course of our discussion, Doxasticism indeed turned out to be a fair description of *some* mental states that we call “intuitions”. However, it also turned out to be a poor characterization of *intuition experiences*. This is simply because doxastic states such as beliefs or judgments do not fit our wish list of features — particularly the *phenomenal* features — outlined in chapter 2. So instead of getting us—as would appear the promise—into the coveted

Samuel/Donald-scenario, Doxasticism gets us into a situation where we try to explain observations about Donald by classifying him as a bat. That explains *some* observations about Donald—but mostly, it *mischaracterises* Donald. Despite being called a bat, Donald can't fly, can't echolocate and is not particularly useful for hitting a baseball. We have seen that Perceptualism takes the phenomenal impotence of Doxasticism as a point of departure and make their *start* from the phenomenal features of intuitions. They proceed by describing and elaborating the features of intuitions, using perceptual experiences as an analogy. Arguably, this approach is more helpful than Doxasticism for acknowledging the phenomenal features of intuitions. The deficit of Perceptualism, I argued, lies in its inability to *explain* these features. Perceptualism claims that intuitions are in some ways *similar to* perceptions but they do not claim that intuitions *are* perceptions: intuitions resemble but are not perceptions. So what *are* intuitions and what explains that intuitions are the way they are? Here Perceptualism does not provide a satisfactory answer. Ultimately, Perceptualism seems to bring us into a situation that can be located somewhere in between a Swampman- and an Orfalo-scenario. We might be able to state that Swampman and Orfalo have such and such features and try to make sense of them by associating them with knowledge about other things, say, Donald, *homo sapiens* or bats. Stating these *similarities*, however, appears to be of relatively little explanatory power to explain the observations made. To recap: Perceptualism does a good job in describing the features that intuitions appear to have. This sets it favourably apart from Doxasticism. A major drawback of Perceptualism is, however, that it either posits intuitions as a *sui generis* kind of states or puts intuitions into a class of mental states (e.g. seemings or presentations) whose postulation is contested. Such a move does not *explain* the features assigned to intuitions in a satisfying way but rather posits another mystery. Most versions of Doxasticism (were they otherwise sound) would not suffer from this drawback since (part of) the specific nature of intuitions could be explained by appeal to the rather well-established kind of state that intuitions are (e.g. judgments or beliefs), about which we arguably know more than about intuitions. Against this background, a theory that 1) allows intuitions to have the right features (as Perceptualism does) *and* 2) *explains* these features by appeal to a kind of state that is relatively well-established and about which we know more than about intuitions (as Doxasticism does) would be preferable to theories that only manage one of the two. This is where Affectivism comes in.

A.3. Chapter 4

Chapter 4 laid the groundwork for a fresh look at intuition experiences in the form of Intellectual Affectivism. In order to do that, it established a first motivation to examine the hypothesis that intuitions are what Affectivism claims them to be: feelings. One part of the chapter was dedicated to clarifying what I mean by feelings: not qualia, not bodily sensations—but affective experiences. After that, I explored what we know about affective experiences, specifically their phenomenal and intentional features. I explained that the central feature of affective experiences is their phenomenal valence, i.e. the felt positivity or negativity of certain experiences. Affective experiences are essentially *valenced* experiences. This is what sets them apart from mere qualia and bodily sensations. Another characteristic phenomenal aspect of affective experiences is felt arousal: during an affective experience, the subject feels a more or less localised increase or decrease (i.e. *change*) in level of activation, energy or excitement. Both valence and arousal, are closely associated with another crucial feature of affective experiences: they are directly motivating and so move us to behave or act in more or less stereotypical ways. So affective experiences are valenced, arousing and motivational. Furthermore, they are gradable along these dimensions: they can be more or less arousing as well as more or less positive or negative and motivating. when it comes to intentionality, feelings have two parts: a *particular object* and a *formal object* where the feeling represents the former as bearing the latter. In this context, the formal object refers to feeling-specific properties such as the painfulness of pain, the fearsomeness of fear, the funniness of amusement, etc. (section 4.4.2) Representing something to be painful or amusing constitutes an evaluation which is *phenomenally grounded* in the *valence* of affective experiences (section 4.4.3). Concerning their particular objects, affective experiences can be in principle about all kinds of things, ranging from bodily events and parts to objects in the physical and imaginary realm (including fictions and propositions themselves), states of affairs in past, present and future to propositional contents (section 4.4.1). To accommodate this broad range of intentional objects, feelings rely on a division of representational labour with other mental states such as bodily sensations, perceptions, judgments, memories, imaginings etc. This *base* of a feeling is a mental state (or a set of states) that supplies a feeling with its particular object: the feeling does not access the particular object directly but through other mental states. Putting these explanatory resources on the table went some way in demonstrating that we know quite a bit about affective experiences. Thus, classifying something as affective experiences would amount to something akin to the Samuel/Donald-scenario. As a consequence, Affectivism about a mental state, if other-

wise successful, emerges as a good thing. The last section of chapter 4 probed the idea that intuition experiences might be affective experiences by juxtaposing their respective features. What we got was this: On a closer look, intuition experiences appear quite similar to affective experiences. This is reminiscent of the reason why Perceptualists exploit perceptual experiences as an analogy to elaborate features of intuitions: intuition experiences are similar to perceptual experiences. There they stop, however. What explains this similarity is not that intuition experiences *are* perceptual experiences. Even Perceptualists dismiss such Literal Perceptualism as counter-intuitive. However, this makes Perceptualist theories—be they Comparative or Quasi— explanatorily impotent for the features of intuitions. In the same vein, Affectivist analogues of Comparative and Quasi-Perceptualism appear as non-starters for they would be haunted by essentially the same problems. This recommends exploring the most straightforward explanation for the similarity between intuitions and affective experiences: it's because intuitions just *are* affective experiences. This got Intellectual Affectivism started.

A.4. Chapter 5

Now, affective experiences appear as a diverse class of mental states with distinct subclasses, the most paradigmatic of which might appear as rather inhospitable places for intuitions. Chapter 5 sought to put this impression into perspective. To do so I introduced a class of mental states that would appear as a rather comfortable home for intuition experiences: epistemic feelings. Epistemic feelings have been broadly described as “feelings that enter into the epistemic processes of inquiry, knowledge and metacognition” (de Sousa 2008, p. 189). I first gave the reader a pre-theoretical grasp for some prominent instances of epistemic feelings such as the feeling of knowing (FOK) and the feeling of familiarity (FOF) (section 5.1 and 5.2). Then I dived into the research of what brings about these feelings. In the course of discussing the FOK we learned about the nowadays unfashionable Direct Access model and then focused on the contemporary majority view: the Heuristic Inference model that posits an inferential process which produces FOKs on the basis of cues. These cues are the familiarity with elements of the information query (cue familiarity, section 5.3.1.3), the amount of retrieved non-target information (accessibility, section 5.3.1.4), as well as the speed with which the implicated perceptual and memory processes operate (processing fluency, section 5.3.1.5). Furthermore, these cues are not mutually exclusive but complimentary when it comes to bringing about the FOK, making it multiply determined (section 5.3.1.7). In the course

of this section we have also made our acquaintance with a functional analogue of FOK in the form of introspective fluency (the introspective observation of some content or response coming quickly to mind, see section 5.3.1.5) and with experiential relatives of the FOK: the feeling of information (un)availability (FOA/U), i.e. the feeling that more information is (un)available for access (section 5.3.1.6), and the TOT (the feeling that one is in possession of a relevant information but is currently unable to produce it, section 5.3.1.8) whose mechanisms can be understood largely analogously to FOK-mechanisms as being multiply determined by cues such as perceptual fluency and the accessibility of information related to a stimulus. Concerning the FOF we saw mechanisms at work similar to the FOK: it appears to be multiply determined by cues such as perceptual fluency and the accessibility of information related to a stimulus. Furthermore, discussing the FOF enabled us to add discrepant processing fluency to our conceptual repertoire of mechanisms. Work on FOF has brought to the fore is that fluency cannot only pertain to different processes but that the fluency pertaining to a specific process can take different forms. Specifically, it has become apparent that the effects of objective fluency, as well as the phenomenology of fluency-based epistemic feelings, tend to be more pronounced and salient if a process is not only fluent but *discrepantly* fluent, i.e. if a process is not “just fast” (absolute fluency) but fast relative to something else (relative fluency). We have seen that a process can be discrepantly fluent in several complementary ways: It can be discrepantly fluent relative to 1) the rest of the currently ongoing processing background, i.e. a specific process can be fast and therefore fluent relative to *other* ongoing processes, or 2) relative to an expectation concerning the speed of the specific processing, whereas this expectation is formed 2.1) outside or 2.2) within the present context. In the special case of FOF it has been argued that the FOF is most pronounced when the perceptual processing of a stimulus is fast in the sense of being faster than expected for comparable stimuli. That is, the processing is fast relative to a stored standard or expectation for how much time perceptual processing usually takes for stimuli of a similar kind. Finally, in section 5.3.3 we touched upon an influential *general* framework which understands (some) epistemic feelings as a form of metacognition. We learned that some epistemic feelings have been proposed to be metacognitive in that they evaluate our cognitive processes, activities and dispositions. This kind of *evaluative* metacognition contrasts with *metarepresentational* metacognition referring to representations of one’s own representations. The emphasis of this section was to point out what the metacognitive framework does *not* imply for epistemic feelings. First, not all epistemic feelings have to be metacognitive feelings, at least not in the same sense. Some epistemic feelings seem to indeed concern the subject’s mental constitution while

others seem to concern things in the external world. Second, even if metacognitive feelings might be somehow about cognitive states and processes, metacognitive feelings do not have to be *explicitly* about these cognitive states or processes. As already noted: they are *not* metarepresentational. In a similar vein, metacognitive feelings also do not have to be explicitly self-referential, i.e. consciously making reference to the subject who is the bearer of the evaluated states and processes. On any count, metacognitive feelings might plausibly be all these things *implicitly*. Surely, from the ways metacognitive feelings evaluate, we can reconstruct implicit references to the evaluated and motivated cognitive processes as well as the subject and its mental constitution.

A.5. Chapter 6

After we acquired a pre-theoretical grasp on the phenomenology of epistemic feelings and accumulated some knowledge about their mechanisms, we wondered: What kind of state are epistemic feelings? Chapter 6 made the case for epistemic feelings as a subclass of affective experiences. Based on empirical findings reviewed in section 6.2 and 6.3 as well as a mix of theoretical and phenomenological considerations interspersed throughout section 6.4 and 6.5.1 I tried to establish that epistemic feelings are affective experiences and analyse them along these lines (section 6.5). In section 6.2 we have seen that epistemic feelings covary with interoceptive changes, variations in SCR and facial muscle activity, all well-established bodily proxies for the affective properties of valence and arousal. Furthermore, positive epistemic feelings lead to increased liking and can serve as positive affective primes—behavioural proxies for the presence of valence. Additionally, performance capitalising on positive epistemic feelings are enhanced in a positive mood. On the negative side, performance that relies on positive epistemic feelings is impaired under conditions of negative mood and depression. Further behavioural evidence comes from reviews of the influences of emotional and epistemic feelings on judgments. They bring to light that both kinds of feelings influence judgments in a similar way and that similar factors moderate the reliance on them when making judgments. Coming from the affective to the epistemic side: Induced affect co-varies with judgments of seemingly non-affective properties such as truth, grammaticality, logical validity and semantic coherence. In section 6.3 I first made the case for the covariation between epistemic feelings and affective properties being not just a correlation but a constitution relationship (section 6.3.1). For that, I presented studies that observe false positives of epistemic properties on the basis of incidentally induced affect. That is, inducing

non-diagnostic affect leads subjects to incorrectly judge that an epistemic property is present. This speaks in favour of an affective constitution of epistemic feelings. Secondly, in section 6.3.2 I made the case that the affect in question is *conscious*. It thus not only causally biases epistemic behaviour but phenomenally constitutes epistemic feelings that provide conscious guidance for the subject's epistemic behaviour. To make this idea plausible I recounted studies where the following happens: the experimenters make the subject believe that the affect they experience in a given epistemic task is *not* diagnostic for the presence of an epistemic property. As a consequence, the subject loses her ability to accurately detect the epistemic property. This does not only indicate that epistemic properties are at times detected on the basis of affect but also that the affect in question is conscious. Section 6.4 considered how Affectivism about *epistemic feelings* can contribute to an understanding of several aspects of them. It can give us a more general perspective on the way some epistemic feelings are metacognitive without being explicitly about other mental states or about the subject. Furthermore, it can help us to solve a puzzle about how process-specific fluency can have quite process-unspecific influences on our behaviour: by giving rise to *transmodal* affect. Against the background of affective experiences being valenced, it also becomes easy to see why the class of epistemic feelings harbours many phenomenal polar opposites. Section 6.5 started out by illustrating the relatively mild affective phenomenology of epistemic feelings (section 6.5.1): that they are usually (but not always) mildly valenced and arousing and that the behaviours they motivate are often of a covert mental or cognitive nature. In doing so it also strived to remedy some worries about the idea that epistemic feelings are affective experiences. Finally, section 6.5.2 was concerned with the affective intentionality of epistemic feelings. As all affective experiences, epistemic feelings have particular objects, formal objects and bases that provide their particular objects under base-specific aspectual shapes. These aspectual shapes, in turn, might themselves vary in their quality such as their determinacy. The feeling-specific formal objects of epistemic feelings are epistemic properties broadly construed: familiarity, uncertainty, understanding, knowing, rightness etc. As feelings in general, epistemic feelings “encode”, “localise” or “embed” value or disvalue in the form of their determinate feeling-specific property within the contents of their bases and motivate us to behave accordingly. These feelings, in other words, epistemically evaluate their particular objects.

A.6. Chapter 7

Chapter 7 marked the conclusion of Intellectual Affectivism by identifying intuitions more precisely with not just affective experiences and not just epistemic feelings but with *specific* epistemic feelings. Its central claim was that positive intuitions and negative intuitions are specific varieties of feelings of rightness (FORs) and feelings of wrongness (FOWs): feelings of truth (FOTs) and feeling of falsity (FOFs, from now on). It proceeded in the following way: section 7.2 provided a first feel for FOR/Ws. In section 7.3, then, we went to a FOR/Ws science exhibition. We have seen that the science that is explicitly concerned with FOR/Ws is as yet in its infancy and that in relation to FOR/Ws various things can be “right” or “wrong” in various ways (section 7.3.1). I then introduced the distinction between performance- and object-directed FOR/Ws which informed our take on explicit FOR/Ws research (section 7.3.2). Making this distinction made us realise that the mentioned research has been so far conducted with a strong metacognitive focus, somewhat neglecting FOR/Ws that are not as obviously concerned with cognitive performance. However, looking closer, we discovered that there are findings on feelings of logical validity that are plausibly construed as a variety of FOR/Ws. In this connection, we learned that kinds of rightness can evoke liking. Then we dealt with a kind of rightness that is central for the present thesis: truth. I introduced research on the truth effect, a tendency to judge a statement as more likely true as a result of various experimental manipulations. This effect is usually studied in relation to ambiguous statements and the most prominent condition under which it occurs is repetition: subjects judge a statement as (more likely) true if they have encountered it before. This repetition-based truth effect is robust across many conditions. In fact, it even occurs if the statements are *not* ambiguous. That is, the truth effect can even be observed when subjects have either knowledge or perfectly reliable testimony on the truth status of a statement. In section 7.3.5 I introduced the dominant theory about the common cause of the truth effect across manipulations of different variables such as previous exposure, readability and rhyming. This theory states that the proximal cause of the truth effect is a familiar construct: processing fluency. Based on the previous discussion of processing fluency as the cause of epistemic feelings and its relation to affect, I argued that, in fact, it is not processing fluency *per se* that is responsible for the truth effect. Rather it is the *result* of processing fluency: affect. Following this line of thought, I tried to make the case that what we observe in the truth effect are actually the workings of two epistemic feelings which are proposition-directed varieties of FOR/Ws: feelings of truth and feelings of falsity. After this, I went on to analyze FOR/Ws drawing on

the resources introduced in the previous chapters. In its course I delineated FOT/Fs among FOR/Ws as those FOR/Ws that take propositions as their particular objects and represent them as right or wrong, amounting to an evaluation of the propositions as true or false. More specifically, FOT/Fs turned out to have the following features: they are phenomenally conscious and intentional states that take propositions as their intentional objects. They represent the content of these propositions assertively, i.e. as true or false and they motivate or push—sometimes stronger, sometimes weaker—the subject to assent or dissent to what they represent as true or false. In doing so, however, they—as affective experiences in general—fall short of fully committing the subject to their contents. Furthermore, FOT/Fs are, depending on their base, sometimes content-gradable. One of the characteristic features of FOT/Fs is that they exhibit a phenomenal epistemic valence, i.e. they can directly represent their contents either as true (FOTs) or as false (FOFs). This phenomenal epistemic valence, as well as their assertiveness, is grounded in the general phenomenal valence of affective experiences. This valence also explains why FOTs and FOFs are phenomenal polar opposites and that FOTs feel genuinely different from FOFs even if their contents are the same. Finally, FOT/Fs—as all affective experiences—are non-voluntary, that is, they are not under voluntary control but happen to one. Call this (part of) the feature profile of FOT/Fs. Here is a condensed version: FOT/Fs are 1) intentional, 2) assertive, 3) motivational, 4) non-committal, 5) gradable in 5.1) (sometimes) content and 5.2) pushiness, 6) phenomenally epistemically valenced and 7) non-voluntary. It came to the fore that this feature profile of FOT/Fs is identical with the one of intuitions. I argued that this is no coincidence and that there is a straightforward explanation for this observation: the two feature profiles are identical because intuition experiences are identical to FOT/Fs. This secures the claim of Intellectual Affectivism.

As mentioned several times by now: a theory of intuition experiences that 1) acknowledges their features (as Perceptualism does and Doxasticism fails to do) *and* 2) *explains* these features by appeal to a kind of state that is relatively well-established and about which we know more than about intuitions (as Doxasticism would be able to do, were it otherwise sound) would be preferable to theories that only manage one of the two. Here I have developed such a theory: Intellectual Affectivism. It acknowledges the features and explains them by appeal to a fairly well-established kind of state: affective experiences. Affectivism is yet more precise: It states that positive and negative intuitions are identical to feelings of truth and feelings of falsity, a variety of feelings of rightness and feelings of wrongness that have bases with propositional contents, such as propositional

imaginings. These feelings, alongside with feelings of familiarity and feelings of knowing, belong to the class of epistemic feelings. Thus, intuitions are epistemic feelings. Alongside with bodily and emotional feelings, epistemic feelings, in turn, are a subclass of affective experiences and can be analysed along similar lines. Drawing on theoretical and empirical resources for FOT/Fs, FOR/Ws, epistemic feelings and affective experiences puts us in a good position to account for the outlined feature profile of intuitions. It turns out that intuition experiences have their features for essentially the same reasons bodily and emotional feelings have them: because they are affective experiences.

A.7. Chapter 8

The goal of chapter 8 was to make good on one of the advantages of Intellectual Affectivism: It sought to go beyond appearances and speculation when it comes to understanding intuitions and tried to fathom their actual causal determinants. The resulting overarching question that guided this chapter was: Do FOT/Fs result from properties that are good guides to actual truth and falsity? Against this background, I considered four questions in section 8.2: 1) Is fluency a good proxy for truth? 2) What causes fluency and is this a good indicator of truth? 3) What is it that is fluent and what does it do? 4) What, other than fluency, causes affect and leads to FOT/Fs? To summarise the tentative emerging answer to question 1) and 2): since fluency is a typical product of prior exposure, fluency emerges as a relatively good proxy for truth. Delving on question 3) brought to the fore that fluency explanations taken for themselves are in general quite thin. They do not tell us much about what it is that is fluent and what it does. But much will depend on this. The search for mechanisms to flesh out such explanations, potentially even independently of fluency, brought us to question 4). I pointed out that FOT/Fs like other epistemic feelings will plausibly be multiply determined—not only by fluency. Considering what we know so far, brought to the fore another plausible candidate property that itself often leads to fluency: coherence. What is coherence? That is, how are we to understand it in terms of psychological processes? Section 8.3.1 went some way in answering this question. We have first considered an alternative, mechanistic explanation of the repetition-based truth effect. This explanation relies on the idea that a previously encountered statement activates previously established corresponding references in memory. If these references are coherent, then this leads the subject to judge the statement as more likely true. On the other hand, if the references are incoherent, then this leads the subject to judge the statement as more

likely false. I suggested a plausible way to understand how the subject becomes aware of her coherent/incoherent corresponding references: coherence or incoherence directly (and indirectly via fluency) triggers positive/negative affect. In other words, the coherence computations over corresponding references lead to FOT/Fs. Section 8.3.2 took this explanation a step further and motivated the idea that the coherence mechanism outlined in the previous section is actually also at work when we *do* have knowledge about the truth of a given sentence. To elaborate this point I introduced the notion of epistemic validation or monitoring found in language comprehension research. Epistemic monitoring describes an automatic, effortless, “routine, nonstrategic validation process that detects knowledge violations” (Isberner and Richter 2014, p. 246). Importantly, next to activation and integration, this validation process is the third of three steps in our routine language comprehension processes. Against the background of the coherence mechanisms introduced in the previous section, epistemic validation can be understood as a coherence computation under conditions of matching knowledge. Finally, 8.3.3 showed that against the background of the outlined mechanism there is a connection between intuitions, on the one hand, and understanding and beliefs, on the other. This connection obtains because the system that enables understanding, regularly produces intuitions and embodies dispositional beliefs is one and the same: memory. More importantly, the section outlined a tentative answer to the guiding question of the chapter: yes, FOT/Fs *regularly* result from properties that are good guides to actual truth and falsity. These psychological properties turn out to be good proxies for actual truth and the feelings that phenomenally represent truth (and falsity) appear as rather reliable trackers of actual truth and falsity. In other words, *phenomenal* truth and falsity seem to covary with the *actual* properties. I went on to qualify this verdict: FOT/Fs are not the exclusive products of the described mechanisms. In fact, in the right circumstances, FOT/Fs can be caused by whatever causes affect in a way that resembles the dynamics of affect resulting from coherence and fluency. Insofar, some FOT/Fs turn out to be indeed more reliable than others. On that score, those intuition theorists that have suggested aetiology and content restrictions were actually onto something. On the other hand, it does not seem straightforward at all to establish whether these restrictions are *actually* satisfied for a given FOT/F—not on the basis of phenomenology alone, at any rate.

References

- Akins, K. (1996). “Of Sensory Systems and the ”Aboutness” of Mental States”. *The Journal of Philosophy* 93.7, pp. 337–372 (cit. on p. 162).
- Alter, A. L. and Oppenheimer, D. M. (2009). “Uniting the Tribes of Fluency to Form a Metacognitive Nation”. *Personality and Social Psychology Review* 13.3, pp. 219–235 (cit. on pp. 135, 156, 157, 166).
- Alvarez, M. (2017). *Reasons for Action: Justification, Motivation, Explanation*. In collab. with E. N. Zalta. URL: <http://plato.stanford.edu/archives/fall2014/entries/intuition/> (visited on 10/11/2015) (cit. on p. 63).
- Alves, H., Unkelbach, C., Burghardt, J., Koch, A. S., Krüger, T., and Becker, V. D. (2015). “A Density Explanation of Valence Asymmetries in Recognition Memory”. *Memory & Cognition* 43.6, pp. 896–909 (cit. on p. 157).
- Anderson, A. K., Christoff, K., Stappen, I., Panitz, D., Ghahremani, D. G., Glover, G., Gabrieli, J. D. E., and Sobel, N. (2003). “Dissociated Neural Representations of Intensity and Valence in Human Olfaction”. *Nature Neuroscience* 6.2, pp. 196–202 (cit. on p. 95).
- Andow, J. (2015). “How “Intuition” Exploded”. *Metaphilosophy* 46.2, pp. 189–212 (cit. on p. 23).
- (2017). “A Partial Defence of Descriptive Evidentialism About Intuitions: A Reply to Molyneux”. *Metaphilosophy* 48.1-2, pp. 183–195 (cit. on p. 47).
- Arango-Muñoz, S. (2013). “Scaffolded Memory and Metacognitive Feelings”. *Review of Philosophy and Psychology* 4.1, pp. 135–152 (cit. on p. 128).
- (2014). “The Nature of Epistemic Feelings”. *Philosophical Psychology* 27.2, pp. 193–211 (cit. on pp. 121, 130, 150, 175).
- Arcangeli, M. (2017). “Interacting with Emotions: Imagination and Supposition”. *Philosophical Quarterly* 67.269, pp. 730–750 (cit. on p. 210).
- Arguel, A., Lockyer, L., Kennedy, G., Lodge, J. M., and Pachman, M. (2019). “Seeking Optimal Confusion: A Review on Epistemic Emotion Management in Interactive Dig-

- ital Learning Environments”. *Interactive Learning Environments* 27.2, pp. 200–210 (cit. on p. 174).
- Arkes, H. R., Boehm, L. E., and Xu, G. (1991). “Determinants of Judged Validity”. *Journal of Experimental Social Psychology* 27.6, pp. 576–605 (cit. on p. 197).
- Arkes, H. R., Hackett, C., and Boehm, L. (1989). “The Generality of the Relation between Familiarity and Judged Validity”. *Journal of Behavioral Decision Making* 2.2, pp. 81–94 (cit. on p. 197).
- Arnold, M. B. (1960). *Emotion and Personality*. Columbia University Press (cit. on p. 89).
- Audi, R. (1973). “Intending”. *Journal of Philosophy* 70.13, pp. 387–403 (cit. on p. 57).
- (1994). “Dispositional Beliefs and Dispositions to Believe”. *Noûs* 28.4, p. 419 (cit. on pp. 60, 243).
- Auvray, M., Myin, E., and Spence, C. (2010). “The Sensory-Discriminative and Affective-Motivational Aspects of Pain”. *Neuroscience and Biobehavioral Reviews*, p. 11 (cit. on p. 85).
- Ayer, A. J. (1956). *The Problem of Knowledge*. Harmondsworth (cit. on p. 49).
- Bach, D. R. and Dolan, R. J. (2012). “Knowing How Much You Don’t Know: A Neural Organization of Uncertainty Estimates”. *Nature Reviews Neuroscience* 13.8, pp. 572–586 (cit. on p. 128).
- Bacon, F. (1979). “Credibility of Repeated Statements: Memory for Trivia”. *Journal of Experimental Psychology: Human Learning and Memory* 5.3, pp. 241–252 (cit. on pp. 197, 198).
- Bain, D. (2013). “What Makes Pains Unpleasant?” *Philosophical Studies* 166.1, pp. 69–89 (cit. on pp. 63, 91, 92, 108, 110, 114).
- (2014). “Pains That Don’t Hurt”. *Australasian Journal of Philosophy* 92.2, pp. 305–320 (cit. on p. 85).
- (2019). “Why Take Painkillers?” *Noûs* 53.2, pp. 462–490 (cit. on p. 92).
- Balas, R., Sweklej, J., Pochwatko, G., and Godlewska, M. (2012). “On the Influence of Affective States on Intuitive Coherence Judgements”. *Cognition & emotion* 26.2, pp. 312–320 (cit. on pp. 154, 159, 205).
- Bar-Anan, Y., Wilson, T. D., and Gilbert, D. T. (2009). “The Feeling of Uncertainty Intensifies Affective Reactions.” *Emotion* 9.1, pp. 123–127 (cit. on p. 126).
- Barrett, L. F. and Bliss-Moreau, E. (2009). “Affect as a Psychological Primitive”. *Advances in Experimental Social Psychology*. Vol. 41, pp. 167–218 (cit. on pp. 94, 95).
- Barrett, L. F. and Simmons, W. K. (2015). “Interoceptive Predictions in the Brain”. *Nature Reviews Neuroscience* 16.7, pp. 419–429 (cit. on p. 165).

- Battaly, H. (2008). “Virtue Epistemology”. *Philosophy Compass* 3.4, pp. 639–663 (cit. on p. 255).
- Baudouin, J.-Y., Gilibert, D., Sansone, S., and Tiberghien, G. (2000). “When the Smile Is a Cue to Familiarity”. *Memory* 8.5, pp. 285–292 (cit. on p. 156).
- Baumann, N. and Kuhl, J. (2002). “Intuition, Affect, and Personality: Unconscious Coherence Judgments and Self-Regulation of Negative Affect.” *Journal of Personality and Social Psychology* 83.5, p. 1213 (cit. on p. 154).
- Bayne, T. (2004). “Self-Consciousness and the Unity of Consciousness.” *Monist* 87.2, pp. 219–236 (cit. on p. 91).
- Bayne, T. and Montague, M. (2011). *Cognitive Phenomenology*. Oxford University Press on Demand (cit. on p. 55).
- Bayne, T. and Pacherie, E. (2004). “Bottom-Up or Top-Down: Campbell’s Rationalist Account of Monothematic Delusions”. *Philosophy, Psychiatry, & Psychology* 11.1, pp. 1–11 (cit. on p. 152).
- (2005). “In Defence of the Doxastic Conception of Delusions”. *Mind and Language* 20.2, pp. 163–188 (cit. on p. 212).
- Bayne, T. J. and Chalmers, D. J. (2003). “What Is the Unity of Consciousness?” *The Unity of Consciousness* (cit. on pp. 31, 52).
- Bealer, G. (1992). “The Incoherence of Empiricism”. *Aristotelian Society Supplementary Volume* 66, pp. 99–138 (cit. on pp. 24, 28, 36, 48, 67, 68).
- (1996). “On the Possibility of Philosophical Knowledge”. *Noûs* 30, p. 1 (cit. on pp. 37, 254).
- (1999). “A Theory of the a Priori”. *Noûs* 33, pp. 29–55 (cit. on pp. 51, 69, 70, 242).
- Begg, I. M., Anas, A., and Farinacci, S. (1992). “Dissociation of Processes in Belief: Source Recollection, Statement Familiarity, and the Illusion of Truth.” *Journal of Experimental Psychology: General* 121.4, p. 446 (cit. on pp. 197, 199).
- Begg, I., Armour, V., and Kerr, T. (1985). “On Believing What We Remember.” *Canadian Journal of Behavioural Science / Revue canadienne des sciences du comportement* 17.3, pp. 199–214 (cit. on p. 198).
- Belloch, A., Fornés, G., Carrasco, A., López-Solá, C., Alonso, P., and Menchón, J. M. (2016). “Incompleteness and Not Just Right Experiences in the Explanation of Obsessive–Compulsive Disorder”. *Psychiatry Research* 236, pp. 1–8 (cit. on p. 187).
- Benacerraf, P. (1981). “Frege: The Last Logician”. *Midwest Studies in Philosophy* 6.1, pp. 17–36 (cit. on p. 242).
- Bengson, J. (2013). “Experimental Attacks on Intuitions and Answers”. *Philosophy and Phenomenological Research* 86.3, pp. 495–532 (cit. on p. 255).

- Bengson, J. (2014). “How Philosophers Use Intuition and ‘Intuition’”. *Philosophical Studies* 171.3, pp. 555–576 (cit. on pp. 24, 47, 72).
- (2015). “The Intellectual Given”. *Mind* 124.495, pp. 707–760 (cit. on pp. 30, 31, 40, 43, 48, 51, 58, 59, 63, 64, 67, 70–72, 114, 213, 253).
- Ben-Sasson, A., Dickstein, N., Lazarovich, L., and Ayalon, N. (2017). “Not Just Right Experiences: Association with Obsessive Compulsive Symptoms and Sensory Over-Responsivity”. *Occupational Therapy in Mental Health* 33.3, pp. 217–234 (cit. on p. 186).
- Berridge, K. C. and Kringelbach, M. L. (2008). “Affective Neuroscience of Pleasure: Reward in Humans and Animals”. *Psychopharmacology* 199.3, pp. 457–480 (cit. on p. 93).
- Berthier, M., Starkstein, S., and Leiguarda, R. (1988). “Asymbolia for Pain: A Sensory-Limbic Disconnection Syndrome”. *Annals of Neurology* 24.1, pp. 41–49 (cit. on p. 85).
- Betsch, C. (2008a). “Chronic Preferences for Intuition and Deliberation in Decision Making: Lessons Learned about Intuition from an Individual Differences Approach”. *Intuition in Judgment and Decision Making*, pp. 231–248 (cit. on p. 24).
- Betsch, T. (2008b). “The Nature of Intuition and Its Neglect in Research on Judgment and Decision Making”. *Intuition in judgment and decision making*, pp. 3–22 (cit. on p. 49).
- Betsch, T. and Glöckner, A. (2010). “Intuition in Judgment and Decision Making: Extensive Thinking Without Effort”. *Psychological Inquiry* 21.4, pp. 279–294 (cit. on p. 232).
- Bisson, T. (1991). “They’re Made out of Meat”. *OMNI* (cit. on p. 16).
- Block, N. (1986). “Advertisement for a Semantics for Psychology”. *Midwest Studies in Philosophy* 10.1, pp. 615–678 (cit. on pp. 30, 217).
- (1995). “On a Confusion About a Function of Consciousness”. *Brain and Behavioral Sciences* 18.2, pp. 227–247 (cit. on p. 52).
- Boghossian, P. (2000). “Knowledge of Logic”. *New Essays on the A Priori* (cit. on pp. 25, 67, 221, 242).
- (2003). “Blind Reasoning”. *Aristotelian Society Supplementary Volume*. Vol. 77, pp. 225–248 (cit. on p. 193).
- (2009). “Virtuous Intuitions: Comments on Lecture 3 of Ernest Sosa’s A Virtue Epistemology”. *Philosophical Studies* 144.1, pp. 111–119 (cit. on p. 48).
- Boghossian, P. and Bonjour, L. (2001). “Inference and Insight” (cit. on pp. 26, 67, 221, 242).

- Bolte, A. and Goschke, T. (2005). “On the Speed of Intuition: Intuitive Judgments of Semantic Coherence under Different Response Deadlines”. *Memory & cognition* 33.7, pp. 1248–1255 (cit. on p. 153).
- (2008). “Intuition in the Context of Object Perception: Intuitive Gestalt Judgments Rest on the Unconscious Activation of Semantic Representations”. *Cognition* 108.3, pp. 608–616 (cit. on p. 153).
- Bolte, A., Goschke, T., and Kuhl, J. (2003). “Emotion and Intuition: Effects of Positive and Negative Mood on Implicit Judgments of Semantic Coherence”. *Psychological science* 14.5, pp. 416–421 (cit. on p. 154).
- BonJour, L. (1998). *In Defense of Pure Reason: A Rationalist Account of a Priori Justification*. Cambridge University Press (cit. on p. 242).
- Bonjour, L. (2001). “Replies”. *Philosophy and Phenomenological Research* 63.3, pp. 673–698 (cit. on p. 30).
- Bornstein, R. F. (1989). “Exposure and Affect: Overview and Meta-Analysis of Research, 1968–1987.” *Psychological bulletin* 106.2, p. 265 (cit. on p. 207).
- Bornstein, R. F. and D’Agostino, P. R. (1994). “The Attribution and Discounting of Perceptual Fluency: Preliminary Tests of a Perceptual Fluency/Attributional Model of the Mere Exposure Effect”. *Social Cognition* 12.2, pp. 103–128 (cit. on p. 207).
- Bourget, D. (2017). “Why Are Some Phenomenal Experiences ‘Vivid’ and Others ‘Faint’? Representationalism, Imagery, and Cognitive Phenomenology”. *Australasian Journal of Philosophy* 95.4, pp. 673–687 (cit. on pp. 40, 105, 176).
- Bouzerda-Wahlen, A., Nahum, L., Ptak, R., and Schnider, A. (2013). “Mechanism of Disorientation: Reality Filtering versus Content Monitoring”. *Cortex* 49.10, pp. 2628–2636 (cit. on p. 243).
- Bowden, E., Jungbeeman, M., Fleck, J., and Kounios, J. (2005). “New Approaches to Demystifying Insight”. *Trends in Cognitive Sciences* 9.7, pp. 322–328 (cit. on p. 174).
- Bowers, K. S., Regehr, G., Balthazard, C., and Parker, K. (1990). “Intuition in the Context of Discovery”. *Cognitive Psychology* 22.1, pp. 72–110 (cit. on p. 153).
- Brady, M. (2009). “The Irrationality of Recalcitrant Emotions”. *Philosophical Studies* 145.3, pp. 413–430 (cit. on pp. 97, 116, 119, 171).
- (2013). *Emotional Insight: The Epistemic Role of Emotional Experience*. Oxford University Press (cit. on pp. 97, 251, 252).
- Bramble, B. (2013). “The Distinctive Feeling Theory of Pleasure”. *Philosophical Studies* 162.2, pp. 201–217 (cit. on p. 172).

- Bransford, J. D. and Johnson, M. K. (1972). “Contextual Prerequisites for Understanding: Some Investigations of Comprehension and Recall”. *Journal of Verbal Learning and Verbal Behavior* 11.6, pp. 717–726 (cit. on p. 174).
- Breen, N., Caine, D., and Coltheart, M. (2000). “Models of Face Recognition and Delusional Misidentification: A Critical Review”. *Cognitive Neuropsychology* 17.1-3, pp. 55–71 (cit. on p. 152).
- Briesen, J. (2015). “Perceptual Justification and Assertively Representing the World”. *Philosophical Studies* 172.8, pp. 2239–2259 (cit. on pp. 33, 34).
- Broad, C. D. (1954). “Emotion and Sentiment”. *The Journal of Aesthetics and Art Criticism* 13.2, pp. 203–214 (cit. on p. 89).
- Brogaard, B. (2013). “Phenomenal Seemings and Sensible Dogmatism”. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, pp. 270–289 (cit. on pp. 51, 69, 242).
- (2014). “Intuitions as Intellectual Seemings”. *Analytic Philosophy* 55.4, pp. 382–393 (cit. on p. 67).
- Brown, A. S. (2003). “A Review of the Déjà vu Experience.” *Psychological Bulletin* 129.3, pp. 394–413 (cit. on p. 124).
- Brown, A. S. and Nix, L. A. (1996). “Turning Lies into Truths: Referential Validation of Falsehoods.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22.5, p. 1088 (cit. on pp. 197, 199).
- Brown, R. and McNeill, D. (1966). “The “Tip of the Tongue” Phenomenon”. *Journal of verbal learning and verbal behavior* 5.4, pp. 325–337 (cit. on pp. 127, 131).
- Brun, G., Dogluoglu, U., and Kuenzle, D. (2008). *Epistemology and Emotions*. Ashgate Publishing Company (cit. on p. 84).
- Byerly, T. R. (2012). “It Seems Like There Aren’t Any Seemings”. *Philosophia* 40.4, pp. 771–782 (cit. on pp. 250, 251).
- Cacioppo, J. T., Gardner, W. L., and Berntson, G. G. (1999). “The Affect System Has Parallel and Integrative Processing Components: Form Follows Function”. *Journal of Personality and Social Psychology* 76.5, p. 839 (cit. on p. 94).
- Candiotto, L. (2017). “Epistemic Emotions. The Building Blocks of Intellectual Virtues”. *Studi di estetica* 4.45, pp. 7–25 (cit. on p. 255).
- Cappelen, H. (2012). *Philosophy Without Intuitions*. OUP (cit. on pp. 24, 47, 49, 67).
- Carruthers, P. (2009). “How We Know Our Own Minds: The Relationship between Mindreading and Metacognition”. *Behavioral and Brain Sciences* 32.02, p. 121 (cit. on pp. 136, 145).

- Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford University Press (cit. on pp. 94, 96, 112, 165, 167).
- (2017a). “Are Epistemic Emotions Metacognitive?” *Philosophical Psychology* 30.1-2, pp. 58–78 (cit. on pp. 146, 171, 190).
 - (2017b). “Valence and Value”. *Philosophy and Phenomenological Research* 97.3, pp. 658–680 (cit. on pp. 90, 96, 97, 107, 121, 150, 165, 168, 239, 256).
 - (2018). “Basic Questions”. *Mind & Language* 33.2, pp. 130–147 (cit. on pp. 92, 128, 175, 239).
- Chalmers, D. J. (2002). “Does Conceivability Entail Possibility”. *Conceivability and Possibility*, pp. 145–200 (cit. on p. 82).
- (2004). “The Representational Character of Experience”. *The Future for Philosophy*, pp. 153–181 (cit. on p. 31).
 - (2006). “Perception and the Fall From Eden”. *Perceptual Experience*, pp. 49–125 (cit. on p. 57).
- Charland, L. (2005). “The Heat of Emotion: Valence and the Demarcation Problem”. *Journal of Consciousness Studies* 12.8-10, pp. 8–10 (cit. on pp. 90, 91).
- Chetverikov, A. (2014). “Warmth of Familiarity and Chill of Error: Affective Consequences of Recognition Decisions”. *Cognition and Emotion* 28.3, pp. 385–415 (cit. on p. 159).
- Childress, A. R., Ehrman, R. N., Wang, Z., Li, Y., Sciortino, N., Hakun, J., Jens, W., Suh, J., Listerud, J., Marquez, K., Franklin, T., Langleben, D., Detre, J., and O’Brien, C. P. (2008). “Prelude to Passion: Limbic Activation by “Unseen” Drug and Sexual Cues”. *PLoS ONE* 3.1, e1506 (cit. on p. 92).
- Chudnoff, E. (2011a). “The Nature of Intuitive Justification”. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 153.2, pp. 313–333 (cit. on p. 72).
- (2011b). “What Intuitions Are Like”. *Philosophy and Phenomenological Research* 82.3, pp. 625–654 (cit. on pp. 30, 34, 40, 47, 48, 51, 67, 70–72, 111, 171).
 - (2013a). “Awareness of Abstract Objects”. *Noûs* 47.4, pp. 706–726 (cit. on p. 30).
 - (2013b). “Intellectual Gestalts”. *Phenomenal Intentionality*, p. 174 (cit. on p. 31).
 - (2013c). *Intuition*. OUP (cit. on pp. 48, 70, 210, 253).
 - (2014a). “Is Intuition Based On Understanding?” *Philosophy and Phenomenological Research* 89.1, pp. 42–67 (cit. on pp. 59, 242).
 - (2014b). “The Rational Roles of Intuition”. *Intuitions*, pp. 9–31 (cit. on pp. 35, 43, 72, 78, 193).

- Clark, A. (2009). “Perception, Action, and Experience: Unraveling the Golden Braid”. *Neuropsychologia* 47.6, pp. 1460–1468 (cit. on p. 31).
- Clavien, C. and FitzGerald, C. (2018). “The Evolution of Moral Intuitions and Their Feeling of Rightness”. *The Routledge Handbook of Evolution and Philosophy* (cit. on pp. 84, 122).
- Cleary, A. M., Konkel, K. E., Nomi, J. S., and McCabe, D. P. (2010). “Odor Recognition without Identification”. *Memory & cognition* 38.4, pp. 452–460 (cit. on p. 128).
- Climenhaga, N. (2018). “Intuitions Are Used as Evidence in Philosophy”. *Mind* 127.505, pp. 69–104 (cit. on p. 47).
- Clore, G. L. (1992). “Cognitive Phenomenology: Feelings and the Construction of Judgment”. *The Construction of Social Judgments*, pp. 133–163 (cit. on p. 121).
- Clore, G. L. and Gasper, K. (2000). “Feeling Is Believing: Some Affective Influences on Belief”. *Emotions and Beliefs: How Feelings Influence Thoughts*, pp. 10–44 (cit. on p. 254).
- Clore, G. L. and Huntsinger, J. R. (2007). “How Emotions Inform Judgment and Regulate Thought”. *Trends in Cognitive Sciences* 11.9, pp. 393–399 (cit. on pp. 116, 118).
- Cohen, J. (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford University Press (cit. on pp. 147, 164).
- Cohen, J. and Fulkerson, M. (2014). “Affect, Rationalization, and Motivation”. *Review of Philosophy and Psychology* 5.1, pp. 103–118 (cit. on p. 97).
- Cohen, L. J. (1992). *An Essay on Belief and Acceptance*. New York: Clarendon Press (cit. on pp. 61, 243).
- Colombetti, G. (2005). “Appraising Valence”. *Journal of Consciousness Studies* 12.8-10, pp. 8–10 (cit. on p. 91).
- (2011). “Varieties of Pre-Reflective Self-Awareness: Foreground and Background Bodily Feelings in Emotion Experience”. *Inquiry* 54.3, pp. 293–313 (cit. on p. 169).
- Colombetti, G. and Harrison, N. (2018). “From Physiology to Experience: Enriching Existing Conceptions of “Arousal” in Affective Science”. *The Interoceptive Mind: From Homeostasis to Awareness*, pp. 245–258 (cit. on p. 94).
- Corder, G., Ahanonu, B., Grewe, B. F., Wang, D., Schnitzer, M. J., and Scherrer, G. (2019). “An Amygdalar Neural Ensemble That Encodes the Unpleasantness of Pain”, p. 7 (cit. on p. 85).
- Corns, J. (2014). “Unpleasantness, Motivational Oomph, and Painfulness”. *Mind & Language* 29.2, pp. 238–254 (cit. on p. 97).

- Cogle, J. R. and Lee, H.-J. (2014). “Pathological and Non-Pathological Features of Obsessive-Compulsive Disorder: Revisiting Basic Assumptions of Cognitive Models”. *Journal of Obsessive-Compulsive and Related Disorders* 3.1, pp. 12–20 (cit. on p. 187).
- Coutanche, M. N. and Thompson-Schill, S. L. (2012). “Reversal without Remapping: What We Can (and Cannot) Conclude about Learned Associations from Training-Induced Behavior Changes”. *Perspectives on Psychological Science* 7.2, pp. 118–134 (cit. on p. 236).
- Cowan, R. (2015). “Perceptual Intuitionism”. *Philosophy and Phenomenological Research* 90.1, pp. 164–193 (cit. on p. 99).
- Craig, A. D. (2009). “How Do You Feel — Now? The Anterior Insula and Human Awareness”. *Nature Reviews Neuroscience* 10.1, pp. 59–70 (cit. on pp. 90, 152).
- Crane, T. (2001). *Elements of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press (cit. on pp. 52, 53, 56).
- (2009). “Intentionalism”. *The Oxford Handbook of Philosophy of Mind* (cit. on pp. 29, 31, 34, 98, 104).
- (2013). “Unconscious Belief and Conscious Thought”. *Phenomenal Intentionality*, p. 156 (cit. on p. 52).
- Currie, G. and Ravenscroft, I. (2002). *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford University Press (cit. on p. 210).
- Damasio, A. R. (1994). *Descartes’ Error: Emotion, Rationality and the Human Brain* (cit. on pp. 85, 98).
- Damasio, A. and Carvalho, G. B. (2013). “The Nature of Feelings: Evolutionary and Neurobiological Origins”. *Nature Reviews Neuroscience* 14.2, pp. 143–152 (cit. on pp. 90, 152).
- Dancy, J. (2014). “Intuition and Emotion”. *Ethics* 124.4, pp. 787–812 (cit. on p. 84).
- D’Arms, J. and Jacobson, D. (2003). “The Significance of Recalcitrant Emotion (or, Anti-Quasijudgmentalism)”. *Royal Institute of Philosophy Supplement* 52, pp. 127–145 (cit. on p. 251).
- Davidson, D. (1987). “Knowing One’s Own Mind”. *Proceedings and Addresses of the American Philosophical Association* 60.3, pp. 441–458 (cit. on p. 16).
- Davies, M., Coltheart, M., Langdon, R., and Breen, N. (2001). “Monothematic Delusions: Towards a Two-Factor Account”. *Philosophy, Psychiatry, and Psychology* 8.2-3, pp. 133–58 (cit. on p. 125).
- Dawson, M. E., Filion, D. L., and Schell, A. M. (1989). “Is Elicitation of the Autonomic Orienting Response Associated with Allocation of Processing Resources?” *Psychophysiology* 26.5, pp. 560–572 (cit. on p. 95).

- De keersmaecker, J., Dunning, D., Pennycook, G., Rand, D. G., Sanchez, C., Unkelbach, C., and Roets, A. (2019). “Investigating the Robustness of the Illusory Truth Effect Across Individual Differences in Cognitive Ability, Need for Cognitive Closure, and Cognitive Style”. *Personality and Social Psychology Bulletin*, p. 0146167219853844 (cit. on p. 198).
- De Neys, W. (2012). “Bias and Conflict: A Case for Logical Intuitions”. *Perspectives on Psychological Science* 7.1, pp. 28–38 (cit. on p. 84).
- (2017). *Dual Process Theory 2.0*. Routledge (cit. on p. 24).
- De Roeck, A., Johnson, R., King, M., Rosner, M., Sampson, G., and Varile, N. (1982). “A Myth about Centre-Embedding”. *Lingua* 58.3, pp. 327–340 (cit. on p. 39).
- De Sousa, R. (1987). *The Rationality of Emotion*. MIT Press (cit. on pp. 97, 102).
- (2008). “Epistemic Feelings”. *Epistemology and Emotions*, pp. 185–203 (cit. on pp. 20, 121, 125, 147, 261).
- De Vignemont, F. (2018). “Was Descartes Right after All? An Affective Background for Bodily Awareness”. *The Interoceptive Mind: From Homeostasis to Awareness* (cit. on pp. 86, 92, 162, 163).
- (2019). “Against Phenomenal Parsimony”. *Metaphysics and Cognitive Science*, pp. 268–284 (cit. on p. 25).
- Dechêne, A., Stahl, C., Hansen, J., and Wänke, M. (2009). “Mix Me a List: Context Moderates the Truth Effect and the Mere-Exposure Effect”. *Journal of Experimental Social Psychology* 45.5, pp. 1117–1122 (cit. on pp. 44, 198, 202).
- (2010). “The Truth About the Truth: A Meta-Analytic Review of the Truth Effect”. *Personality and Social Psychology Review* 14.2, pp. 238–257 (cit. on pp. 197, 198).
- Deigh, J. (1994). “Cognitivism in the Theory of Emotions”. *Ethics* 104.4, pp. 824–854 (cit. on p. 251).
- DeLong, K. A., Urbach, T. P., and Kutas, M. (2005). “Probabilistic Word Pre-Activation during Language Comprehension Inferred from Electrical Brain Activity”. *Nature Neuroscience* 8.8, pp. 1117–1121 (cit. on p. 197).
- Deonna, J. A. (2006). “Emotion, Perception and Perspective”. *dialectica* 60.1, pp. 29–46 (cit. on pp. 15, 147, 251).
- Deonna, J. A. and Teroni, F. (2012). *The Emotions: A Philosophical Introduction*. Routledge (cit. on pp. 104, 108, 109, 252).
- (2015). “Emotions as Attitudes: Emotions as Attitudes”. *Dialectica* 69.3, pp. 293–311 (cit. on pp. 102, 104, 108, 218, 252).
- DePaul, M. R. and Ramsey, W. (1998). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Rowman & Littlefield Publishers (cit. on p. 47).

- Descartes, R. (1649). *Passions of the Soul* (cit. on p. 86).
- Deutsch, M. (2015). *The Myth of the Intuitive*. The MIT Press (cit. on p. 47).
- Dewey, J. (1929). *Nature and Experience*. George Allen & Unwin (cit. on p. 122).
- Dijksterhuis, A. (2006). “On Making the Right Choice: The Deliberation-Without-Attention Effect”. *Science* 311.5763, pp. 1005–1007 (cit. on p. 196).
- Dijksterhuis, A. and Nordgren, L. F. (2006). “A Theory of Unconscious Thought”. *Perspectives on Psychological Science* 1.2, pp. 95–109 (cit. on p. 196).
- Dijksterhuis, A. and Van Olden, Z. (2006). “On the Benefits of Thinking Unconsciously: Unconscious Thought Can Increase Post-Choice Satisfaction”. *Journal of Experimental Social Psychology* 42.5, pp. 627–631 (cit. on p. 196).
- Dodd, J. (2014). “Realism and Anti-Realism about Experiences of Understanding”. *Philosophical Studies* 168.3, pp. 745–767 (cit. on pp. 128, 174).
- Dokic, J. (2012). “Seeds of Self-Knowledge: Noetic Feelings and Metacognition”. *The Foundations of Metacognition*, pp. 302–321 (cit. on pp. 121, 126, 130, 150, 164, 177).
- (2014a). “Feeling the Past: A Two-Tiered Account of Episodic Memory”. *Review of Philosophy and Psychology* 5.3, pp. 413–426 (cit. on p. 143).
- (2014b). “Feelings of (Un)Certainty and Margins for Error”. *Philosophical Inquiries* 2.1, pp. 123–144 (cit. on p. 128).
- (2016). “Aesthetic Experience as a Metacognitive Feeling? A Dual-Aspect View”. *Proceedings of the Aristotelian Society* 116.1, pp. 69–88 (cit. on pp. 122, 169).
- Dokic, J. and Lemaire, S. (2013). “Are Emotions Perceptions of Value?” *Canadian Journal of Philosophy* 43.2, pp. 227–247 (cit. on pp. 99, 251).
- Dokic, J. and Martin, J.-R. (2015a). “Felt Reality and the Opacity of Perception”. *Topoi* 36.2, pp. 299–309 (cit. on pp. 117, 243).
- (2015b). “Looks the Same but Feels Different’: A Metacognitive Approach to Cognitive Penetrability”. *Cognitive Effects of Perception: New Perspectives* (cit. on p. 250).
- Döring, S. A. (2004). “Gründe Und Gefühle. Rationale Motivation Durch Emotionale Vernunft”. Habilitationsschrift. Universität Essen (cit. on p. 251).
- (2007). “Seeing What to Do: Affective Perception and Rational Motivation”. *dialectica* 61.3, pp. 363–394 (cit. on p. 251).
- (2009). “Why Be Emotional?” *The Oxford Handbook of Philosophy of Emotion* (cit. on p. 36).
- Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press (cit. on pp. 29, 30, 227).
- (1986). “Misrepresentation”. *Belief: Form, Content, and Function*, pp. 17–36 (cit. on p. 29).

- Dub, R. (2015). “Delusions, Acceptances, and Cognitive Feelings”. *Philosophy and Phenomenological Research* 94.1, pp. 27–60 (cit. on p. 122).
- Duke, D., Fiacconi, C. M., and Köhler, S. (2014). “Parallel Effects of Processing Fluency and Positive Affect on Familiarity-Based Recognition Decisions for Faces”. *Frontiers in Psychology* 5 (cit. on pp. 156–159).
- Earlenbaugh, J. and Molyneux, B. (2009). “Intuitions Are Inclinations to Believe”. *Philosophical Studies* 145.1, pp. 89–109 (cit. on pp. 47, 48).
- Echeverri, S. (2017). “Emotional Justification”. *Philosophy and Phenomenological Research* (cit. on pp. 109, 111).
- Efklides, A. (2014). “The Blank-in-the-Mind Experience: Another Manifestation of the Tip-of-the-Tongue State or Something Else”. *Tip-of-the-Tongue States and Related Phenomena*, pp. 232–263 (cit. on p. 139).
- Egan, A. (2011). “Comments on Gendler’s, “the Epistemic Costs of Implicit Bias””. *Philosophical Studies* 156.1, p. 65 (cit. on p. 25).
- Egler, M. (2019). “Testing for the Phenomenal: Intuition, Metacognition, and Philosophical Methodology”. *Mind & Language* (cit. on p. 245).
- Eisenberger, N. I. (2012). “The Pain of Social Disconnection: Examining the Shared Neural Underpinnings of Physical and Social Pain”. *Nature Reviews Neuroscience* 13.6, pp. 421–434 (cit. on p. 102).
- Elga, A. and Rayo, A. (2015). “Fragmentation and Information Access” (cit. on pp. 237, 239, 240).
- Elliot, A. J. and Devine, P. G. (1994). “On the Motivational Nature of Cognitive Dissonance: Dissonance as Psychological Discomfort.” *Journal of personality and social psychology* 67.3, p. 382 (cit. on p. 229).
- Ellis, H. D., Young, A. W., Quayle, A. H., and De Pauw, K. W. (1997). “Reduced Autonomic Responses to Faces in Capgras Delusion”. *Proceedings of the Royal Society B: Biological Sciences* 264.1384, pp. 1085–1092 (cit. on p. 152).
- Ellis, H. D. and Lewis, M. B. (2001). “Capgras Delusion: A Window on Face Recognition”. *Trends in cognitive sciences* 5.4, pp. 149–156 (cit. on p. 152).
- Epstein, S. (2008). “Intuition from the Perspective of Cognitive-Experiential Self-Theory”. *Intuition in Judgment and Decision Making*, pp. 23–37 (cit. on pp. 24, 49).
- Erle, T. M., Reber, R., and Topolinski, S. (2017). “Affect from Mere Perception: Illusory Contour Perception Feels Good”. *Emotion* 17.5, pp. 856–866 (cit. on p. 157).
- Erle, T. M. and Topolinski, S. (2018). “Disillusionment”. *Experimental Psychology* 65.6, pp. 332–344 (cit. on p. 157).
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press (cit. on p. 192).

- Evans, J. S. B. T. (2008). “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition”. *Annual Review of Psychology* 59.1, pp. 255–278 (cit. on p. 24).
- Evans, J. S. B. T. and Stanovich, K. E. (2013). “Dual-Process Theories of Higher Cognition Advancing the Debate”. *Perspectives on Psychological Science* 8.3, pp. 223–241 (cit. on pp. 24, 84).
- Fairweather, A. and Zagzebski, L. (2001). *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford University Press on Demand (cit. on p. 255).
- Fang, X., Singh, S., and Ahluwalia, R. (2007). “An Examination of Different Explanations for the Mere Exposure Effect”. *Journal of Consumer Research* 34.1, pp. 97–103 (cit. on p. 207).
- Faramarzi, S. (2018). *The Odd Psychology behind Oddly Satisfying Slime Videos*. URL: <https://www.wired.co.uk/article/oddly-satisfying-videos-explained-psychology-youtube> (visited on 08/22/2019) (cit. on p. 183).
- Fazio, L. K., Brashier, N. M., Payne, B. K., and Marsh, E. J. (2015). “Knowledge Does Not Protect against Illusory Truth.” *Journal of Experimental Psychology: General* 144.5, pp. 993–1002 (cit. on pp. 199, 212, 237).
- Fazio, L. K., Rand, D. G., and Pennycook, G. (2019). “Repetition Increases Perceived Truth Equally for Plausible and Implausible Statements”. *Psychonomic Bulletin & Review* (cit. on pp. 200, 201, 242).
- Fernandez Cruz, A. L., Arango-Muñoz, S., and Volz, K. G. (2016). “Oops, Scratch That! Monitoring One’s Own Errors during Mental Calculation”. *Cognition* 146, pp. 110–120 (cit. on pp. 183, 190, 191).
- Festinger, L. (1962). *A Theory of Cognitive Dissonance*. Stanford university press (cit. on p. 229).
- Fiacconi, C. M., Kouptsova, J. E., and Köhler, S. (2017). “A Role for Visceral Feedback and Interoception in Feelings-of-Knowing”. *Consciousness and Cognition* 53, pp. 70–80 (cit. on p. 152).
- Fiacconi, C. M., Peter, E. L., Owais, S., and Köhler, S. (2016). “Knowing by Heart: Visceral Feedback Shapes Recognition Memory Judgments.” *Journal of Experimental Psychology: General* 145.5, pp. 559–572 (cit. on p. 152).
- Filion, D. L., Dawson, M. E., Schell, A. M., and Hazlett, E. A. (1991). “The Relationship between Skin Conductance Orienting the Allocation of Processing Resources”. *Psychophysiology* 28.4, pp. 410–424 (cit. on p. 95).
- Flavell, J. H. (1979). “Metacognition and Cognitive Monitoring: A New Area of Cognitive–Developmental Inquiry.” *American psychologist* 34.10, p. 906 (cit. on p. 145).

- Flavell, J. C., Tipper, S. P., and Over, H. (2018). “Preference for Illusory Contours: Beyond Object Symmetry, Familiarity, and Nameability”. *Emotion* 18.5, pp. 736–738 (cit. on p. 157).
- Forster, M., Leder, H., and Ansorge, U. (2013). “It Felt Fluent, and I Liked It: Subjective Feeling of Fluency Rather than Objective Fluency Determines Liking.” *Emotion* 13.2, pp. 280–289 (cit. on pp. 135, 153, 157, 226).
- (2016). “Exploring the Subjective Feeling of Fluency”. *Experimental Psychology* 63.1, pp. 45–58 (cit. on p. 152).
- Friedman, J. (2013). “Question-Directed Attitudes”. *Philosophical Perspectives* 27.1, pp. 145–174 (cit. on p. 175).
- Frijda, N. H. (1986). *The Emotions*. Cambridge University Press (cit. on pp. 89, 162, 173).
- (2008). “The Psychologists’ Point of View”. *Handbook of Emotions*. Vol. 3rd, pp. 68–87 (cit. on pp. 90, 116, 162, 177).
- (2009). “Emotion Experience and Its Varieties”. *Emotion Review* 1.3, pp. 264–271 (cit. on p. 162).
- Fulkerson, M. (2019). “Emotional Perception”. *Australasian Journal of Philosophy*, pp. 1–15 (cit. on pp. 90, 111, 113, 162–165).
- Fumerton, R. A. (1990). *Reason and Morality: A Defense of the Egocentric Perspective*. Cornell University Press (cit. on p. 49).
- Gangemi, A., Bourgeois-Gironde, S., and Mancini, F. (2015). “Feelings of Error in Reasoning—in Search of a Phenomenon”. *Thinking & Reasoning* 21.4, pp. 383–396 (cit. on pp. 183, 191).
- García-Carpintero, M. (2013). “Norms of Fiction-Making”. *British Journal of Aesthetics* 53.3, pp. 339–357 (cit. on p. 210).
- Garcia-Marques, T., Mackie, D. M., Claypool, H. M., and Garcia-Marques, L. (2004). “Positivity Can Cue Familiarity”. *Personality and Social Psychology Bulletin* 30.5, pp. 585–593 (cit. on pp. 156, 205).
- Garcia-Marques, T., Silva, R. R., and Mello, J. (2016). “Judging the Truth-Value of a Statement in and out of a Deep Processing Context”. *social cognition* 34.1, pp. 40–54 (cit. on p. 198).
- Garcia-Marques, T., Silva, R. R., Mello, J., and Hansen, J. (2019). “Relative to What? Dynamic Updating of Fluency Standards and between-Participants Illusions of Truth”. *Acta Psychologica* 195, pp. 71–79 (cit. on pp. 142, 198, 202).

- Garcia-Marques, T., Silva, R. R., Reber, R., and Unkelbach, C. (2015). “Hearing a Statement Now and Believing the Opposite Later”. *Journal of Experimental Social Psychology* 56, pp. 126–129 (cit. on pp. 198, 237).
- Gawronski, B. and Strack, F. (2012). *Cognitive Consistency: A Fundamental Principle in Social Cognition*. Guilford press (cit. on p. 229).
- Geach, P. (1957). *Mental Acts*. Routledge and Kegan Paul (cit. on p. 60).
- Gendler, T. S. (2006). “Imaginative Contagion”. *Metaphilosophy* 37.2, pp. 183–203 (cit. on p. 167).
- Gerrans, P. (2015). “All the Self We Need”. *Open MIND* (cit. on pp. 86, 87).
- Gettier, E. (1963). “Is Justified True Belief Knowledge?” *Analysis* 23.6, pp. 121–123 (cit. on pp. 27, 40).
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Taylor Francis Group, LLC (cit. on p. 162).
- Gilbert, D. (1991). “How Mental Systems Believe”. *American Psychologist* 46.2, pp. 107–119 (cit. on p. 235).
- Gilead, M., Sela, M., and Maril, A. (2019). “That’s My Truth: Evidence for Involuntary Opinion Confirmation”. *Social Psychological and Personality Science* 10.3, pp. 393–401 (cit. on p. 240).
- Glöckner, A. and Wittman, C. (2010). “Beyond Dual-Process Models: A Categorisation of Processes Underlying Intuitive Judgement and Decision Making”. *Thinking & Reasoning* 16.1, pp. 1–25 (cit. on pp. 24, 84).
- Glucksberg, S. and McCloskey, M. (1981). “Decisions about Ignorance: Knowing That You Don’t Know.” *Journal of Experimental Psychology: Human Learning and Memory* 7.5, p. 311 (cit. on p. 126).
- Glüer, K. (2009). “In Defence of a Doxastic Account of Experience”. *Mind and Language* 24.3, pp. 297–327 (cit. on pp. 56, 57).
- Goffin, K. (2019). “The Affective Experience of Aesthetic Properties”. *Pacific Philosophical Quarterly* 100.1, pp. 283–300 (cit. on p. 169).
- Goldie, P. (2000). *The Emotions: A Philosophical Exploration*. Oxford University Press (cit. on pp. 88, 98, 102, 105, 165, 252).
- (2002). “Emotions, Feelings and Intentionality”. *Phenomenology and the Cognitive Sciences* 1.3, pp. 235–254 (cit. on pp. 98, 111).
- (2012). “Loss of Affect in Intellectual Activity”. *Emotion Review* 4.2, pp. 122–126 (cit. on pp. 121, 154).
- Goldstein, I. (1981). “Cognitive Pleasure and Distress”. *Philosophical Studies* 39 (January), pp. 15–23 (cit. on p. 102).

- Goldstein, I. (2002). “Are Emotions Feelings? A Further Look at Hedonic Theories of Emotions”. *Consciousness and Emotion* 3.1, pp. 21–33 (cit. on p. 111).
- Gordon, R. M. (1995). “Simulation Without Introspection or Inference From Me to You”. *Mental Simulation* (cit. on p. 192).
- Greene, J. D. (2008). “The Secret Joke of Kant’s Soul”. *Moral psychology* 3, pp. 35–79 (cit. on p. 84).
- Greifeneder, R. and Bless, H. (2008). “Depression and Reliance on Ease-of-Retrieval Experiences”. *European Journal of Social Psychology* 38.2, pp. 213–230 (cit. on p. 154).
- Greifeneder, R., Bless, H., and Pham, M. T. (2011). “When Do People Rely on Affective and Cognitive Feelings in Judgment? A Review”. *Personality and Social Psychology Review* 15.2, pp. 107–141 (cit. on p. 154).
- Grice, H. P. (1975). “Logic and Conversation”. *Syntax and Semantics: Speech Acts*. Vol. 3, pp. 41–58 (cit. on p. 225).
- Griffiths, P. E. and Scarantino, A. (2005). “Emotions in the Wild: The Situated Perspective on Emotion”. *The Cambridge Handbook of Situated Cognition* (cit. on p. 162).
- Grzankowski, A. (2016). “Attitudes Towards Objects: Attitudes Towards Objects”. *Noûs* 50.2, pp. 314–328 (cit. on p. 101).
- Haidt, J. (2001). “The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment.” *Psychological review* 108.4, p. 814 (cit. on p. 215).
- Halamish, V., McGillivray, S., and Castel, A. D. (2011). “Monitoring One’s Own Forgetting in Younger and Older Adults.” *Psychology and Aging* 26.3, pp. 631–635 (cit. on p. 128).
- Halberstadt, J. and Catty, S. (2008). “Analytic Thought Disrupts Familiarity-Based Decision Making”. *Social Cognition* 26.6, pp. 755–765 (cit. on p. 196).
- Hansen, J. and Wänke, M. (2013). “Fluency in Context: Discrepancy Makes Processing Experiences Informative”. *The Experience of Thinking: How the Fluency of Mental Processes Influences Cognition and Behavior*, pp. 70–84 (cit. on pp. 142, 203, 225).
- Hansen, J., Dechêne, A., and Wänke, M. (2008). “Discrepant Fluency Increases Subjective Truth”. *Journal of Experimental Social Psychology* 44.3, pp. 687–691 (cit. on pp. 153, 202).
- Hansen, J. and Wänke, M. (2010). “Truth From Language and Truth From Fit: The Impact of Linguistic Concreteness and Level of Construal on Subjective Truth”. *Personality and Social Psychology Bulletin* 36.11, pp. 1576–1588 (cit. on p. 202).
- Harman, G. (1987). “(Nonsolipsistic) Conceptual Role Semantics”. *New Directions in Semantics*, pp. 55–81 (cit. on p. 30).

- Harman, G. (1990). "The Intrinsic Quality of Experience". *Philosophical perspectives* 4, pp. 31–52 (cit. on p. 40).
- Harmon-Jones, E. and Allen, J. J. B. (2001). "The Role of Affect in the Mere Exposure Effect: Evidence from Psychophysiological and Individual Differences Approaches". *Personality and Social Psychology Bulletin* 27.7, pp. 889–898 (cit. on p. 207).
- Harmon-Jones, E., Amodio, D. M., and Harmon-Jones, C. (2009). "Action-Based Model of Dissonance: A Review, Integration, and Expansion of Conceptions of Cognitive Conflict". *Advances in experimental social psychology* 41, pp. 119–166 (cit. on p. 229).
- Hart, J. T. (1965). "Memory and the Feeling-of-Knowing Experience." *Journal of Educational Psychology* 56.4, pp. 208–216 (cit. on pp. 122, 126, 131).
- Hasher, L., Goldstein, D., and Toppino, T. (1977). "Frequency and the Conference of Referential Validity". *Journal of Verbal Learning and Verbal Behavior* 16.1, pp. 107–112 (cit. on pp. 197, 198).
- Hauser, M., Cushman, F., Young, L., Kang-Xing Jin, R., and Mikhail, J. (2007). "A Dissociation between Moral Judgments and Justifications". *Mind & language* 22.1, pp. 1–21 (cit. on p. 215).
- Haybron, D. M. (2008). *The Pursuit of Unhappiness: The Elusive Psychology of Well-Being*. Oxford University Press on Demand (cit. on p. 171).
- Helm, B. W. (2001). *Emotional Reason: Deliberation, Motivation, and the Nature of Value*. Cambridge University Press (cit. on pp. 102, 114, 115, 117, 163, 252).
- Henkel, L. A. and Mattson, M. E. (2011). "Reading Is Believing: The Truth Effect and Source Credibility". *Consciousness and Cognition* 20.4, pp. 1705–1721 (cit. on p. 199).
- Henrich, J., Heine, S. J., and Norenzayan, A. (2010). "The Weirdest People in the World?" *Behavioral and Brain Sciences* 33.2-3, pp. 61–83 (cit. on p. 38).
- Hertwig, R., Herzog, S. M., Schooler, L. J., and Reimer, T. (2008). "Fluency Heuristic: A Model of How the Mind Exploits a by-Product of Information Retrieval." *Journal of Experimental Psychology: Learning, memory, and cognition* 34.5, p. 1191 (cit. on p. 135).
- Herzberg, L. A. (2012). "To Blend or to Compose: A Debate About Emotion Structure". *Dynamicality in Emotion Concepts* (cit. on pp. 108, 111, 113).
- Herzog, S. M. and Hertwig, R. (2013). "The Ecological Validity of Fluency". *The experience of thinking: How the fluency of mental processes influences cognition and behaviour*, p. 190 (cit. on pp. 199, 223, 225, 238, 244).
- Hicks, J. A., Cicero, D. C., Trent, J., Burton, C. M., and King, L. A. (2010). "Positive Affect, Intuition, and Feelings of Meaning." *Journal of Personality and Social Psychology* 98.6, pp. 967–979 (cit. on p. 128).

- Hogarth, R. M. (2001). *Educating Intuition*. University of Chicago Press (cit. on p. 255).
- (2011). “On the Learning of Intuition”. *Intuition in Judgment and Decision Making*, pp. 111–126 (cit. on p. 255).
- Holroyd, J. (2012). “Responsibility for Implicit Bias”. *Journal of Social Philosophy* 43.3, pp. 274–306 (cit. on p. 43).
- Hookway, C. (1993). “Mimicking Foundationalism: On Sentiment and Self-Control”. *European Journal of Philosophy* 1.2, pp. 156–174 (cit. on pp. 236, 244).
- (1998). “Doubt: Affective States and the Regulation of Inquiry”. *Canadian Journal of Philosophy* 28 (sup1), pp. 203–225 (cit. on pp. 121, 128, 239).
- (2002). “Emotions and Epistemic Evaluations”. *The Cognitive Basis of Science*, p. 251 (cit. on pp. 121, 241).
- (2003). “Affective States and Epistemic Immediacy”. *Metaphilosophy* 34.1-2, pp. 78–96 (cit. on p. 121).
- Horgan, T. and Tienson, J. (2002). “The Intentionality of Phenomenology and the Phenomenology of Intentionality”. *Philosophy of Mind: Classical and Contemporary Readings*, pp. 520–533 (cit. on pp. 30, 34, 217).
- Horr, N. K., Braun, C., and Volz, K. G. (2014). “Feeling before Knowing Why: The Role of the Orbitofrontal Cortex in Intuitive Judgments—an MEG Study”. *Cognitive, Affective, & Behavioral Neuroscience* 14.4, pp. 1271–1285 (cit. on p. 153).
- Howe, C. Q. and Purves, D. (2005). “The Muller-Lyer Illusion Explained by the Statistics of Image-Source Relationships”. *Proceedings of the National Academy of Sciences* 102.4, pp. 1234–1239 (cit. on p. 223).
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield (cit. on p. 251).
- (2007). “Compassionate Phenomenal Conservatism”. *Philosophy and Phenomenological Research* 74.1, pp. 30–55 (cit. on pp. 67, 68).
- Isberner, M.-B. and Richter, T. (2014). “Comprehension and Validation: Separable Stages of Information Processing? A Case for Epistemic Monitoring in Language Comprehension”. *Processing Inaccurate Information: Theoretical and Applied Perspectives from Cognitive Science and the Educational Sciences*, pp. 245–276 (cit. on pp. 237, 239–241, 268).
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. MIT Press (cit. on p. 214).
- Jackson, M. B. (2016). “On the Epistemic Value of Imagining, Supposing, and Conceiving”. *Knowledge Through Imagination* (cit. on p. 210).

- Jacobs, R. A. (2002). “What Determines Visual Cue Reliability?” *Trends in Cognitive Sciences* 6.8, pp. 345–350 (cit. on pp. 166, 223).
- Jacoby, L. L. and Dallas, M. (1981). “On the Relationship between Autobiographical Memory and Perceptual Learning.” *Journal of Experimental Psychology: General* 110.3, p. 306 (cit. on pp. 135, 137, 142).
- Jacoby, L. L. and Whitehouse, K. (1989). “An Illusion of Memory: False Recognition Influenced by Unconscious Perception.” *Journal of Experimental Psychology: General* 118.2, p. 126 (cit. on pp. 124, 207).
- James, W. (1882). “On Some Hegelisms”. *Mind* 7.26, pp. 186–208 (cit. on pp. 29, 185).
 – (1884). “What Is an Emotion?” *Mind* 9.34, pp. 188–205 (cit. on pp. 85, 88, 111, 161).
 – (1890). *The Principles of Psychology*. Henry Holt and Company (cit. on pp. 124, 127, 143, 171, 176, 189).
- Jönsson, F. U. and Olsson, M. J. (2003). “Olfactory Metacognition”. *Chemical Senses* 28.7, pp. 651–658 (cit. on p. 128).
- Jönsson, F. U. and Stevenson, R. J. (2014). “Odor Knowledge, Odor Naming and the “Tip of the Nose” Experience”. *Tip-of-the-Tongue States and Related Phenomena*, p. 305 (cit. on p. 128).
- Jung, C. G. (1971). *Psychological Types*. Princeton University Press (cit. on p. 23).
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Penguin (cit. on p. 24).
- Kauppinen, A. (2013). “A Humean Theory of Moral Intuition”. *Canadian Journal of Philosophy* 43.3, pp. 360–381 (cit. on pp. 47, 84).
- Kelley, C. M. and Jacoby, L. L. (1998). “Subjective Reports and Process Dissociation: Fluency, Knowing, and Feeling”. *Acta Psychologica* 98.2-3, pp. 127–140 (cit. on pp. 135, 236).
- Kelley, C. M. and Lindsay, D. S. (1993). “Remembering Mistaken for Knowing: Ease of Retrieval as a Basis for Confidence in Answers to General Knowledge Questions”. *Journal of Memory and Language* 32.1, pp. 1–24 (cit. on p. 202).
- Kenny, A. (1963). *Action, Emotion and Will*. London: Routledge & Kegan Paul (cit. on p. 102).
- Kensinger, E. A. and Corkin, S. (2004). “Two Routes to Emotional Memory: Distinct Neural Processes for Valence and Arousal”. *Proceedings of the National Academy of Sciences* 101.9, pp. 3310–3315 (cit. on pp. 119, 153).
- Kind, A. (2016). “Introduction: Exploring Imagination”. *The Routledge Handbook of Philosophy of Imagination*, pp. 1–12 (cit. on pp. 210, 213).
- Klauer, K., Musch, J., and Naumer, B. (2000). “On Belief Bias in Syllogistic Reasoning”. *Psychological Review* 107.4, pp. 852–884 (cit. on p. 192).

- Klausen, S. H. (2013). “Losing Belief, While Keeping up the Attitudes: The Case for Cognitive Phenomenology”. *New Essays on Belief*, pp. 188–208 (cit. on p. 54).
- Klein, C. (2015). “What Pain Asymbolia Really Shows”. *Mind* 124.494, pp. 493–516 (cit. on p. 85).
- Klein, S. (2013). “Making the Case That Episodic Recollection Is Attributable to Operations Occurring at Retrieval Rather than to Content Stored in a Dedicated Subsystem of Long-Term Memory”. *Frontiers in Behavioral Neuroscience* 7 (cit. on p. 143).
- Klein, S. B. (2014). “Autonoesis and Belief in a Personal Past: An Evolutionary Theory of Episodic Memory Indices”. *Review of Philosophy and Psychology* 5.3, pp. 427–447 (cit. on p. 143).
- Klin, C. M., Guzmán, A. E., and Levine, W. H. (1997). “Knowing That You Don’t Know: Metamemory and Discourse Processing.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23.6, p. 1378 (cit. on p. 126).
- Knight, E. J., Klepac, K. M., and Kralik, J. D. (2013). “Too Good to Be True: Rhesus Monkeys React Negatively to Better-than-Expected Offers”. *PLoS ONE* 8.10, e75768 (cit. on p. 94).
- Koch, A. S. and Forgas, J. P. (2012). “Feeling Good and Feeling Truth: The Interactive Effects of Mood and Processing Fluency on Truth Judgments”. *Journal of Experimental Social Psychology* 48.2, pp. 481–485 (cit. on pp. 204, 205).
- Koch, A., Alves, H., Krüger, T., and Unkelbach, C. (2016). “A General Valence Asymmetry in Similarity: Good Is More Alike than Bad.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 42.8, p. 1171 (cit. on p. 157).
- Koksvik, O. (forthcoming). “Intuition: A Brief Introduction”. *Methods in Analytic Philosophy: A Contemporary Reader* (cit. on pp. 51, 67).
- (2011). “Intuition”. PhD Thesis. Australian National University (cit. on pp. 23, 28, 30–32, 35, 36, 38–41, 43, 48, 52, 57, 59, 62, 67, 73, 79, 117, 194, 253).
 - (2013). “Intuition and Conscious Reasoning”. *The Philosophical Quarterly* 63.253, pp. 709–715 (cit. on pp. 59, 72).
 - (2014). “Three Models of Phenomenal Unity”. *Journal of Consciousness Studies* 21.7–8, pp. 105–131 (cit. on p. 91).
 - (2017). “The Phenomenology of Intuition”. *Philosophy Compass* 12.1, e12387 (cit. on pp. 29, 30, 40, 73, 119, 163).
- Kolers, P. A. and Paley, S. R. (1976). “Knowing Not”. *Memory & Cognition* 4.5, pp. 553–558 (cit. on p. 126).

- Koriat, A. and Levy-Sadot, R. (2001). “The Combined Contributions of the Cue-Familiarity and Accessibility Heuristics to Feelings of Knowing”. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27.1, pp. 34–53 (cit. on p. 140).
- Koriat, A. and Lieblich, I. (1974). “What Does a Person in a ”TOT” State Know That a Person in a ”Don’t Know” State Doesn’t Know”. *Memory & Cognition* 2.4, pp. 647–655 (cit. on p. 126).
- Koriat, A. (1993). “How Do We Know That We Know? The Accessibility Model of the Feeling of Knowing”. *Psychological Review* 100.4, pp. 609–639 (cit. on pp. 134, 137).
- (1994). “Memory’s Knowledge of Its Own Knowledge: The Accessibility Account of the Feeling of Knowing”. *Metacognition: Knowing about Knowing*, pp. 115–135 (cit. on pp. 126, 134).
- (1995). “Dissociating Knowing and the Feeling of Knowing: Further Evidence for the Accessibility Model.” *Journal of Experimental Psychology: General* 124.3, pp. 311–333 (cit. on p. 175).
- (2000). “The Feeling of Knowing: Some Metatheoretical Implications for Consciousness and Control”. *Consciousness and Cognition* 9.2, pp. 149–171 (cit. on pp. 121, 122, 126).
- (2007). “Metacognition and Consciousness”. *The Cambridge Handbook of Consciousness*, pp. 289–325 (cit. on p. 134).
- (2012). “The Self-Consistency Model of Subjective Confidence.” *Psychological Review* 119.1, pp. 80–113 (cit. on pp. 128, 224).
- Koriat, A. and Goldsmith, M. (1994). “Memory in Naturalistic and Laboratory Contexts: Distinguishing the Accuracy-Oriented and Quantity-Oriented Approaches to Memory Assessment.” *Journal of Experimental Psychology: General* 123.3, p. 297 (cit. on p. 134).
- Koriat, A. and Levy-Sadot, R. (1999). “Processes Underlying Metacognitive Judgments: Information-Based and Experience-Based Monitoring of One’s Own Knowledge.” (cit. on pp. 126, 133).
- Koriat, A., Levy-Sadot, R., Edry, E., and de Marcas, S. (2003). “What Do We Know About What We Cannot Remember? Accessing the Semantic Attributes of Words That Cannot Be Recalled.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29.6, pp. 1095–1105 (cit. on p. 175).
- Kozuch, B. (2018). “No Pain, No Gain (in Darwinian Fitness): A Representational Account of Affective Experience”. *Erkenntnis* (cit. on p. 177).
- Kriegel, U. (2013). *Phenomenal Intentionality*. Oxford University Press (cit. on p. 30).

- Kriegel, U. (2014). “Towards a New Feeling Theory of Emotion: Towards a New Feeling Theory of Emotion”. *European Journal of Philosophy* 22.3, pp. 420–442 (cit. on pp. 89, 98).
- Kuhbandner, C. and Zehetleitner, M. (2011). “Dissociable Effects of Valence and Arousal in Adaptive Executive Control”. *PLoS ONE* 6.12, e29287 (cit. on p. 95).
- Kumkale, G. T. and Albarracín, D. (2004). “The Sleeper Effect in Persuasion: A Meta-Analytic Review.” *Psychological bulletin* 130.1, p. 143 (cit. on p. 198).
- Kunda, Z. and Thagard, P. (1996). “Forming Impressions from Stereotypes, Traits, and Behaviors: A Parallel-Constraint-Satisfaction Theory.” *Psychological review* 103.2, p. 284 (cit. on p. 232).
- Kurth, C., Crosby, H., and Basse, J. (2018). “Are Emotions Perceptions of Value (and Why This Matters)? A Review Essay of Christine Tappolet’s *Emotions, Values, and Agency*”. *Philosophical Psychology* 31.4, pp. 483–499 (cit. on p. 251).
- Lamarque, P. (1981). “How Can We Fear and Pity Fictions?” *The British Journal of Aesthetics* 21.4, pp. 291–304 (cit. on p. 114).
- Lander, K. and Metcalfe, S. (2007). “The Influence of Positive and Negative Facial Expressions on Face Familiarity”. *Memory* 15.1, pp. 63–69 (cit. on p. 156).
- Landwehr, J. R., Golla, B., and Reber, R. (2017). “Processing Fluency: An Inevitable Side Effect of Evaluative Conditioning”. *Journal of Experimental Social Psychology* 70, pp. 124–128 (cit. on p. 229).
- Lang, P. J., Bradley, M. M., Fitzsimmons, J. R., Cuthbert, B. N., Scott, J. D., Moulder, B., and Nangia, V. (1998). “Emotional Arousal and Activation of the Visual Cortex: An fMRI Analysis”. *Psychophysiology* 35.2, pp. 199–210 (cit. on p. 119).
- Lange, C. G. (1887). *Über Gemütsbewegungen. Ihr Wesen Und Ihr Einfluß Auf Körperliche, Besonders Auf Krankhafte Lebenserscheinungen. Ein Medizinisch-Psychologische Studie.* T. Thomas (cit. on p. 85).
- Larsen, J. T. and McGraw, A. P. (2014). “The Case for Mixed Emotions”. *Social and Personality Psychology Compass* 8.6, pp. 263–274 (cit. on p. 94).
- Lazarus, R. S. (1991). *Emotion and Adaptation.* Oxford University Press (cit. on pp. 89, 102).
- Le Bon, G. (2001). *The Crowd: A Study of the Popular Mind.* Dover (cit. on p. 238).
- Lev-Ari, S. and Keysar, B. (2010). “Why Don’t We Believe Non-Native Speakers? The Influence of Accent on Credibility”. *Journal of experimental social psychology* 46.6, pp. 1093–1096 (cit. on p. 202).

- Levy, N. (2015). “Neither Fish nor Fowl: Implicit Attitudes as Patchy Endorsements: Neither Fish nor Fowl: Implicit Attitudes as Patchy Endorsements”. *Noûs* 49.4, pp. 800–823 (cit. on p. 239).
- (2016). “Have I Turned the Stove Off? Explaining Everyday Anxiety”. *Philosopher’s Imprint* 16.2 (cit. on pp. 25, 83).
- Lewis, D. (1978). “Truth in Fiction”. *American Philosophical Quarterly* 15.1, pp. 37–46 (cit. on p. 114).
- (1983). *Philosophical Papers Vol. I*. Oxford University Press (cit. on pp. 48, 61).
- Lewis, P., Critchley, H., Rotshtein, P., and Dolan, R. (2006). “Neural Correlates of Processing Valence and Arousal in Affective Words”. *Cerebral Cortex* 17.3, pp. 742–748 (cit. on p. 95).
- Lodge, J. M., Kennedy, G., Lockyer, L., Arguel, A., and Pachman, M. (2018). “Understanding Difficulties and Resulting Confusion in Learning: An Integrative Review”. *Frontiers in Education* 3 (cit. on p. 174).
- Ludwig, K. (2007). “The Epistemology of Thought Experiments: First Person versus Third Person Approaches”. *Midwest Studies in Philosophy* 31.1, pp. 128–159 (cit. on pp. 48, 51, 242).
- Lynch, M. P. (2006). “Trusting Intuitions”. *Truth and Realism*, pp. 227–238 (cit. on pp. 25, 48, 59).
- Lyons, J. C. (2018). “Perception and Intuition of Evaluative Properties”. *Evaluative Perception* (cit. on p. 54).
- Mandelbaum, E. (2016). “Attitude, Inference, Association: On the Propositional Structure of Implicit Bias”. *Noûs* 50.3, pp. 629–658 (cit. on p. 239).
- Mandik, P. (2006). “The Introspectibility of Brain States as Such”. *Paul Churchland*, pp. 66–87 (cit. on p. 136).
- Mangan, B. (2000). “What Feeling Is the “Feeling of Knowing?””. *Consciousness and Cognition* 9.4, pp. 538–544 (cit. on p. 189).
- (2001). “Sensation’s Ghost”. *Psyche* 7, p. 18 (cit. on pp. 121, 125, 139, 171, 172, 183, 189).
- Marks, J. (1982). “A Theory of Emotion”. *Philosophical Studies* 42.2, pp. 227–242 (cit. on p. 57).
- Marsh, E. J., Cantor, A. D., and Brashier, N. M. (2016). “Believing That Humans Swallow Spiders in Their Sleep: False Beliefs as Side Effects of the Processes That Support Accurate Knowledge”. *Psychology of Learning and Motivation*. Vol. 64, pp. 93–132 (cit. on p. 238).

- Martin, J.-R. and Dokic, J. (2013). “Seeing Absence or Absence of Seeing?” *Thought: A Journal of Philosophy* 2.2, pp. 117–125 (cit. on p. 117).
- Martínez, M. (2015). “Pains as Reasons”. *Philosophical Studies* 172.9, pp. 2261–2274 (cit. on p. 109).
- Matchar, E. (2019). *Finding What’s ‘Oddly Satisfying’ on the Internet*. URL: <https://www.nytimes.com/2019/02/22/opinion/sunday/oddly-satisfying-videos-internet.html> (visited on 08/22/2019) (cit. on p. 183).
- Matthen, M. (2010). “On the Diversity of Auditory Objects”. *Review of Philosophy and Psychology* 1.1, pp. 63–89 (cit. on p. 166).
- Mauss, I. B. and Robinson, M. D. (2009). “Measures of Emotion: A Review”. *Cognition & Emotion* 23.2, pp. 209–237 (cit. on p. 151).
- Maynes, J. and Gross, S. (2013). “Linguistic Intuitions: Linguistic Intuitions”. *Philosophy Compass* 8.8, pp. 714–730 (cit. on p. 77).
- McAllister, B. (2018). “Seemings as Sui Generis”. *Synthese* 195.7, pp. 3079–3096 (cit. on p. 68).
- McGahhey, M. and Van Leeuwen, N. (2018). “Interpreting Intuitions”. *Third-Person Self-Knowledge, Self-Interpretation, and Narrative*, pp. 73–98 (cit. on pp. 30, 83).
- McGlone, M. S. and Tofighbakhsh, J. (2000). “Birds of a Feather Flock Conjointly (?): Rhyme as Reason in Aphorisms”. *Psychological science* 11.5, pp. 424–428 (cit. on p. 202).
- McLaughlin, B. P. (2010). “Monothematic Delusions and Existential Feelings”. *Delusion and Self-Deception: Affective and Motivational Influences on Belief Formation*, pp. 139–164 (cit. on pp. 115, 122, 254).
- Mendelovici, A. (2014). “Pure Intentionalism about Moods and Emotions”. *Current Controversies in Philosophy of Mind*, pp. 135–157 (cit. on pp. 98–100, 103).
- Merker, B. (2007). “Consciousness without a Cerebral Cortex: A Challenge for Neuroscience and Medicine”. *Behavioral and Brain Sciences* 30.01 (cit. on p. 93).
- Metcalf, J., Schwartz, B. L., and Joaquim, S. G. (1993). “The Cue-Familiarity Heuristic in Metacognition.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19.4, p. 851 (cit. on p. 133).
- Meyer, W.-U., Niepel, M., Rudolph, U., and Schützwohl, A. (1991). “An Experimental Analysis of Surprise”. *Cognition & Emotion* 5.4, pp. 295–311 (cit. on p. 128).
- Meyer, W.-U., Reisenzein, R., and Schützwohl, A. (1997). “Toward a Process Analysis of Emotions: The Case of Surprise”. *Motivation and Emotion* 21.3, pp. 251–274 (cit. on p. 128).

- Michaelian, K. (2012). “Metacognition and Endorsement”. *Mind & Language* 27.3, pp. 284–307 (cit. on p. 176).
- Michaelian, K. and Arango-Muñoz, S. (2014). “Epistemic Feelings, Epistemic Emotions: Review and Introduction to the Focus Section”. *Philosophical Inquiries* 2.1, pp. 97–122 (cit. on p. 126).
- Michal, M., Reuchlein, B., Adler, J., Reiner, I., Beutel, M. E., Vögele, C., Schächinger, H., and Schulz, A. (2014). “Striking Discrepancy of Anomalous Body Experiences with Normal Interoceptive Accuracy in Depersonalization-Derealization Disorder”. *PLOS ONE* 9.2, e89823 (cit. on p. 87).
- Millikan, R. G. (1984). *Language, Thought and Other Biological Objects*. MIT Press (cit. on p. 30).
- Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. MIT Press (cit. on p. 30).
- (1995). “Pushmi-Pullyu Representations”. *Philosophical perspectives* 9, pp. 185–200 (cit. on p. 72).
- Mitchell, J. (2018). “The Irreducibility of Emotional Phenomenology”. *Erkenntnis* 1, pp. 1–28 (cit. on pp. 54, 86, 89, 173, 252).
- (2019). “Emotional Intentionality and the Attitude-Content Distinction”. *Pacific Philosophical Quarterly* 0.0 (cit. on pp. 103, 105, 111, 165, 252).
- Mitchell, K. J. and Johnson, M. K. (2000). “Source Monitoring: Attributing Mental Experiences”. *The Oxford Handbook of Memory*, pp. 179–195 (cit. on pp. 199, 243).
- Molyneux, B. (2014). “New Arguments That Philosophers Don’t Treat Intuitions as Evidence”. *Metaphilosophy* 45.3, pp. 441–461 (cit. on pp. 24, 30, 47, 72).
- Monin, B. (2003). “The Warm Glow Heuristic: When Liking Leads to Familiarity.” *Journal of Personality and Social Psychology* 85.6, pp. 1035–1048 (cit. on p. 156).
- Montague, M. (2007). “Against Propositionalism”. *Nous* 41.3, pp. 503–518 (cit. on p. 30).
- Morris, A. L., Cleary, A. M., and Still, M. L. (2008). “The Role of Autonomic Arousal in Feelings of Familiarity”. *Consciousness and Cognition* 17.4, pp. 1378–1385 (cit. on p. 153).
- Morsanyi, K. and Handley, S. J. (2012). “Logic Feels so Good—I like It! Evidence for Intuitive Detection of Logicality in Syllogistic Reasoning.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38.3, pp. 596–616 (cit. on pp. 153, 195, 196, 209).
- Morton, A. (2009). “Epistemic Emotions”. *The Oxford Handbook of Philosophy of Emotion* (cit. on p. 121).

- Mulligan, K. (1998). “From Appropriate Emotions to Values”. *The Monist* 81.1, pp. 161–188 (cit. on p. 108).
- (2007). “Intentionality, Knowledge and Formal Objects”. *Disputatio* 2.23, pp. 205–228 (cit. on p. 42).
- Myers, I. B. (1962). *The Myers-Briggs Type Indicator: Manual (1962)*. Consulting Psychologists Press (cit. on p. 23).
- Nadarevic, L. and Aßfalg, A. (2017). “Unveiling the Truth: Warnings Reduce the Repetition-Based Truth Effect”. *Psychological Research* 81.4, pp. 814–826 (cit. on p. 197).
- Nadarevic, L. and Erdfelder, E. (2014). “Initial Judgment Task and Delay of the Final Validity-Rating Task Moderate the Truth Effect”. *Consciousness and Cognition* 23, pp. 74–84 (cit. on pp. 197, 198).
- Nado, J. (2014). “Why Intuition?” *Philosophy and Phenomenological Research* 89.1, pp. 15–41 (cit. on p. 242).
- (2016). “The Intuition Deniers”. *Philosophical Studies* 173.3, pp. 781–800 (cit. on p. 47).
- (2017). “Demythologizing Intuition”. *Inquiry* 60.4, pp. 386–402 (cit. on p. 47).
- Nanay, B. (2015). “Perceptual Content and the Content of Mental Imagery”. *Philosophical Studies* 172.7, pp. 1723–1736 (cit. on p. 79).
- Nelson, T. O. and Narens, L. (1994). “Why Investigate Metacognition”. *Metacognition: Knowing about Knowing*, pp. 1–25 (cit. on p. 145).
- Newman, E. J., Garry, M., Bernstein, D. M., Kantner, J., and Lindsay, D. S. (2012). “Nonprobative Photographs (or Words) Inflate Truthiness”. *Psychonomic Bulletin & Review* 19.5, pp. 969–974 (cit. on pp. 202, 204).
- Nimtz, C. (2010). “Saving the Doxastic Account of Intuitions”. *Philosophical Psychology* 23.3, pp. 357–375 (cit. on pp. 38, 53, 62).
- Noordewier, M. K. and Breugelmans, S. M. (2013). “On the Valence of Surprise”. *Cognition & Emotion* 27.7, pp. 1326–1334 (cit. on p. 94).
- Noordewier, M. K., Topolinski, S., and Van Dijk, E. (2016). “The Temporal Dynamics of Surprise: The Temporal Dynamics of Surprise”. *Social and Personality Psychology Compass* 10.3, pp. 136–149 (cit. on p. 94).
- Nussbaum, M. C. (2001). *Upheavals of Thought: A Theory of the Emotions*. Cambridge: Cambridge University Press (cit. on p. 251).
- O’Brien, E. J. and Cook, A. E. (2016). “Coherence Threshold and the Continuity of Processing: The RI-Val Model of Comprehension”. *Discourse Processes* 53.5-6, pp. 326–338 (cit. on p. 239).

- Olds, J. M. and Westerman, D. L. (2012). “Can Fluency Be Interpreted as Novelty? Retraining the Interpretation of Fluency in Recognition Memory”. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38.3, pp. 653–664 (cit. on p. 236).
- Oliver-Skuse, T. (2016). “Emotions and Representation”. University of Melbourne (cit. on pp. 75, 108, 252).
- Öllinger, M. and von Müller, A. (2017). “Search and Coherence-Building in Intuition and Insight Problem Solving”. *Frontiers in Psychology* 8 (cit. on pp. 197, 232).
- Oppenheimer, D. M. (2008). “The Secret Life of Fluency”. *Trends in cognitive sciences* 12.6, pp. 237–241 (cit. on p. 157).
- Ozubko, J. D. and Fugelsang, J. (2011). “Remembering Makes Evidence Compelling: Retrieval from Memory Can Give Rise to the Illusion of Truth.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 37.1, p. 270 (cit. on p. 198).
- Pacherie, E. (2002). “Emotion and Action”. *European Review of Philosophy* 5, pp. 55–90 (cit. on pp. 162, 163).
- (2010). “Perception, Emotions, and Delusions: The Case of the Capgras Delusion”. *Delusion and Self-Deception: Affective and Motivational Influences on Belief Formation*, pp. 107–125 (cit. on p. 152).
- Pennycook, G., Cannon, T. D., and Rand, D. G. (2018). “Prior Exposure Increases Perceived Accuracy of Fake News.” *Journal of Experimental Psychology: General* 147.12, pp. 1865–1880 (cit. on pp. 200, 242).
- Perry, J. (1979). “The Problem of the Essential Indexical”. *Nous*, pp. 3–21 (cit. on p. 163).
- Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., and Frith, C. D. (2007). “How the Brain Translates Money into Force: A Neuroimaging Study of Subliminal Motivation”. *Science* 316.5826, pp. 904–906 (cit. on p. 92).
- Pessoa, L. (2008). “On the Relationship between Emotion and Cognition”. *Nature reviews neuroscience* 9.2, p. 148 (cit. on p. 97).
- (2011). “Reprint of: Emotion and Cognition and the Amygdala: From “What Is It?” To “What’s to Be Done?””. *Neuropsychologia* 49.4, pp. 681–694 (cit. on p. 153).
- Pessoa, L. and Adolphs, R. (2010). “Emotion Processing and the Amygdala: From a ‘low Road’ to ‘Many Roads’ of Evaluating Biological Significance”. *Nature Reviews Neuroscience* 11, pp. 773–782 (cit. on p. 97).
- Phaf, R. H. and Rotteveel, M. (2005). “Affective Modulation of Recognition Bias.” *Emotion* 5.3, pp. 309–318 (cit. on p. 156).

- Phelps, E. A., Ling, S., and Carrasco, M. (2006). “Emotion Facilitates Perception and Potentiates the Perceptual Benefits of Attention”. *Psychological Science* 17.4, pp. 292–299 (cit. on pp. 119, 153).
- Picard, F. (2013). “State of Belief, Subjective Certainty and Bliss as a Product of Cortical Dysfunction”. *Cortex* 49.9, pp. 2494–2500 (cit. on p. 184).
- Plailly, J., Tillmann, B., and Royet, J.-P. (2007). “The Feeling of Familiarity of Music and Odors: The Same Neural Signature?” *Cerebral Cortex* 17.11, pp. 2650–2658 (cit. on p. 124).
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford University Press (cit. on p. 48).
- Ploner, M. (1999). “Pain Affect without Pain Sensation in a Patient with a Postcentral Lesion”, p. 4 (cit. on pp. 85, 86).
- Poher, J. (2018). “What Emotions Really Are (In the Theory of Constructed Emotions)”. *Philosophy of Science* 85.4, pp. 640–659 (cit. on p. 165).
- Price, D. D. (2000). “Psychological and Neural Mechanisms of the Affective Dimension of Pain”. *Science* 288.5472, pp. 1769–1772 (cit. on p. 85).
- Price, H. H. (1969). *Belief: The Gifford Lectures Delivered at the University of Aberdeen in 1960*. New York: Humanities P. (cit. on p. 60).
- Prinz, J. (2004a). “Embodied Emotions”. *Thinking about Feeling: Contemporary Philosophers on Emotions*, pp. 44–58 (cit. on p. 108).
- (2005). “Are Emotions Feelings?” *Journal of Consciousness Studies* 12.8-9, pp. 9–25 (cit. on p. 93).
- (2014). “Seeing with Feeling”. *Aesthetics and the Sciences of Mind*, pp. 143–158 (cit. on pp. 113, 164, 169).
- Prinz, J. J. (2004b). *Gut Reactions: A Perceptual Theory of Emotion*. Oxford University Press (cit. on pp. 85, 93, 94, 98, 102, 173, 251).
- (2006a). “Is Emotion a Form of Perception?” *Canadian Journal of Philosophy* 36 (sup1), pp. 137–160 (cit. on pp. 99, 102, 223).
- (2006b). “Is Emotion a Form of Perception?” *Canadian Journal of Philosophy* 36 (sup1), pp. 137–160 (cit. on p. 251).
- (2007). “All Consciousness Is Perceptual”. *Contemporary Debates in Philosophy of Mind* (cit. on pp. 60, 84, 121).
- (2011). “The Sensory Basis of Cognitive Phenomenology”. *Cognitive Phenomenology*, pp. 174–196 (cit. on pp. 121, 150).

- Proust, J. (2008). “Epistemic Agency and Metacognition: An Externalist View”. *Proceedings of the Aristotelian Society* 108 (1pt3), pp. 241–268 (cit. on pp. 173, 177, 191).
- (2010). “Metacognition”. *Philosophy Compass* 5.11, pp. 989–998 (cit. on pp. 146, 162).
 - (2013). “Mental Acts as Natural Kinds”. *Decomposing the Will*, pp. 262–282 (cit. on pp. 122, 146).
 - (2014). “Précis of The Philosophy of Metacognition”. *Philosophy and Phenomenological Research* 89.3, pp. 703–709 (cit. on p. 130).
 - (2015). “The Representational Structure of Feelings”. *Open MIND* (cit. on pp. 86, 94, 96, 150, 162, 190).
- Pryor, J. (2013). “There Is Immediate Justification”. *Contemporary Debates in Epistemology*. Second Edition, pp. 202–222 (cit. on p. 33).
- Pust, J. (2000). *Intuitions as Evidence*. Routledge (cit. on p. 67).
- (2014). *Intuition*. In collab. with E. N. Zalta. URL: <http://plato.stanford.edu/archives/fall2014/entries/intuition/> (visited on 10/11/2015) (cit. on pp. 29, 62).
- Quilty-Dunn, J. (2015). “Believing Our Eyes: The Role of False Belief in the Experience of Cinema”. *The British Journal of Aesthetics* 55.3, pp. 269–283 (cit. on p. 25).
- Quinn, W. (1993). “Putting Rationality in Its Place”. *Morality and Action*, pp. 228–255 (cit. on p. 63).
- Radford, C. and Weston, M. (1975). “How Can We Be Moved by the Fate of Anna Karenina?” *Aristotelian Society Supplementary Volume* 49.1, pp. 67–93 (cit. on p. 114).
- Rafal, R., Ward, R., and Danziger, S. (2006). “Selection for Action and Selection for Awareness: Evidence from Hemispatial Neglect”. *Brain Research* 1080.1, pp. 2–8 (cit. on p. 31).
- Rapp, B. and Goldrick, M. (2000). “Discreteness and Interactivity in Spoken Word Production.” *Psychological review* 107.3, p. 460 (cit. on p. 128).
- Ratcliffe, M. (2005). “The Feeling of Being”. *Journal of Consciousness Studies* 12.8-10, pp. 43–60 (cit. on p. 99).
- (2007a). “What Is a Feeling of Unfamiliarity?” *Philosophy, Psychiatry, & Psychology* 14.1, pp. 43–49 (cit. on p. 125).
 - (2007b). “What Is a Feeling of Unfamiliarity?” *Philosophy, Psychiatry, & Psychology* 14.1, pp. 43–49 (cit. on p. 152).
- Reber, R., Winkielman, P., and Schwarz, N. (1998). “Effects of Perceptual Fluency on Affective Judgments”. *Psychological Science* 9.1, pp. 45–48 (cit. on p. 153).

- Reber, R., Fazendeiro, T. A., and Winkielman, P. (2002). “Processing Fluency as the Source of Experiences at the Fringe of Consciousness”. *Psyche* 8.10, pp. 1–21 (cit. on p. 160).
- Reber, R. and Schwarz, N. (1999). “Effects of Perceptual Fluency on Judgments of Truth”. *Consciousness and cognition* 8.3, pp. 338–342 (cit. on pp. 153, 202).
- Reber, R., Schwarz, N., and Winkielman, P. (2004). “Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver’s Processing Experience?” *Personality and social psychology review* 8.4, pp. 364–382 (cit. on p. 169).
- Reber, R. and Unkelbach, C. (2010). “The Epistemic Status of Processing Fluency as Source for Judgments of Truth”. *Review of Philosophy and Psychology* 1.4, pp. 563–581 (cit. on pp. 199, 224, 225, 238).
- Reber, R., Wurtz, P., and Zimmermann, T. D. (2004). “Exploring “Fringe” Consciousness: The Subjective Experience of Perceptual Fluency and Its Objective Bases”. *Consciousness and cognition* 13.1, pp. 47–60 (cit. on pp. 135, 227).
- Recanati, F. (2012). *Mental Files*. Oxford University Press (cit. on pp. 138, 231).
- (2016). *Mental Files in Flux*. Oxford University Press (cit. on pp. 138, 231).
- Reder, L. M. and Ritter, F. E. (1992). “What Determines Initial Feeling of Knowing? Familiarity with Question Terms, Not with the Answer.” *Journal of Experimental Psychology: Learning, memory, and cognition* 18.3, p. 435 (cit. on pp. 133, 175).
- Reimer, M. (2010). “Only a Philosopher or a Madman: Impractical Delusions in Philosophy and Psychiatry”. *Philosophy, Psychiatry, & Psychology* 17.4, pp. 315–328 (cit. on p. 59).
- Reisenzein, R. (2012). “What Is an Emotion in the Belief-Desire Theory of Emotion?” *The Goals of Cognition: Essays in Honor of Cristiano Castelfranchi* (cit. on p. 57).
- Remmers, C., Topolinski, S., Buxton, A., Dietrich, D. E., and Michalak, J. (2017). “The Beneficial and Detrimental Effects of Major Depression on Intuitive Decision-Making”. *Cognition and Emotion* 31.4, pp. 799–805 (cit. on p. 154).
- Remmers, C., Topolinski, S., Dietrich, D. E., and Michalak, J. (2015). “Impaired Intuition in Patients with Major Depressive Disorder”. *British Journal of Clinical Psychology* 54.2, pp. 200–213 (cit. on p. 154).
- Remmers, C. and Zander, T. (2018). “Why You Don’t See the Forest for the Trees When You Are Anxious: Anxiety Impairs Intuitive Decision Making”. *Clinical Psychological Science* 6.1, pp. 48–62 (cit. on p. 154).
- Renoult, J. P. (2016). “The Evolution of Aesthetics: A Review of Models”. *Aesthetics and Neuroscience: Scientific and Artistic Perspectives*, pp. 271–299 (cit. on p. 169).

- Richter, T. (2015). “Validation and Comprehension of Text Information: Two Sides of the Same Coin”. *Discourse Processes* 52.5-6, pp. 337–355 (cit. on p. 239).
- Richter, T., Schroeder, S., and Wöhrmann, B. (2009). “You Don’t Have to Believe Everything You Read: Background Knowledge Permits Fast and Efficient Validation of Information.” *Journal of Personality and Social Psychology* 96.3, pp. 538–558 (cit. on p. 239).
- Roberts, R. C. (2003). *Emotions: An Essay in Aid of Moral Psychology*. Cambridge University Press (cit. on pp. 103, 252).
- (2009). “The Sophistication of Non-Human Emotion”. *The Philosophy of Animal Minds*, pp. 145–164 (cit. on p. 103).
- Rubins, J. L. and Friedman, E. D. (1948). “ASYMBOLIA FOR PAIN”. *Archives of Neurology & Psychiatry* 60.6, pp. 554–573 (cit. on p. 85).
- Russell, B. (1921). *The Analysis of Mind*. George Allen & Unwin (cit. on p. 143).
- Russell, J. A. (2003). “Core Affect and the Psychological Construction of Emotion”. *Psychological Review* 110.1, pp. 145–172 (cit. on p. 165).
- Russell, J. A. and Carroll, J. M. (1999). “On the Bipolarity of Positive and Negative Affect”. *Psychological bulletin* 125.1, p. 3 (cit. on p. 94).
- Ryle, G. (1949). *The Concept of Mind*. University of Chicago Press (cit. on p. 60).
- Salmela, M. (2011). “Can Emotion Be Modelled on Perception?” *Dialectica* 65.1, pp. 1–29 (cit. on p. 251).
- Sander, D., Grafman, J., and Zalla, T. (2003). “The Human Amygdala: An Evolved System for Relevance Detection”. *Reviews in the Neurosciences* 14.4, pp. 303–316 (cit. on p. 97).
- Scarantino, A. (2010). “Insights and Blindspots of the Cognitivist Theory of Emotions”. *The British Journal for the Philosophy of Science* 61.4, pp. 729–768 (cit. on pp. 59, 109, 215, 251).
- (2014). “The Motivational Theory of Emotions”. *Moral Psychology and Human Agency*, pp. 156–185 (cit. on p. 96).
- (2017). “Do Emotions Cause Actions, and If so How?” *Emotion Review* 9.4, pp. 326–334 (cit. on p. 96).
- Schapiro, T. (2009). “The Nature of Inclination”. *Ethics* 119.2, pp. 229–256 (cit. on p. 64).
- Scherer, K. R., Schorr, A., and Johnstone, T. (2001). *Appraisal Processes in Emotion: Theory, Methods, Research*. Oxford University Press (cit. on p. 165).
- Schneider, I. K. and Schwarz, N. (2017). “Mixed Feelings: The Case of Ambivalence”. *Current Opinion in Behavioral Sciences* 15, pp. 39–45 (cit. on p. 94).

- Schnider, A. (2013). “Orbitofrontal Reality Filtering”. *Frontiers in Behavioral Neuroscience* 7 (cit. on p. 243).
- Scholl, S. G., Greifeneder, R., and Bless, H. (2014). “When Fluency Signals Truth: Prior Successful Reliance on Fluency Moderates the Impact of Fluency on Truth Judgments”. *Journal of Behavioral Decision Making* 27.3, pp. 268–280 (cit. on p. 202).
- Schupp, H. T., Junghöfer, M., Weike, A. I., and Hamm, A. O. (2003). “Emotional Facilitation of Sensory Processing in the Visual Cortex”. *Psychological Science* 14.1, pp. 7–13 (cit. on pp. 119, 153).
- Schwartz, B. L. (1994). “Sources of Information in Metamemory: Judgments of Learning and Feelings of Knowing”. *Psychonomic Bulletin & Review* 1.3, pp. 357–375 (cit. on pp. 131, 227).
- Schwartz, B. L. and Brown, A. S. (2014). *Tip-of-the-Tongue States and Related Phenomena*. Cambridge University Press (cit. on p. 127).
- Schwartz, B. L. and Metcalfe, J. (2011). “Tip-of-the-Tongue (TOT) States: Retrieval, Behavior, and Experience”. *Memory & Cognition* 39.5, pp. 737–749 (cit. on pp. 127, 131).
- (2014). “Tip-of-the-Tongue (TOT) States: Mechanisms and Metacognitive Control”. *Tip-of-the-tongue states and related phenomena* 39.5, pp. 737–749 (cit. on pp. 131, 140, 178).
- Schwartz, B. L., Travis, D. M., Castro, A. M., and Smith, S. M. (2000). “The Phenomenology of Real and Illusory Tip-of-the-Tongue States”. *Memory & Cognition* 28.1, pp. 18–27 (cit. on pp. 139, 227).
- Schwartz, M. (1982). “Repetition and Rated Truth Value of Statements”. *The American Journal of Psychology* 95.3, pp. 393–407 (cit. on pp. 197, 198).
- Schwarz, N. (2004). “Metacognitive Experiences in Consumer Judgment and Decision Making”. *Journal of Consumer Psychology* 14.4, pp. 332–348 (cit. on p. 136).
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., and Simons, A. (1991). “Ease of Retrieval as Information: Another Look at the Availability Heuristic.” *Journal of Personality and Social Psychology* 61.2, p. 195 (cit. on pp. 137, 139).
- Schwarz, N. and Clore, G. L. (2007). “Feelings and Phenomenal Experiences”. *Social Psychology: Handbook of Basic Principles*. 2nd ed. Vol. 2, pp. 385–407 (cit. on p. 154).
- Schwitzgebel, E. (2008). “The Unreliability of Naive Introspection”. *Philosophical Review* 117.2, pp. 245–273 (cit. on p. 171).
- Schwitzgebel, E. (2002). “A Phenomenal, Dispositional Account of Belief”. *Noûs* 36.2, pp. 249–275 (cit. on p. 60).

- Segall, M. H., Campbell, D. T., and Herskovits, M. J. (1966). *The Influence of Culture on Visual Perception*. Bobbs-Merrill Indianapolis (cit. on p. 38).
- Shand, A. F. (F. (1914). *The Foundations of Character; Being a Study of the Tendencies of the Emotions and Sentiments*. In collab. with University of California Libraries. London : Macmillan and co., limited (cit. on p. 89).
- Sharot, T., Delgado, M. R., and Phelps, E. A. (2004). “How Emotion Enhances the Feeling of Remembering”. *Nature Neuroscience* 7.12, pp. 1376–1380 (cit. on pp. 119, 153).
- Shea, N. (2014). “Reward Prediction Error Signals Are Meta-Representational: Reward Prediction Error Signals Are Meta-Representational”. *Noûs* 48.2, pp. 314–341 (cit. on p. 146).
- Siegel, S. (2013). “The Epistemic Impact of the Etiology of Experience”. *Philosophical Studies* 162.3, pp. 697–722 (cit. on p. 109).
- (2014). “Affordances and the Contents of Perception”. *Does Perception Have Content?*, pp. 39–76 (cit. on p. 35).
- (2017). “How Is Wishful Seeing Like Wishful Thinking?” *Philosophy and Phenomenological Research* 95.2, pp. 408–435 (cit. on p. 110).
- Silva, R. R., Garcia-Marques, T., and Mello, J. (2016). “The Differential Effects of Fluency Due to Repetition and Fluency Due to Color Contrast on Judgments of Truth”. *Psychological research* 80.5, pp. 821–837 (cit. on pp. 226, 227, 234, 236).
- Silva, R. R., Garcia-Marques, T., and Reber, R. (2017). “The Informative Value of Type of Repetition: Perceptual and Conceptual Fluency Influences on Judgments of Truth”. *Consciousness and Cognition* 51, pp. 53–67 (cit. on pp. 198, 207, 234, 237).
- Silvia, P. J. (2010). “Confusion and Interest: The Role of Knowledge Emotions in Aesthetic Experience.” *Psychology of Aesthetics, Creativity, and the Arts* 4.2, pp. 75–80 (cit. on p. 174).
- Singer, M. (2013). “Validation in Reading Comprehension”. *Current Directions in Psychological Science* 22.5, pp. 361–366 (cit. on pp. 239, 243).
- Sinhababu, N. (2013). “The Desire-Belief Account of Intention Explains Everything”. *Noûs* 47.4, pp. 680–696 (cit. on p. 57).
- Sizer, L. (2006). “What Feelings Can’t Do”. *Mind and Language* 21.1, pp. 108–135 (cit. on p. 100).
- Smith, S. M. (1994). “Frustrated Feelings of Imminent Recall: On the Tip of the Tongue”. *Metacognition: Knowing about Knowing*, pp. 27–46 (cit. on p. 139).
- Smith, T. (2000). *Viable Values: A Study of Life as the Root and Reward of Morality*. Rowman & Littlefield Publishers (cit. on pp. 48, 49).

- Solomon, R. C. (1993). *The Passions: Emotions and the Meaning of Life*. Hackett (cit. on p. 251).
- (2003a). “Against Valence (‘Positive’ and ‘Negative’ Emotions)”. *Not Passion’s Slave: Emotions and Choice*, pp. 162–177 (cit. on p. 94).
 - (2003b). *Not Passion’s Slave: Emotions and Choice*. Oxford University Press (cit. on p. 251).
- Solomon, R. C. and Stone, L. D. (2002). “On “Positive” and “Negative” Emotions”. *Journal for the Theory of Social Behaviour* 32.4, pp. 417–435 (cit. on p. 94).
- Sosa, E. (1996). “Rational Intuition: Bealer on Its Nature and Epistemic Status”. *Philosophical Studies* 81.2-3, pp. 151–162 (cit. on pp. 35, 48).
- (1998). “Minimal Intuition”. *Rethinking Intuition*, pp. 257–269 (cit. on p. 48).
 - (2006). “Intuitions and Truth”. *Truth and Realism*, pp. 208–26 (cit. on p. 48).
 - (2007a). *A Virtue Epistemology*. Oxford University Press (cit. on pp. 30, 37, 48, 51, 63, 64, 242, 245).
 - (2007b). “Experimental Philosophy and Philosophical Intuition”. *Philosophical Studies* 132.1, pp. 99–107 (cit. on pp. 60, 255).
 - (2007c). “Intuitions: Their Nature and Epistemic Efficacy”. *Grazer Philosophische Studien* 74.1, pp. 51–67 (cit. on pp. 35, 37, 48, 64, 213, 242, 254, 255).
- Spaulding, S. (2016). “Imagination Through Knowledge”. *Knowledge Through Imagination*, pp. 207–226 (cit. on p. 213).
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., and Wilson, D. (2010). “Epistemic Vigilance”. *Mind & Language* 25.4, pp. 359–393 (cit. on pp. 177, 199, 218, 241).
- Sperber, D. and Wilson, D. (1987). “Précis of Relevance: Communication and Cognition”. *Behavioral and brain sciences* 10.4, pp. 697–710 (cit. on p. 225).
- Sperduti, M., Arcangeli, M., Makowski, D., Wantzen, P., Zalla, T., Lemaire, S., Dokic, J., Pelletier, J., and Piolino, P. (2016). “The Paradox of Fiction: Emotional Response toward Fiction and the Modulatory Role of Self-Relevance”. *Acta Psychologica* 165, pp. 53–59 (cit. on p. 115).
- Stock, K. (2017). *Only Imagine: Fiction, Interpretation and Imagination*. Oxford University Press (cit. on pp. 210, 213).
- Stocker, M. (1983). “Psychic Feelings: Their Importance and Irreducibility”. *Australasian Journal of Philosophy* 61.1, pp. 5–26 (cit. on pp. 84, 86, 89, 252).
- (2004). “Some Considerations About Intellectual Desire and Emotions”. *Thinking About Feeling: Contemporary Philosophers on Emotions* (cit. on p. 121).

- Stocker, M. (2009). “Intellectual and Other Non-Standard Emotions”. *The Oxford Handbook of Philosophy of Emotion* (cit. on p. 88).
- Stoerig, P. and Cowey, A. (2007). “Blindsight”. *Current Biology* 17.19, R822–R824 (cit. on p. 31).
- Strawson, G. (2009). *Mental Reality*. MIT Press (cit. on pp. 53, 62).
- Sullivan-Bissett, E. (2018). “Biased by Our Imaginings”. *Mind & Language* (cit. on p. 215).
- Sun, Y., Vinson, D. P., and Vigliocco, G. (1998). “Tip-of-the-Tongue and Tip-of-the-Pen in Chinese”. *Abstracts of the Psychonomic Society* 3, p. 32 (cit. on p. 128).
- Sweklej, J., Balas, R., Pochwatko, G., and Godlewska, M. (2014). “Intuitive (in)Coherence Judgments Are Guided by Processing Fluency, Mood and Affect”. *Psychological Research* 78.1, pp. 141–149 (cit. on pp. 154, 204).
- Talarico, J. M. and Rubin, D. C. (2007). “Flashbulb Memories Are Special after All; in Phenomenology, Not Accuracy”. *Applied Cognitive Psychology* 21.5, pp. 557–578 (cit. on p. 119).
- Talbot, B. (2013). “Reforming Intuition Pumps: When Are the Old Ways the Best?” *Philosophical Studies* 165.2, pp. 315–334 (cit. on p. 29).
- Tappolet, C. (2000). *Emotions et Valeurs*. FeniXX (cit. on p. 251).
- (2016). *Emotions, Value, and Agency*. Oxford University Press UK (cit. on p. 251).
- Teng, L. (2018). “Is Phenomenal Force Sufficient for Immediate Perceptual Justification?” *Synthese* 195.2, pp. 637–656 (cit. on p. 72).
- Teroni, F. (2007). “Emotions and Formal Objects”. *Dialectica* 61 (3,) (cit. on p. 102).
- (2018). “Emotionally Charged—The Puzzle of Affective Valence”. *Shadows of the Soul: Philosophical Perspectives on Negative Emotions*, pp. 10–19 (cit. on pp. 116, 166, 216).
- Thagard, P. (1989). “Explanatory Coherence”. *Behavioral and brain sciences* 12.3, pp. 435–467 (cit. on pp. 232, 241).
- (2004). “What Is Doubt and When Is It Reasonable?” *Canadian Journal of Philosophy* 34 (sup1), pp. 391–406 (cit. on pp. 121, 239).
- (2008). “How Cognition Meets Emotion: Beliefs, Desires, and Feelings as Neural Activity”. *Epistemology and emotions*, pp. 167–184 (cit. on p. 101).
- Thomas, A. K., Bulevich, J. B., and Dubois, S. J. (2011). “Context Affects Feeling-of-Knowing Accuracy in Younger and Older Adults.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 37.1, p. 96 (cit. on pp. 135, 228).
- (2012). “An Analysis of the Determinants of the Feeling of Knowing”. *Consciousness and cognition* 21.4, pp. 1681–1694 (cit. on p. 135).

- Thomas, A. K., Lee, M., and Hughes, G. (2016). “Introspecting on the Elusive: The Uncanny State of the Feeling of Knowing”. *Oxford Handbook of Metamemory*, pp. 81–94 (cit. on pp. 126, 131, 139, 226).
- Thompson, R., Emmorey, K., and Gollan, T. H. (2005). ““Tip of the Fingers” Experiences by Deaf Signers: Insights into the Organization of a Sign-Based Lexicon”. *Psychological Science* 16.11, pp. 856–860 (cit. on p. 128).
- Thompson, V. A. (2014). “What Intuitions Are... and Are Not”. *Psychology of Learning and Motivation*. Vol. 60, pp. 35–75 (cit. on pp. 160, 187).
- Thompson, V. A. and Johnson, S. C. (2014a). “Conflict, Metacognition, and Analytic Thinking”. *Thinking & Reasoning* 20.2, pp. 215–244 (cit. on p. 194).
- (2014b). “Conflict, Metacognition, and Analytic Thinking”. *Thinking & Reasoning* 20.2, pp. 215–244 (cit. on p. 196).
- Thompson, V. A., Prowse Turner, J. A., and Pennycook, G. (2011). “Intuition, Reason, and Metacognition”. *Cognitive Psychology* 63.3, pp. 107–140 (cit. on pp. 122, 183, 190, 193, 194, 196, 209).
- Thompson, V. A., Turner, J. A. P., Pennycook, G., Ball, L. J., Brack, H., Ophir, Y., and Ackerman, R. (2013). “The Role of Answer Fluency and Perceptual Fluency as Metacognitive Cues for Initiating Analytic Thinking”. *Cognition* 128.2, pp. 237–251 (cit. on pp. 194, 209).
- Thompson, V. and Morsanyi, K. (2012). “Analytic Thinking: Do You Feel like It?” *Mind & Society* 11.1, pp. 93–105 (cit. on pp. 183, 194, 209).
- Topolinski, S. and Strack, F. (2008). “Where There’s a Will—There’s No Intuition. The Unintentional Basis of Semantic Coherence Judgments”. *Journal of Memory and Language* 58.4, pp. 1032–1048 (cit. on p. 196).
- (2009a). “Scanning the “Fringe” of Consciousness: What Is Felt and What Is Not Felt in Intuitions about Semantic Coherence”. *Consciousness and Cognition* 18.3, pp. 608–618 (cit. on pp. 153, 158, 159, 228).
- (2009b). “The Analysis of Intuition: Processing Fluency and Affect in Judgements of Semantic Coherence”. *Cognition & Emotion* 23.8, pp. 1465–1503 (cit. on pp. 152, 153, 156–159, 228, 230, 242).
- (2009c). “The Architecture of Intuition: Fluency and Affect Determine Intuitive Judgments of Semantic and Visual Coherence and Judgments of Grammaticality in Artificial Grammar Learning.” *Journal of Experimental Psychology: General* 138.1, pp. 39–63 (cit. on pp. 153, 154, 158, 230).
- Trippas, D., Handley, S. J., Verde, M. F., and Morsanyi, K. (2016). “Logic Brightens My Day: Evidence for Implicit Sensitivity to Logical Validity.” *Journal of Experimental*

- Psychology: Learning, Memory, and Cognition* 42.9, pp. 1448–1457 (cit. on pp. 153, 195, 209).
- Trout, J. D. (2002). “Scientific Explanation And The Sense Of Understanding*”. *Philosophy of Science* 69.2, pp. 212–233 (cit. on p. 174).
- Tucker, C. (2013). “Seemings and Justification: An Introduction”. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, pp. 1–29 (cit. on pp. 68, 69).
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press (cit. on pp. 31, 98, 176).
- (2008). “The Experience of Emotion: An Intentionalist Theory”. *Revue internationale de philosophie* 243.1, pp. 25–50 (cit. on pp. 98, 114).
- Uhelski, M. L., Davis, M. A., and Fuchs, P. N. (2012). “Pain Affect in the Absence of Pain Sensation: Evidence of Asomaesthesia after Somatosensory Cortex Lesions in the Rat”. *PAIN®* 153.4, pp. 885–892 (cit. on p. 85).
- Unkelbach, C. (2006). “The Learned Interpretation of Cognitive Fluency”. *Psychological Science* 17.4, pp. 339–345 (cit. on p. 236).
- (2007). “Reversing the Truth Effect: Learning the Interpretation of Processing Fluency in Judgments of Truth.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 33.1, pp. 219–230 (cit. on pp. 153, 166, 199, 225, 226, 236, 238).
- Unkelbach, C., Bayer, M., Alves, H., Koch, A., and Stahl, C. (2011a). “Fluency and Positivity as Possible Causes of the Truth Effect”. *Consciousness and Cognition* 20.3, pp. 594–602 (cit. on p. 153).
- (2011b). “Fluency and Positivity as Possible Causes of the Truth Effect”. *Consciousness and Cognition* 20.3, pp. 594–602 (cit. on pp. 202, 204, 206).
- Unkelbach, C., Fiedler, K., Bayer, M., Stegmüller, M., and Danner, D. (2008). “Why Positive Information Is Processed Faster: The Density Hypothesis.” *Journal of personality and social psychology* 95.1, p. 36 (cit. on p. 157).
- Unkelbach, C. and Greifeneder, R. (2013a). *The Experience of Thinking: How the Fluency of Mental Processes Influences Cognition and Behaviour*. Psychology Press (cit. on pp. 135, 157).
- (2013b). *The Experience of Thinking: How the Fluency of Mental Processes Influences Cognition and Behaviour*. Psychology Press (cit. on pp. 136, 142, 202, 203).
- (2018). “Experiential Fluency and Declarative Advice Jointly Inform Judgments of Truth”. *Journal of Experimental Social Psychology* 79, pp. 78–86 (cit. on pp. 199, 228).

- Unkelbach, C. and Rom, S. C. (2017). “A Referential Theory of the Repetition-Induced Truth Effect”. *Cognition* 160, pp. 110–126 (cit. on pp. 226, 231, 233–235, 241).
- Unkelbach, C. and Stahl, C. (2009). “A Multinomial Modeling Approach to Dissociate Different Components of the Truth Effect”. *Consciousness and Cognition* 18.1, pp. 22–38 (cit. on pp. 202, 208, 227).
- Vaidya, A. J. (2010). “Philosophical Methodology: The Current Debate”. *Philosophical Psychology* 23.3, pp. 391–417 (cit. on pp. 72, 215).
- Van Berkum, J. J., Holleman, B., Nieuwland, M., Otten, M., and Murre, J. (2009). “Right or Wrong? The Brain’s Fast Response to Morally Objectionable Statements”. *Psychological Science* 20.9, pp. 1092–1099 (cit. on p. 240).
- Van De Cruys, S. (2017). “Affective Value in the Predictive Mind”. *Philosophy and Predictive Processing* (cit. on p. 157).
- Van Inwagen, P. (1997). “Materialism and the Psychological-Continuity Account of Personal Identity”. *Noûs* 31.1, pp. 305–319 (cit. on pp. 35, 36, 48, 52).
- Verde, M. F., Stone, L. K., Hatch, H. S., and Schnall, S. (2010). “Distinguishing between Attributional and Mnemonic Sources of Familiarity: The Case of Positive Emotion Bias”. *Memory & Cognition* 38.2, pp. 142–153 (cit. on pp. 141, 144, 145, 157).
- Volz, K. G., RübSamen, R., and von Cramon, D. Y. (2008). “Cortical Regions Activated by the Subjective Sense of Perceptual Coherence of Environmental Sounds: A Proposal for a Neuroscience of Intuition”. *Cognitive, Affective, & Behavioral Neuroscience* 8.3, pp. 318–328 (cit. on pp. 122, 153).
- Volz, K. G. and von Cramon, D. Y. (2006). “What Neuroscience Can Tell about Intuitive Processes in the Context of Perceptual Discovery”. *Journal of Cognitive Neuroscience* 18.12, pp. 2077–2087 (cit. on p. 153).
- Von Helversen, B., Gendolla, G. H. E., Winkielman, P., and Schmidt, R. E. (2008). “Exploring the Hardship of Ease: Subjective and Objective Effort in the Ease-of-Processing Paradigm”. *Motivation and Emotion* 32.1, pp. 1–10 (cit. on p. 152).
- Vuilleumier, P. (2005). “How Brains Beware: Neural Mechanisms of Emotional Attention”. *Trends in cognitive sciences* 9.12, pp. 585–594 (cit. on p. 97).
- Walton, K. (1994). “Listening with Imagination: Is Music Representational?” *The Journal of Aesthetics and Art Criticism* 52.1, pp. 47–61 (cit. on p. 213).
- Weinberg, J. M. (2007). “How to Challenge Intuitions Empirically without Risking Skepticism”. *Midwest Studies in Philosophy* 31.1, pp. 318–343 (cit. on pp. 40, 242, 255).
- Weinberg, J. M. and Alexander, J. (2014). “Intuitions through Thick and Thin”. *Intuitions*, pp. 186–209 (cit. on p. 255).

- Weiss, J. (2016). “A Feeling Theory of Feelings”. PhD Thesis. The Ohio State University (cit. on pp. 34, 90, 91, 98, 100, 103–105, 107, 109, 116, 216, 252).
- Werner, P. J. (2014). “Seemings: Still Dispositions to Believe”. *Synthese* 191.8, pp. 1761–1774 (cit. on pp. 62, 64).
- Whittlesea, B. W. A. (1993). “Illusions of Familiarity”. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19.6, pp. 1235–1253 (cit. on p. 144).
- (2002). “Two Routes to Remembering (and Another to Remembering Not).” *Journal of Experimental Psychology: General* 131.3, pp. 325–348 (cit. on p. 166).
- Whittlesea, B. W. A. and Williams, L. D. (2000). “The Source of Feelings of Familiarity: The Discrepancy-Attribution Hypothesis.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26.3, pp. 547–565 (cit. on pp. 125, 143, 166).
- (2001a). “The Discrepancy-Attribution Hypothesis: II. Expectation, Uncertainty, Surprise, and Feelings of Familiarity.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27.1, pp. 14–33 (cit. on p. 125).
- Whittlesea, B. W. and Leboe, J. P. (2003). “Two Fluency Heuristics (and How to Tell Them Apart)”. *Journal of Memory and Language* 49.1, pp. 62–79 (cit. on p. 142).
- Whittlesea, B. W. and Williams, L. D. (1998). “Why Do Strangers Feel Familiar, but Friends Don’t? A Discrepancy-Attribution Account of Feelings of Familiarity”. *Acta psychologica* 98.2-3, pp. 141–165 (cit. on pp. 124, 125).
- (2001b). “The Discrepancy-Attribution Hypothesis: I. The Heuristic Basis of Feelings and Familiarity.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27.1, p. 3 (cit. on p. 125).
- Wiese, W. (2018). *Experienced Wholeness: Integrating Insights from Gestalt Theory, Cognitive Neuroscience, and Predictive Processing*. MIT Press (cit. on p. 166).
- Williams, B. (1970). “Deciding to Believe”. *Problems of the Self*, pp. 136–51 (cit. on pp. 43, 218).
- Williamson, T. (2004). “Philosophical ‘Intuitions’ and Scepticism about Judgement”. *dialectica* 58.1, pp. 109–153 (cit. on p. 48).
- (2007). *The Philosophy of Philosophy*. Blackwell Pub (cit. on pp. 36, 48, 67, 210).
- Wilson, T. D., Lisle, D. J., Schooler, J. W., Hodges, S. D., Klaaren, K. J., and LaFleur, S. J. (1993). “Introspecting about Reasons Can Reduce Post-Choice Satisfaction”. *Personality and Social Psychology Bulletin* 19.3, pp. 331–339 (cit. on p. 196).
- Wilson, T. D. and Schooler, J. W. (1991). “Thinking Too Much: Introspection Can Reduce the Quality of Preferences and Decisions.” *Journal of personality and social psychology* 60.2, p. 181 (cit. on p. 196).

- Winkielman, P. and Berridge, K. C. (2004). “Unconscious Emotion”. *Current directions in psychological science* 13.3, pp. 120–123 (cit. on p. 93).
- Winkielman, P., Berridge, K. C., and Wilbarger, J. L. (2005). “Unconscious Affective Reactions to Masked Happy Versus Angry Faces Influence Consumption Behavior and Judgments of Value”. *Personality and Social Psychology Bulletin* 31.1, pp. 121–135 (cit. on p. 92).
- Winkielman, P., Huber, D. E., Kavanagh, L., and Schwarz, N. (2012). “Fluency of Consistency: When Thoughts Fit Nicely and Flow Smoothly”. *Cognitive Consistency: A Fundamental Principle in Social Cognition*, pp. 89–111 (cit. on p. 229).
- Winkielman, P., Schwarz, N., Fazendeiro, T., and Reber, R. (2003). “The Hedonic Marking of Processing Fluency: Implications for Evaluative Judgment”. *The Psychology of Evaluation: Affective Processes in Cognition and Emotion*, pp. 189–217 (cit. on pp. 152, 153, 207).
- Winkielman, P., Ziembowicz, M., and Nowak, A. (2015). “The Coherent and Fluent Mind: How Unified Consciousness Is Constructed from Cross-Modal Inputs via Integrated Processing Experiences”. *Frontiers in Psychology* 6.83 (cit. on pp. 159, 228).
- Wittgenstein, L. (2009). *Philosophical Investigations*. John Wiley & Sons (cit. on p. 224).
- Wringe, B. (2015). “The Contents of Perception and the Contents of Emotion: The Contents of Perception and the Contents of Emotion”. *Nous* 49.2, pp. 275–297 (cit. on p. 211).
- Yablo, S. (1993). “Is Conceivability a Guide to Possibility?” *Philosophy and Phenomenological Research* 53.1, pp. 1–42 (cit. on p. 211).
- Young, G. (2007a). “Clarifying “Familiarity”: Examining Differences in the Phenomenal Experiences of Patients Suffering From Prosopagnosia and Capgras Delusion”. *Philosophy, Psychiatry, & Psychology* 14.1, pp. 29–37 (cit. on p. 125).
- (2007b). “In Defense of Estrangement”. *Philosophy, Psychiatry, and Psychology* 14.1, pp. 51–56 (cit. on p. 125).
- (2009). “In What Sense ‘Familiar’? Examining Experiential Differences within Pathologies of Facial Recognition”. *Consciousness and Cognition* 18.3, pp. 628–638 (cit. on p. 125).
- Young, J. O. (2018). *The Coherence Theory of Truth*. In collab. with E. N. Zalta. URL: <https://plato.stanford.edu/archives/fall2018/entries/ruth-coherence/> (visited on 09/23/2019) (cit. on p. 241).
- Zahavi, D. and Kriegel, U. (2015). “For-Me-Ness: What It Is and What It Is Not”. *Philosophy of Mind and Phenomenology*, pp. 36–53 (cit. on p. 164).

- Zajonc, R. B. (1968). "Attitudinal Effects of Mere Exposure." *Journal of personality and social psychology* 9 (2p2), p. 1 (cit. on p. 207).
- Zanna, M. P. and Cooper, J. (1974). "Dissonance and the Pill: An Attribution Approach to Studying the Arousal Properties of Dissonance." *Journal of personality and social psychology* 29.5, p. 703 (cit. on p. 229).
- Zeelenberg, M., Nelissen, R. M., Breugelmans, S. M., and Pieters, R. (2008). "On Emotion Specificity in Decision Making: Why Feeling Is for Doing". *Judgment and Decision making* 3.1, p. 18 (cit. on p. 96).

RÉSUMÉ

La question directrice de la présente thèse est : « Quel type d'états sont des intuitions ? » Nous appelons la réponse développée ici Affectivisme Intellectuel. L'affectivisme soutient que les intuitions sont des expériences affectives, ou plus précisément : ce sont des cas spécifiques de sentiments épistémiques, de sentiments de vérité et de sentiments de fausseté. Dans un premier temps, l'état cible dont l'affectivisme est une théorie est distingué d'autres choses que nous appelons aussi « intuition ». Le profil caractéristique des « intuitions » au sens des expériences intuitives ciblées par la théorie est ensuite décrit : Les expériences d'intuition sont des états mentaux conscients qui sont (partiellement) caractérisés par leur phénoménologie caractéristique (mais pas nécessairement sui generis). Elles sont en outre 1) intentionnelles, 2) assertives, 3) motivationnelles, 4) non engageantes, 5) susceptibles de degrés dans 5.1) leur contenu et 5.2) leur insistance, 6) phénoménologiquement épistémiquement valuées (autrement dit, il existe des intuitions positives concernant la vérité et négatives concernant la fausseté) et 7) non volontaires. Nous soutenons qu'une bonne théorie des expériences intuitives doit rendre compte de ce profil de caractéristiques et ainsi qu'une bonne réponse à la question directrice devrait être capable de reconnaître et d'expliquer ces caractéristiques. Les théories de l'intuition existantes donnent les réponses suivantes : L'Eliminativisme soutient que le terme « intuition » n'a pas d'extension—les intuitions n'existent pas. Le doxasticisme soutient que les intuitions sont des états doxastiques. Le perceptualisme soutient que les intuitions sont semblables aux expériences perceptuelles. Nous démontrons que toutes les approches existantes sont insatisfaisantes. Soit elles ne sont pas en mesure de reconnaître les caractéristiques des intuitions, soit elles ne sont pas en mesure de les expliquer (soit elle ne peuvent ni l'un ni l'autre). Le reste de la thèse est consacré au développement d'une nouvelle théorie de l'intuition : l'Affectivisme Intellectuel. La réponse à la question directrice est la suivante : les intuitions sont des expériences affectives, ou plus précisément : il s'agit de cas spécifiques de sentiments épistémiques, de sentiments de vérité et de sentiments de fausseté. Pour motiver cette réponse, le type psychologique des sentiments ou expériences affectives est introduit et caractérisé : les expériences affectives, dont les sentiments corporels tels que la douleur ou le plaisir corporel et les sentiments émotionnels tels que la peur ou la joie sont des sous-classes paradigmatiques, sont valuées, excitatoires, motivationnelles et riches intentionnellement en s'engageant dans une division du travail représentationnel avec les autres états mentaux. Ensuite, la classe des sentiments épistémiques est introduite et caractérisée. La thèse défend ensuite l'idée que les sentiments épistémiques sont des expériences affectives. Une fois cela établi, elle identifie et analyse ensuite des sentiments épistémiques spécifiques comme candidats prometteurs pour une identification avec des expériences intuitives : sentiments de justesse et leurs contraires. Il s'avère qu'une variété propositionnelle de ces sentiments, les sentiments de vérité et les sentiments de fausseté, a le même profil caractéristique que les expériences intuitives. Nous soutenons ainsi que les intuitions positives doivent être identifiées aux sentiments de vérité et les intuitions négatives doivent être identifiées aux sentiments de fausseté. En vertu du fait que ces sentiments sont des expériences affectives, l'affectivisme non seulement reconnaît les caractéristiques des intuitions, mais les explique aussi. Les intuitions ont les caractéristiques qui sont les leurs pour essentiellement les mêmes raisons que les sentiments corporels et émotionnels les ont- parce qu'elles sont des expériences affectives (spécifiques).

MOTS CLÉS

intuitions ; sentiments épistémiques ; émotions ; affect ; phénoménologie

ABSTRACT

The guiding question of the present thesis is: “What kind of states are intuitions?” The answer developed here is Intellectual Affectivism or Affectivism (about intuitions). Affectivism claims that intuitions are affective experiences, or more precisely: they are specific instances of epistemic feelings, feelings of truth and feelings of falsity. First, the target state of which Affectivism is a theory is delineated from other things we call “intuition”. Then the feature profile of “intuitions” in the target sense of intuition experiences is outlined: Intuition experiences are occurrent conscious mental states that are (partially) characterised by their characteristic (but not necessarily sui generis) phenomenology. They are furthermore 1) intentional, 2) assertive, 3) motivational, 4) noncommittal, 5) gradable in 5.1) content and 5.2) pushiness, 6) phenomenally epistemically valenced (i.e. there are positive intuitions concerning truth and negative intuitions concerning falsity) and 7) nonvoluntary. It is argued that this feature profile needs to be accommodated by a good theory of intuition experiences, i.e. a good answer to the guiding question should be able to acknowledge and explain these features. Extant intuition theories provide the following answers: Eliminativism claims the term “intuition” has no extension — intuitions do not exist. Doxasticism claims intuitions are doxastic states. Perceptualism claims intuitions are similar to perceptual experiences. It is shown that all the existing answers are unsatisfactory. Either they cannot acknowledge the features of intuitions or they cannot explain them (or both). The rest of the thesis is dedicated to the development of a new intuition theory: Intellectual Affectivism . The answer it gives to the guiding question is the following: intuitions are affective experiences, or more precisely: they are specific instances of epistemic feelings, feelings of truth and feelings of falsity. To motivate this answer, the psychological kind of feelings or affective experiences is introduced and characterised: affective experiences, of which bodily feelings such as bodily pain or pleasure and emotional feelings such as fear or joy are paradigmatic subclasses, are valenced, arousing, motivational and richly intentional by engaging in a division of representational labour with other mental states. Then the class of epistemic feelings is introduced and characterised. The thesis proceeds to make a case for epistemic feelings being affective experiences. Having established that, it goes on to identify and analyse specific epistemic feelings as promising candidates for an identification with intuition experiences: feelings of rightness and feelings of wrongness. It turns out that a propositional variety of these feelings, feelings of truth and feelings of falsity, has the same feature profile as intuition experiences. Thus, the claim goes, positive intuitions are to be identified with feelings of truth and negative intuitions are to be identified with feelings of falsity. In virtue of these feelings being affective experiences, Affectivism cannot only acknowledge the features of intuitions but also explain them. Intuitions have the features they have for essentially the same reasons as bodily and emotional feelings have them —because they are (specific) affective experiences.

KEYWORDS

Intuitions; epistemic feelings; emotions; affect; phenomenology