



HAL
open science

Primitive introspection

Anna Giustina

► **To cite this version:**

Anna Giustina. Primitive introspection. Philosophy. Université Paris sciences et lettres, 2018. English. NNT : 2018PSLEE033 . tel-02133944

HAL Id: tel-02133944

<https://theses.hal.science/tel-02133944v1>

Submitted on 20 May 2019

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 l'École Normale Supérieure

Primitive Introspection

Soutenue par

Anna GIUSTINA

Le 17 décembre 2018

Ecole doctorale n° 540

**Transdisciplinaire
lettres/sciences**

Spécialité

Philosophie

Composition du jury :

François RÉCANATI Directeur d'études, ENS, IJN	<i>Président</i>
Susanna SCHELLENBERG Professeure, Rutgers University	<i>Rapporteur</i>
Charles SIEWERT Professeur, Rice University	<i>Rapporteur</i>
Dorothea DEBUS Senior Lecturer, University of York	<i>Examineur</i>
Uriah KRIEGEL Directeur de recherche, ENS, IJN	<i>Directeur de thèse</i>

TABLE OF CONTENTS

Acknowledgments	5
Résumé substantiel de la thèse en français	7
INTRODUCTION	13
Existence	14
Nature	15
Epistemology.....	15
PART 1: EXISTENCE.....	17
CHAPTER 1. What is primitive introspection?	19
1. The place of primitive introspection in the debate about self-knowledge	20
2. Primitive introspection.....	26
2.1. A first pass	26
2.2. Preliminary motivation for primitive introspection.....	28
3. Primitive introspection between Cartesianism and skepticism	30
Conclusion.....	34
CHAPTER 2. The phenomenal-concept acquisition argument for primitive introspection ...	35
1. Classification and concepts	35
2. Phenomenal concepts	42
3. The argument from phenomenal-concept acquisition	43
3.1. The argument: first step.....	44
3.2. Defending (P2)	45
3.3. Defending (P3)	61
4. Pure demonstratives, direct phenomenal concepts, and primitive introspection	63
Conclusion.....	64
PART 2: NATURE	65
CHAPTER 3. Metaphysical structure: preliminaries.....	67
1. Four conditions for primitive introspection.....	68
2. Act, target, and state of primitive introspection	71
3. The act of primitive introspection.....	74
3.1. Attention	74
3.2. Apprehension	80
4. The target of primitive introspective acts: preliminary considerations.....	82
4.1. Restriction on target	82
4.2. Temporality	83
5. The target of primitive introspective acts: ontology	85
5.1. The Phenomenal Property Instance View	86

5.2. The Particular View.....	88
5.3. The Phenomenal Event View	89
5.4. The Phenomenal State View	93
5.5. The Mereological View.....	93
Conclusion.....	96
CHAPTER 4. The integration account of the metaphysics of primitive introspection	97
1. Two desiderata for an account of the metaphysics of primitive introspection	98
1.1. The no introspective hallucination desideratum.....	98
1.2. The phenomenal modification desideratum	101
2. The inadequacy of the inner sense account	106
2.1. The inner sense view.....	106
2.2. Rejoinder attempt n. 1: mere nomological necessity?	108
2.3. Rejoinder attempt n. 2: can inner sense deliver metaphysical necessity?	110
3. The inadequacy of extant acquaintance accounts	112
4. A proposed solution: the integration account of primitive introspection.....	115
4.1. Preliminary assumptions	116
4.2. Integration.....	126
Conclusion.....	131
Appendix 1. Unshared essential properties	131
Appendix 2. Sketch for an alternative account of primitive introspection’s metaphysics	135
PART 3: EPISTEMOLOGY.....	139
CHAPTER 5. Primitive introspection as <i>sui generis</i> knowledge by acquaintance	141
1. Knowledge by acquaintance.....	142
2. Kinds of knowledge.....	148
3. Propositional knowledge does not exhaust knowledge by acquaintance.....	153
4. Challenging intuitions	158
5. Knowledge by acquaintance is <i>sui generis</i> : rebutting <i>disunity</i> and <i>mysteriousness</i> ..	164
5.1. The objection from disunity.....	164
5.2. The objection from mysteriousness	167
6. Why <i>sui generis</i> knowledge by acquaintance matters.....	169
Conclusion.....	170
CHAPTER 6. Knowledge by acquaintance is epistemically significant.....	173
1. Two marks of epistemic significance	174
1.1. Some epistemically significant phenomena	175
1.2. Marks of epistemic significance	178
2. Information acquisition	183
3. Information acquisition and epistemic significance.....	187
4. Knowledge by acquaintance involves information acquisition.....	190
4.1. Knowledge by acquaintance and information acquisition	190

4.2. The format of information in knowledge by acquaintance	192
4.3. Knowledge by acquaintance and amount of information	194
4.4. Complete and perfect grasp.....	201
5. Epistemic evaluability	205
6. Epistemic evaluability and epistemic significance.....	209
7. Epistemic evaluability and knowledge by acquaintance	212
7.1. Knowledge by acquaintance is epistemically evaluable	213
7.2. Knowledge by acquaintance is epistemically valuable	217
Conclusion.....	217
CONCLUSION.....	219
REFERENCES	227

ACKNOWLEDGMENTS

This work, or any of its proper parts, would have never seen the light if it were not for my supervisor, and friend, Uriah Kriegel, whose thorough intellectual guidance, across-the-board advice, and permanent support have constituted a fundamental contribution to the development of this dissertation's ideas and of myself as a philosopher and as a person. I could not have hoped for a better supervision for my PhD.

I would like to thank Dorothea Debus, François Récanati, Susanna Schellenberg, and Charles Siewert for accepting to participate in my dissertation defense: it is a pleasure and an honor to have them as members of my defense jury. Particular thanks to Susanna Schellenberg and Charles Siewert, who have been willing to write a report on the dissertation.

I am grateful to David Chalmers, Andrew Lee, John Morrison, Jim Pryor, and Charles Siewert for comments on earlier drafts of parts of this dissertation, and to Susanna Schellenberg, to whom I owe the label "primitive introspection" that gives this dissertation its title.

My own mind would have never been able to conceive the present work's ideas in isolation. Discussion with other philosophers, friends and colleagues, has been essential for me to articulate the details of my view more deeply, become aware of objections I had overlooked, improve my arguments, and discover potential consequences and further developments of my claims. For this I am deeply indebted to Alberto Barbieri, Romain Bourdoncle, Géraldine Carranante, David Chalmers, Bertille de Vlieger, Paul Egré, Luca Gasparri, François Kammerer, Andrew Lee, Jorge Morales, John Morrison, Michael Murez, Takuya Niikawa, Christopher Peacocke, Jim Pryor, François Récanati, Susanna Siegel, Charles Siewert, Enrico Terrone, and Alfredo Tomasetta. Some of these people have given inputs and chatted with me about primitive introspection several times.

Along the way, I have been lucky enough to meet and discuss with a huge number of scholars at various conferences and seminars around the world. I hope the following list does not leave out anyone of them: John Barnden, Alma Barner, Silvia Bianchi, David Bourget, Johannes Brandl, Julien Bugnon, Eli Chudnoff, Filippo Contesi, Tim Crane, Dorothea Debus, Arnauld Dewalque, Michele di Francesco, Jérôme Dokic, Jordi Fernandez, Pablo Fernandez Velasco, Katalin Farkas, Sascha Fink, Keith Frankish, Guillaume Fréchette, Marie Guillot, Louis Gularte, Matthew Heeney, Aaron Henry, Ben Holguin, Dan Hoek, Yarran Hominh, Liz Irvine, Ethan Jacobs, Arden Koehler, Stacy

Kohls, Ben Koons, Benjamin Kozuch, Geoff Lee, Yifan Li, Yaojun Lu, Pascal Ludwig, Philippe Lusson, Deborah Marber, Farid Masrouf, Olivier Massin, Manuel Martínez Merino, Angela Mendelovici, Shimon Mercer-Wood, Trenton Merricks, Matthias Michel, Cheryl Misak, Bernard Molyneux, Devin Morse, Daniel Munro, Bence Nanay, Michele Palmira, David Papineau, Theodore Parise, Adam Patterson, L.A. Paul, Giulia Piredda, Nicolas Porot, Melissa Rees, Adriana Renero, Kara Richardson, Andrew Richmond, Luke Roelofs, Sophie Roux, Léa Salje, Katia Samoiloova, Kranti Saran, Chris Scambler, Tobias Schlicht, Eric Schwitzgebel, Andrea Sereni, Denis Seron, Julia Smith, Maja Spener, Jan Swiderski, Hamid Taieb, Christine Tappolet, Lucas Thorpe, Emile Thalabard, Reinaldo Bernal Velazquez, Markus Werning, and Helen Zhao.

I am grateful to all the other PhD students and post-docs at the Institut Jean Nicod, for their affection and support, and for enriching discussions on various philosophical topics, particularly to Romain Bourdoncle, Géraldine Carranante, Luca Gasparri, Benjamin Icard, François Kammerer, Hye Young Kim, Sandra Lasry, Slawa Loev, Tricia Magalotti, Michael Murez, Takuya Niikawa, Joanna Patriarca, Andrés Soria Ruiz, Tristan Thommen, Enrico Terrone, and Hualin Xiao. Special thanks to Romain Bourdoncle and Camille Sallé for helping me with the French summary of the dissertation.

Finally, thanks to my friends and my family, particularly to my mother and my sister Isabella, for supporting me and staying by my side all this time. Thanks to Raffa, for his inexhaustible intellectual aliveness, for his calm and playful care, and for having been beside me all along independently of the physical distance.

RÉSUMÉ SUBSTANTIEL DE LA THÈSE EN FRANÇAIS

Chatouillez doucement votre coude intérieur avec une plume ou avec votre doigt. Soyez très attentif à la sensation qu'un tel chatouillement provoque en vous. Vous découvrirez peut-être que la sensation est un peu agaçante. Ou, peut-être, vous découvrirez qu'en effet, elle est assez agréable. Maintenant, essayez de laisser de côté les considérations sur le caractère éventuellement agréable ou non de la sensation. Essayez d'éviter de porter *tout* jugement sur la façon dont votre sensation actuelle pourrait être décrite. Gardez votre attention concentrée sur la sensation, tout en résistant à la tentation de la classer dans une catégorie familière. Tâche difficile, je le sais. Cependant, si vous réussissez, vous vous rendrez peut-être compte que cette activité de concentration de l'attention exempte de tout jugement ou classification vous fournit néanmoins des informations (peut-être même beaucoup d'informations) sur vos sensations. Vous apprendrez ainsi quelque chose sur votre sensation indépendamment de vos capacités de classification ou de reconnaissance.

Si vous avez accompli la tâche avec succès, vous avez accompli ce que j'appelle l'*introspection primitive*. Cette thèse est une étude de l'existence, de la nature et de l'épistémologie du phénomène introspectif que vous venez d'expérimenter.

L'introspection primitive est, en première approximation, un type non classificatoire d'introspection d'état phénoménal. Par « introspection d'état phénoménal » j'entends la méthode, propre à la perspective en première personne, par laquelle on peut acquérir la connaissance de la phénoménologie de sa propre expérience consciente présente. « Non classificatoire » signifie ici que l'introspection primitive n'implique aucune reconnaissance ou classification : pour introspecter primitivement, on ne doit pas reconnaître l'état phénoménal introspecté comme un exemple de type d'expérience précédemment rencontrée. L'introspection primitive est donc une méthode introspective, propre à la perspective en première personne, par laquelle on peut connaître la phénoménologie de sa propre expérience consciente présente sans la classer ou la reconnaître comme une instance d'expérience précédemment rencontrée.

Je défends principalement trois thèses sur l'introspection primitive. Premièrement, elle *existe* : il y a un phénomène mental qui a les caractéristiques que j'attribue à l'introspection primitive et un tel phénomène est un processus introspectif réel. Deuxièmement, sa *nature* est mieux

expliquée par une version de la théorie de *l'accointance* – ce que j'appelle la *théorie de l'intégration*. Troisièmement, elle a une *valeur épistémique* : elle permet au sujet d'acquérir une connaissance de ses états phénoménaux. Cette connaissance est un type de connaissance *sui generis* : la *connaissance par accointance*. La connaissance par accointance a une propriété épistémique spéciale qui est, on pourrait dire, analogue à *l'infailibilité*.

En conséquence, cette thèse est divisée en trois parties, concernant respectivement l'existence, la nature et l'épistémologie de l'introspection primitive. Chaque partie est composée de deux chapitres. J'en résume ici brièvement le contenu.

*

Existence

Au chapitre 1, j'explique ce qu'est l'introspection primitive. Je contextualise le phénomène en esquissant l'espace logique autour de la connaissance introspective de soi et en montrant la place qu'y occupe l'introspection primitive. En passant en revue certaines des cibles potentielles de la connaissance de soi, et certaines méthodes pour atteindre la connaissance de soi, je présente l'introspection primitive comme une méthode introspective dont les cibles sont des états phénoménaux, c'est-à-dire des états conscients avec phénoménologie. Je soutiens qu'il faut distinguer deux types d'introspection d'état phénoménal, l'une qui consiste à classer l'état phénoménal introspecté comme une instance d'un certain type d'expérience (par exemple « ceci [cette expérience que j'ai maintenant] est [une expérience de] douleur »), et une autre qui n'implique aucune classification. J'appelle le premier type d'introspection « introspection réflexive » et le seconde type « introspection primitive ». La notion d'introspection primitive est ensuite décrite à l'aide de quelques exemples et par contraste avec l'introspection réflexive. L'introspection primitive doit être distinguée de la simple conscience (c'est-à-dire, du simple fait d'*avoir* une expérience consciente) et de la simple attention à l'expérience consciente. De plus elle a une valeur avant tout épistémique. Une motivation préalable à l'existence de l'introspection primitive vient de l'introspection des états phénoménaux que l'on a pour la première fois : même si l'on ne peut pas classer un état phénoménal lorsqu'on l'a la première fois, on peut néanmoins l'introspecter (autrement dit, on peut l'introspecter même si on ne peut pas introspecter *qu'il est comme-ci-et-comme-ça*).

Au chapitre 2, je développe un argument de l'acquisition des concepts phénoménaux pour l'existence de l'introspection primitive. D'abord, je présuppose que la capacité de classer ou de reconnaître un état phénoménal comme une instance d'un certain type d'expérience déjà rencontrée (par exemple l'expérience de douleur) implique la possession du concept phénoménal qui est associé à ce type d'expérience (par exemple douleur). Je soutiens, en suite, que si toute introspection implique une classification, la plupart des concepts phénoménaux ne pourraient pas être acquis. Je conclus que, pour éviter un nativisme radical sur des concepts phénoménaux, nous devons accepter l'existence d'une introspection non classificatoire (c'est-à-dire l'existence de l'introspection primitive).

Nature

Au chapitre 3, j'explique plus en détail ce qu'est l'introspection primitive. Je la caractérise comme une appréhension attentive non classificatoire des états phénoménaux et je spécifie les notions d'attention et d'appréhension sur lesquelles repose cette caractérisation. J'analyse également le processus d'introspection primitive en y distinguant trois éléments fondamentaux : *l'acte*, la *cible* et *l'état* de l'introspection primitive. Le cœur du chapitre est consacré à *l'acte* et à la *cible* de l'introspection primitive. L'acte d'introspection primitive est caractérisé comme un acte d'appréhension attentive qui (i) modèle la structure centre-périphérie de l'expérience globale de l'individu de telle sorte que l'état phénoménal cible en devient l'aspect le plus saillant et (ii) permet au sujet de faire référence à et d'acquérir des informations sur la phénoménologie de l'état phénoménal cible. Différents points de vue concernant le statut ontologique de la cible de l'introspection primitive sont explorés. Ma théorie de la métaphysique de l'introspection primitive peut les accommoder tous.

Au chapitre 4, je me concentre sur la métaphysique de *l'état* d'introspection primitive, c'est-à-dire l'état mental dans lequel se trouve le sujet lorsqu'il introspecte primitivement. Plus précisément, je développe une théorie de la relation entre l'état d'introspection primitive et sa cible (c'est-à-dire, l'état phénoménal introspecté). Je soutiens qu'il y a au moins deux *desiderata* qu'une théorie de la nature de l'introspection primitive doit satisfaire, ce que j'appelle le *desideratum de l'absence d'hallucination introspective* et le *desideratum de la modification phénoménale*. Je montre qu'aucune des meilleures théories disponibles ne satisfait ces deux *desiderata*. La théorie du sens interne a peut-être les ressources pour expliquer la modification phénoménale, mais elle ne peut pas satisfaire le *desideratum* de l'absence

d'hallucination introspective. La théorie de l'accointance rend compte de l'absence d'hallucination introspective. Cependant, les versions existantes de la théorie de l'accointance n'expliquent pas la modification phénoménale. Je propose ma propre version de la théorie de l'accointance, ce que j'appelle la *théorie de l'intégration* de l'introspection primitive, comme une meilleure version de la théorie de l'accointance. Je montre que la théorie de l'intégration satisfait les deux *desiderata*. Cela est une bonne raison pour préférer une explication de la nature de l'introspection primitive en termes de la théorie de l'intégration plutôt qu'en termes des théories concurrentes.

Épistémologie

Au chapitre 5, je défends l'idée que la connaissance de soi directement fondée sur l'introspection primitive, c'est-à-dire la *connaissance par accointance*, est un type de connaissance *sui generis*. La connaissance par accointance est un type de connaissance qui est constituée par la relation d'accointance introspective qui se réalise quand, et en vertu du fait que, on introspecte primitivement son expérience. Je soutiens que la connaissance par accointance est une sorte de connaissance qui est irréductible à la connaissance propositionnelle – ou, d'ailleurs, à tout autre type de connaissance, comme le savoir-faire ou la connaissance d'un domaine. Je présente quelques exemples qui donnent à première vue des raisons de croire que la connaissance par accointance est un type de connaissance *sui generis*. Je montre ensuite qu'il est possible de répondre aux objections et que l'idée que la connaissance par accointance est un type de connaissance *sui generis* reste une option prometteuse, qu'il faut considérer.

Au chapitre 6, je soutiens que, même si elle n'est pas réductible à la connaissance propositionnelle, la connaissance par accointance n'en est pas moins *épistémologiquement significative*. Je propose qu'il y a deux grandes symptômes d'importance épistémique : d'un côté, le fait qu'elle permet l'acquisition d'information et, de l'autre côté, le fait de pouvoir être évalué du point de vue proprement épistémique. Je montre que la connaissance par accointance possède les deux marqueurs. Chemin faisant, je soutiens également que la connaissance par accointance introspective a une propriété épistémique spéciale, ce que j'appelle le *saisi complet et parfait* – une propriété épistémique qui, dans un sens, est analogue à celle qui a été traditionnellement attribuée à certaine connaissance propositionnelle introspective, à savoir *l'infailibilité*. Je conclus qu'il y a de bonnes raisons de penser que la connaissance par accointance a une importance épistémique.

La principale contribution de ma thèse au débat contemporain sur l'introspection est triple. Tout d'abord, dans la plupart des théories de l'accointance les plus récentes, la relation d'accointance implique nécessairement la formation d'un type particulier de concepts phénoménaux (des concepts phénoménaux qui sont en partie constitués par l'état phénoménal cible). Ce que je propose, au contraire, c'est l'existence d'une sorte d'état introspectif qui *précède* théoriquement la formation de tout concept phénoménal. Deuxièmement, je soutiens que même si l'on ne classe pas ce que l'on introspecte, on peut acquérir une connaissance complète de son état phénoménal (une sorte de connaissance qui est fondamentalement différente de la connaissance propositionnelle). Troisièmement, tout en reconnaissant les limites de l'introspection (mises en évidence par de la littérature critique récente), je revendique une part de vérité dans l'intuition cartésienne que certaines connaissances introspectives sont épistémologiquement spéciales. Tout au moins, la connaissance introspective que constitue l'introspection primitive a une propriété épistémique particulière, qui s'apparente en quelque sorte à l'infailibilité : elle donne au sujet un *saisi complet et parfait* de la phénoménologie de son expérience.

INTRODUCTION

Tickle your inner elbow softly with a feather, or with your finger. Attend very carefully to the sensation such a tickling provokes in you. You may find out that the sensation is a bit annoying. Or, perhaps, you may discover that it is pleasurable instead. Try to leave considerations about the sensation's agreeableness aside now. Indeed, try and avoid *any* judgment about how your present sensation could be described. Keep your attention focused on the sensation, while resisting the temptation to classify it under a familiar category. Hard task—I know. Yet, if you succeed, you may realize that such an unjudgmental, non-classificatory attentional activity provides you with some (perhaps quite a lot of) information about your sensation. You thereby come to know something about your sensation that cannot be captured by your classificatory or recognitional capacities.

If you have successfully accomplished the task, you have carried out what I call *primitive introspection*.¹ This dissertation is a study of the existence, nature and epistemology of the introspective phenomenon you have just instantiated.

Primitive introspection is, at a first approximation, non-classificatory phenomenal-state introspection. By 'phenomenal-state introspection' I mean the distinctively first-personal method through which one can get knowledge of the phenomenology of one's own current conscious experience. By 'non-classificatory' I mean not involving recognizing the introspected phenomenal state as an instance of any previously encountered experience type. Primitive introspection is thus a distinctive first-personal method through which you can acquire knowledge about the phenomenology of your current conscious experience without classifying or recognizing it as an instance of any previously encountered experience type.

I defend three main claims about primitive introspection. First, it *exists*: there is a mental phenomenon that has the features I attribute to primitive introspection and such a phenomenon is a full-fledged introspective process. Second, its *nature* is best accounted for by a version of the *acquaintance* theory—what I call the *integration account*. Third, it has a *distinct epistemic value*: it provides the subject with knowledge of their phenomenal states. Such knowledge is a *sui generis* kind of knowledge: knowledge by

¹ I am indebted to Susanna Schellenberg for suggesting the label 'primitive introspection' to me.

acquaintance. Knowledge by acquaintance displays a distinctive epistemic property which is somewhat analogous to *infallibility*.

Accordingly, this dissertation is divided into three parts, concerning, respectively, the existence, the nature, and the epistemology of primitive introspection. Each part is composed of two chapters. Here I briefly summarize their contents.

Existence

In *Chapter 1* I explain what primitive introspection is. I contextualize the phenomenon by outlining the introspective self-knowledge logical space and pointing at the place primitive introspection occupies in it. Upon surveying some potential targets of self-knowledge and some candidate methods to achieve self-knowledge, I present primitive introspection as an introspective method whose target are phenomenal states, i.e. conscious states with phenomenology. I argue that two kinds of phenomenal state introspection should be distinguished, one which involves classifying the introspected phenomenal state as an instance of a certain experience type (e.g. “this [experience I am now having] is [a] pain [experience]”), and one which does not. I call the former *reflective introspection* and the latter *primitive introspection*. The notion of primitive introspection is then elucidated by way of some examples and by contrast with reflective introspection. Primitive introspection is full-fledged introspection: it has an eminently epistemic significance and should be distinguished from both mere consciousness and mere attention to consciousness. Preliminary motivation for the existence of primitive introspection comes from introspection of phenomenal states one has for the first time: even if one cannot classify a phenomenal state the first time one has it, one can still introspect it (roughly, one can introspect *it* even if one cannot introspect *that it is thus-and-so*).

In *Chapter 2* I develop an argument from phenomenal-concept acquisition for the existence of primitive introspection. By assuming that the capacity to classify or recognize a phenomenal state as an instance of an experience type (e.g. pain experience) maps into one’s possession of the relevant phenomenal concept (e.g. PAIN), I argue that if all introspection involved classification, most phenomenal concepts could not be acquired. I conclude that, if we are to avoid radical nativism about phenomenal concepts, we must accept the existence of non-classificatory introspection (i.e. primitive introspection).

Nature

In *Chapter 3* I explain more thoroughly what primitive introspection is. I characterize it as non-classificatory *attentive apprehension* of phenomenal states and I elucidate the notions of attention and apprehension this characterization relies on. I also analyze the process of primitive introspection by distinguishing three basic elements in it: the *act*, the *target* and the *state* of primitive introspection. The bulk of the chapter is devoted to the *act* and the *target* of primitive introspection. The act of primitive introspection is characterized as an act of attentive apprehension which (i) shapes the center-periphery structure of one's overall experience in such a way that the target phenomenal state becomes the most prominent aspect in it and (ii) enables the subject to refer to and acquire information about the phenomenology of the target phenomenal state. Different views concerning the ontological status of the target of primitive introspection are explored. My account of the metaphysics of primitive introspection can accommodate them all.

In *Chapter 4* I focus on the metaphysics of the *state* of primitive introspection, i.e. the mental state the subject is in when they primitively introspect. More precisely, I provide an account of the relationship between the state and the target of primitive introspection. I argue that there are at least two desiderata that an account of the nature of primitive introspection should satisfy, which I call the *no introspective hallucination desideratum* and the *phenomenal modification desideratum*. I show that none of the best available theories satisfies both desiderata. The *inner sense* theory of introspection may have the resources to explain *phenomenal modification*, but it cannot satisfy the *no introspective hallucination* desideratum. The *acquaintance* theory does fit the *no introspective hallucination* desideratum. However, extant versions of the acquaintance theory do not explain *phenomenal modification*. I propose my own version of the acquaintance view, what I call the *integration account of primitive introspection*, as an improvement on extant acquaintance accounts. I show that the integration account satisfies both desiderata. I take this to be a compelling reason to prefer the integration account over its competitors as an account of the nature of primitive introspection.

Epistemology

In *Chapter 5* I argue that self-knowledge directly based on primitive introspection, that is, *knowledge by acquaintance*, is a *sui generis* kind of knowledge. Knowledge by acquaintance is knowledge which is constituted by the relation of introspective acquaintance one bears to one's experience in

virtue of primitively introspecting it. I argue that knowledge by acquaintance is a kind of knowledge which is irreducible to propositional knowledge—or, for that matter, to *any* other (putative) kind of knowledge, such as knowing-how or knowledge of a subject matter. I present some examples providing *prima facie* reason to believe that knowledge by acquaintance is a *sui generis* kind of knowledge. I then show that potentially threatening objections can be answered and that knowledge by acquaintance being a *sui generis* kind of knowledge remains a live option on the table.

In *Chapter 6* I argue that, even though it is not reducible to propositional knowledge, knowledge by acquaintance is nonetheless *epistemically significant*. I propose that there are two main marks of epistemic significance: information acquisition and epistemic evaluability. I show that knowledge by acquaintance displays both marks. Along the way I also argue that knowledge by introspective acquaintance has a special epistemic property, what I call *complete and perfect grasp*—an epistemic property that, in a sense, is analogous to an epistemic property that has traditionally been attributed to some introspective propositional knowledge, namely, *infallibility*. I conclude that there are good reasons for thinking that knowledge by acquaintance is epistemically significant.

*

My dissertation's main contribution to the current debate on introspection is threefold. First, on most current acquaintance accounts of phenomenal-state introspection, acquaintance necessarily involves the formation of (a special kind of) phenomenal concepts (concepts that are partly constituted by the target phenomenal state). What I argue for, instead, is the existence of a kind of introspective state which theoretically *precedes* the formation of any phenomenal concept. Secondly, I argue that even if one does not classify what one introspects, one can get full-fledged knowledge of one's phenomenal state (a kind of knowledge that is fundamentally different from propositional knowledge). Thirdly, while acknowledging the limits of introspection emphasized by recent critical literature, I vindicate a grain of truth in the Cartesian intuition that some introspective knowledge is epistemically special. At the very least, the introspective knowledge which is constituted by primitive introspection has a special epistemic property, which is in some respect akin to infallibility: it provides the subject with a *complete and perfect grasp* of the phenomenology of their experience.

PART 1: EXISTENCE

CHAPTER 1

WHAT IS PRIMITIVE INTROSPECTION?

The focus of this dissertation is an introspective phenomenon which I call *primitive introspection*.¹ Primitive introspection is non-classificatory introspection of phenomenal states. By ‘introspection of phenomenal states’ I mean the distinctively first-personal method through which one acquires knowledge of the phenomenology of one’s own current conscious experience. By ‘non-classificatory’ I mean not involving any classification or recognition of what is introspected as an instance of any previously encountered experience type.

In this chapter, I am going to introduce the phenomenon of primitive introspection and make a first pass at an elucidation of what it is. First, I contextualize primitive introspection in the more general debate about self-knowledge. I survey a number of potential *methods* to acquire self-knowledge, as well as different *targets* of self-knowledge. Primitive introspection is an introspective method and its targets are phenomenal states. Second, I explain more thoroughly what primitive introspection is by way of some examples and by contrasting it with a different kind of introspection of phenomenal states, what I call *reflective introspection*. I also provide some preliminary motivation for thinking that primitive introspection is a psychologically real phenomenon. Finally, I briefly elaborate on how the apparently anachronistic quasi-Cartesian idea that primitive introspection is epistemically special interacts with recent criticism about introspection’s reliability.

*

¹ As mentioned in fn. 1 of the Introduction, I owe the label ‘primitive introspection’ to Susanna Schellenberg. My original label for primitive introspection was ‘thing-introspection’ because of some analogies it bears to Fred Dretske’s (1993) notion of *thing-awareness* (Giustina and Kriegel 2017). The analogy with Dretske’s notion, though, tended to confuse some receivers of my view and somewhat prevented them from getting the core characterization of the introspective phenomenon which constitutes the object of my study. Rather than being elucidated by the analogy, the phenomenon was often conflated with Dretske’s thing-awareness, whereas, notwithstanding some important similarities, the two notions are significantly different. Whence the urge for a new, less theoretically-burdened label.

1. The place of primitive introspection in the debate about self-knowledge

Self-knowledge, as the name quite transparently suggests, is knowledge of *oneself*. There are several aspects of yourself you may know or come to know. Your physical appearance, the physiology of your organs and organ-systems, your size, your weight, your height are aspects of yourself, as well as your nationality, your date of birth, your DNA, and perhaps your genealogic tree (your ‘origin’, what you ‘belong to’). Your behavior, your posture, the way you relate to others, your statements, your declarations, in short, what you say and do, are aspects of yourself. There are also some more ‘inner’ aspects of yourself: your character traits, your personal preferences, your aptitudes; your beliefs, your desires, wishes and hopes, your intentions; your reasons for past actions, your grounds for a decision; your memories, your states of imagination; your feelings, your sensations, pains, and pleasures, your perceptual experiences, your emotions, your moods. All these aspects of yourself—and, arguably, many more—are potential targets of self-knowledge. Some of these aspects are perhaps more important or interesting than others. Coming to know that you want to have a child is, arguably, more important to you than coming to know that you believe you are wearing socks, or that you believe it is raining outside (Cassam 2015).

You may come to know (some of) these aspects of yourself in different ways. Here are some methods you may adopt: looking at yourself in the mirror, looking at the parts of your body that are visible to you, asking your friends, being told by your analyst, observing your behavior, making inferences, reasoning, remembering, *introspecting*. Sometimes, a certain method is available for getting knowledge about some kinds of target but not others. For instance, looking at yourself in the mirror provides you with some knowledge about your physical appearance, but, arguably, it typically does not provide you with knowledge about the way you feel.² Similarly, by introspection you may come to know that you have a burning pain in your stomach but not that you weight 2 kilos more than you did two months ago.³ At least sometimes, some aspects of yourself are such that you need to combine different methods to get knowledge of them. You may discover that you love Thai food partly by introspecting the taste experience you have when

² It may occur that by discovering a sad face in the mirror you realize that there might be a sadness component in your overall current emotional state which you have overlooked. However, in my opinion, although this may constitute part of your grounds for believing that you are sad, this is not yet sufficient for you to come to know that you are sad—some extra introspective, observational, or inferential work needs to be done.

³ Although by introspection you may come to know that you feel heavier than before.

you eat Thai and partly by realizing that you often drive all the way to the other side of town to get Thai food (behavior observation) and inferring from that that you must really like Thai food (inference or reasoning).⁴

Sometimes you may come to know some aspects of yourself by *direct inference*. If you believe that there is some soup left in the fridge then, plausibly, you can thereby know (or, at least, believe) that you believe that there is some soup left in the fridge (Dretske 1995, Byrne 2005). Sometimes, the process of coming to know a certain aspect of yourself coincides with the very process through which such aspect of yourself is *shaped*, or comes into existence (Moran 2001). You may come to know where you want to live throughout the next few years by *deciding* where you want to live throughout the next few years. The decision process may involve a combination of the aforementioned methods: imagining yourself in a certain place and *introspecting* how you would feel, listening to *friends'* suggestions (they sometimes know what is best for you better than you do), acquiring information about the work environment you would have in that place and make *inferences* to the estimated level of happiness that would yield in you, and so on. Sometimes, it may occur that you discover something about yourself by finding yourself *expressing* it (Bar-On 2004), as when you find yourself crying out "I'm sick of this job!" (whence you discover that you do not really like your current job), or when you catch yourself yawning during a talk (whence you find out that this talk is boring you to death).

You may have noticed that I have numbered introspection among the methods to acquire self-knowledge, as if I somewhat took for granted that introspection is a *sui generis* method or process (perhaps an inner-observation method, or a self-scanning internal process, or something like that) among others (behavior observation, inference, testimony, and so on). However, this should *not* be taken for granted. On the one hand, some philosophers reject the idea that introspection consists in a sort of inner observation or self-scanning process and try to reduce introspection to, or explain it in terms of, one of the other methods mentioned above. For example, Fred Dretske (1995) defends the idea that introspection just is *external-world* observation plus inference: you come to believe that you hear a sound by inferring such a belief from your hearing a sound, just like you come to believe that your tank is empty by seeing your fuel gauge pointing at 'E' and thereby inferring that the tank is empty. No 'inner observation' or 'self-scanning' process is involved.

⁴ The Thai food example is by Eric Schwitzgebel (2012: 31).

On the other hand, other philosophers maintain that there is no such thing as *the* introspective process (or *the* introspective method). What we call introspection is nothing but a *combination* of different methods. Eric Schwitzgebel, who is the eminent representative of this kind of stance, argues that no single introspective process can account for the acquisition of self-knowledge:

I doubt that we can draw sharp lines through this snarl, cleanly isolating some genuinely introspective process from related, adjoining, and overlapping processes. What we have, or seem to have, is a cognitive confluence of crazy spaghetti, with aspects of self-detection, self-shaping, self-fulfillment, spontaneous expression, priming and association, categorical assumptions, outward perception, memory, inference, hypothesis testing, bodily activity, and who only knows what else, all feeding into our judgments about current states of mind. (Schwitzgebel 2012: 41)

This somewhat overlaps with our above considerations about the plurality of methods one may adopt to achieve self-knowledge. From these considerations, however, Schwitzgebel concludes that *introspection* should not be described as one single process:

Introspection is not a single process but a plurality of processes. It's a plurality both *within* and *between* cases: most individual introspective judgments arise from a plurality of processes (that's the within-case claim), and the collection of processes issuing in introspective judgments differs from case to case (that's the between-case claim). Introspection is not the operation of a single cognitive mechanism or small collection of mechanisms. Introspective judgments arise from a shifting confluence of many processes, recruited opportunistically. (Schwitzgebel 2012: 29)

Schwitzgebel then seems to conceive introspection very *broadly*, as an activity involving a variety of different processes or methods, possibly directed at a variety of different targets.

Whether Schwitzgebel's *broad* or compound notion of introspection or some *narrower*, more specific notion corresponds to what we ultimately want to call 'introspection' is likely to turn out a merely verbal or terminological question, which, at any rate, I do not intend to address here.

Regardless of whether this is what ultimately deserves the name 'introspection', the notion of introspection I will work with here is a *narrow* one. For one thing, my notion of introspection imposes a restriction on targets. On this notion, only one's *own current conscious mental states and processes* are targets of introspection. Other aspects of oneself are *not* known through introspection—as well as, obviously, things *other than* oneself. Accordingly,

are excluded from the potential targets of introspection not only the non-mental aspects of oneself (such as one's look, one's height, or one's internal physiology), which virtually nobody would ever reasonably claim that could be *introspected*, but also some aspects that both some men of the street and some professional philosophers or psychologists *would* sometimes number among the targets of introspection. On my narrow notion, character traits, personal preferences, and aptitudes are not potential targets of introspection because they are *unconscious*. For the same reason, standing (and thereby unconscious) attitudes such as beliefs, desires, hopes, and intentions are not introspectible, nor are repressed (and thereby unconscious) emotions. *Past* conscious states are also excluded from introspection's targets on this view (arguably, knowledge about them is achieved, at least partly, through memory): to be introspected, a mental state must be *present*. What *can* be introspected, on this view, are one's current conscious states and processes. Among those are one's occurrent (and thereby conscious) attitudes such as beliefs, desires, hopes, and intentions; one's conscious memories and states of imagination; one's conscious feelings, sensations, perceptual experiences, emotions, and moods.

Introspection is thus here understood as a method through which one can get knowledge of, or at least form beliefs about, one's own current conscious mental states. Importantly, this method is *distinctively first-personal*. You may come to know that you are nervous now by being told by your partner. Even if what you thereby come to know is a current conscious state of yours, you have not come to know that through introspection. Introspection is a first-personal method in the sense that it is a method that can only be adopted by the subject, to acquire knowledge about the subject themselves: it cannot be used to acquire knowledge about others.

So, as conceived of here, *introspection* is a distinctively first-personal method through which we can get knowledge of, or form beliefs about, our own current conscious states. The focus of this dissertation is a yet *narrower* notion of introspection, what is sometimes called *phenomenal-state introspection*, or *introspection of phenomenal states*. Phenomenal-state introspection is introspection directed to the *phenomenology* of experience: it is the distinctively first-personal method through which one can get knowledge of, or form beliefs about, the *phenomenology* of one's current conscious experience. All and only the subject's current conscious mental states with phenomenology (i.e. those which have a subjective qualitative character, in virtue of which there is something *it is like* for the subject to have them) are eligible targets of phenomenal-state introspection. The unpleasant character of your pain, the greenish character of your visual

experience of the grass, the way you feel when you are angry, elated, or surprised are all potential targets of phenomenal-state introspection. The scope of phenomenal-state introspection partly depends on which conscious mental states have phenomenology. One example concerns cognitive states such as thoughts, judgments, and occurrent beliefs. Whether these states can be the target of phenomenal-state introspection depends, at least partly, on one's stance about the existence and nature of cognitive phenomenology. Here I remain neutral about this point. For those who believe that there is cognitive phenomenology, what I say will apply to cognitive states too.

Phenomenal-state introspection may not exhaust (what I have characterized above as) *introspection*—there may be other introspective methods that are not phenomenal-state introspection. In other words, phenomenal-state introspection might not be the only introspective method: there may be other distinctively first-personal methods we can use to get knowledge of our current conscious states. For instance, you may think that there are at least two (at least conceptually distinct) aspects of our conscious states that can be known through introspection. On the one hand there is the *phenomenology* of one's current phenomenally conscious states, that is, the eminently subjective 'feel' associated with one's experiences (say, the reddish character of your visual experience as of a ripe tomato). On the other hand, there is the propositional *content* of occurrent propositional attitudes (occurrent beliefs, desires, hopes, intentions, and so on). For instance, the content of your (now) occurrent belief that French revolution occurred in 1789. If so, we may be inclined to favor a view that distinguishes between introspection of the content of propositional attitudes, let us call it *content-introspection*, and introspection of the phenomenology of phenomenally conscious states, that is, *phenomenal-state introspection*. To be sure, this distinction does not exclude that some conscious mental states may fall both in the 'content' category and in the 'phenomenology' category. If, for instance, you believe that cognitive states have phenomenology, then occurrent beliefs may be the target of both content-introspection and phenomenal-state introspection. In this framework, their content is known through content-introspection; their phenomenology is known through phenomenal-state introspection. Similarly, if you think that perceptual experience has propositional content, then the relevant content is the target of content-introspection and the experience's phenomenology is the target of phenomenal-state introspection.

The process of content-introspection and the process of phenomenal-state introspection may come apart: they may be two distinct kinds of process. For instance, you may have a view such that, whereas phenomenal-state

introspection is a matter of self-scanning, content-introspection is a fundamentally inferential process (like the one involved in Dretske's notion of introspection mentioned above). Moreover, each kind of introspection (i.e. content-introspection and phenomenal-state introspection respectively) may involve more than one process. For instance, you may think that content-introspection is sometimes a matter of inference and reasoning (you infer that you believe that 'there is some soup left' is what you believe from the fact that you believe that there is some soup left), sometimes a matter of self-shaping (you come to know where you want to go by undertaking a decision about where you want to go). Arguably, these would be two different kinds of content-introspection.

Similar considerations may apply to other potential aspects of conscious states, such as *attitudes* (or *modes*). If attitude is a conscious aspect of (some of) our conscious states and is (at least conceptually) distinct from both content and phenomenology, then it might be that attitudes are known through a yet different variety of introspection—*attitude introspection*.

So, depending partly on one's assumptions about the structure of conscious states, it may or may not be the case that all introspection is phenomenal-state introspection. I will remain neutral about this issue. Regardless of whether there exist other varieties of introspection, my present focus is on *introspection of phenomenal states*, that is, the distinctively first-personal method through which one can get knowledge of, or form beliefs about, the *phenomenology* of one's current conscious experience.

I suggest that a distinction should be drawn between two kinds of phenomenal-state introspection. On the one hand, there is a kind of introspection which involves recognizing the introspected phenomenal state and classifying it as an instance of a certain experience type. I call this kind of phenomenal-state introspection *reflective introspection*. On the other hand, there is a kind of phenomenal-state introspection that does not involve recognition or classification of what is introspected. It merely involves attending to the phenomenal state and non-descriptively acquiring information about its phenomenology, where by 'non-descriptively' I mean without recognizing or classifying the introspected phenomenal state as an instance of a certain *experience type*. The latter is what I call *primitive introspection*. The difference between these two kinds of phenomenal-state introspection is spelled out more fully in the next section.

So, at a first approximation, primitive introspection is non-classificatory phenomenal-state introspection: it is a distinctively first personal method

through which one can non-descriptively acquire knowledge of the phenomenology of one's own current conscious experience.

2. Primitive introspection

2.1. A first pass

Suppose you are tasting papaya and considering the gustatory experience you have while doing that. You know what papaya tastes like, therefore you can recognize a papaya-taste experience and distinguish it from other types of gustatory experience. When you attend to your current experience, then, you can recognize that experience as a papaya-taste experience. You thereby introspect that this experience you are having is a papaya-taste experience. Call this case (1). Now, suppose instead that you have never tasted papaya before, and unexpectedly bump into a papaya-taste experience (perhaps you are blindfolded and someone administers a papaya slice to you). In this case—call it case (2)—when you attend to your experience, you surely do not introspect *that you are having a papaya-taste experience*. You have no idea what papaya tastes like, so you cannot recognize your current experience as a papaya-taste experience. However, you can certainly introspect the weird gustatory experience you are having: you can introspect *the experience* even if you do not introspect *that* it is an instance of any particular experience type.

There is thus an important difference between the introspective state you have in (1) and the one you have in (2). In (1), you have a kind of introspective state which involves your recognizing the introspected phenomenal state and classifying it under a known experience type. Case (2) features a kind of introspective state that does not involve any such classification: there you simply attend to your experience's phenomenology, and thereby introspect it, without recognizing it or classifying it as an instance of any previously encountered experience type.

You may object that, even though in case (2) you cannot classify the experience you introspect as a papaya-taste experience, you can nevertheless classify it as an instance of other experience types. For example, you can classify it as a *fruit-taste experience*, or as a *sweet-taste experience*, or, at the very least, as a *taste experience*. And indeed, it may well be that, in case (2), you do introspect that this experience you are having is a fruit-taste experience, or that it is a sweet-taste experience, or that it is a taste experience. But consider now a modification of case (2); call it (2*). Here, not only you have never tasted papaya before—you have never tasted any fruit. Accordingly, you cannot classify what you introspect as a *fruit-taste*

experience. Furthermore, you have never tasted sweet food before, so you cannot classify it as a *sweet-taste experience* either. In fact, you have never had *any* gustatory experience in your life before tasting this papaya, say because of a dysfunction in your taste buds from which you have just recovered. Thus, you cannot even classify this experience you are having as a *taste experience*. In (2*), then, not only you cannot introspect that your experience is a papaya-taste experience: there is *no experience type* such that you could classify your experience as an instance of it. Therefore, in (2*), you introspect your experience without classifying it.

I argue that the introspective state you have in (1) is importantly different from that which you have in (2*). Indeed, I suggest, they are two radically different kinds of introspective state. The first kind is what I call *reflective introspection* and the second kind is what I call *primitive introspection*.

One may object that, even in case (2*), you *could* classify your experience as an instance of *some* experience type. Even if you cannot classify it as a gustatory experience, you certainly *can* classify it, say, as a *non-visual* experience, or as a *non-auditory* experience (provided that you have had visual or auditory experiences before).⁵ However, even if you do possess such a classificatory ability, you do not need to *exercise* it when you introspect: although you *can* classify your current experience as, say, non-visual, you do not *need* to so classify it in order to be introspectively aware of it. In fact, there are innumerable ways we could classify the experiences we introspect (as *non-reddish*, as *non-pain*, as *non-itch*, as *non-hunger*, as *non-anger*, and so on and so forth), most of which we do not actually exercise when we introspect (even when we reflectively introspect). Primitive introspection occurs independently of the subject's exercising the relevant classificatory ability—although it is compatible with the subject's possessing the relevant ability. So, one may primitively introspect even in cases in which one *can* classify or recognize the experience. What is essential to primitive introspection is that the subject does *not exercise* such an ability.

So, the most fundamental difference between primitive introspection and reflective introspection is that the latter involves classifying what is introspected and recognize it as an instance of a previously encountered experience type, whereas the former does not involve any such classification. Primitive introspection involves introspectively attending to the

⁵ Thanks to an anonymous referee for *Erkenntnis* for this objection.

phenomenology of the experience without the exercise of any classificatory or recognitional ability.

2.2. Preliminary motivation for primitive introspection

The papaya example above might not have shaken your skepticism about the psychological reality of primitive introspection. That example, you might think, is too artificial, especially when it comes to imagining that you have never had any taste experience in your life. Does anything like what I call ‘primitive introspection’ ever feature in a human being’s *ordinary* life? Are there any more down-to-heart cases of primitive introspection, you may wonder, cases that you can find in your own—*real*—experience? In this subsection I point at three classes of more familiar cases of primitive introspection.

The first class includes cases of experiences one has for the first time. In such cases, you *do not* recognize or classify the introspected experience because you *cannot* so classify or recognize it: you cannot recognize the experience because you have never instantiated that experience type. So, take any experience at the first time you had it: if you introspected it, your introspecting was an instance of primitive introspection. The first time you have an orgasm, for example, the experience you have is completely new to you: there is no already-encountered experience type under which you could classify the experience you are having.⁶ Nonetheless, you can certainly introspect your orgasm experience, even on the very first time you have it. Similarly, the first time you take a ride on a rollercoaster you have such an idiosyncratic experience that you cannot classify it as anything you have ever experienced before (or so it seemed to me the first time I took a rollercoaster ride). Once it is over, perhaps, you come to realize how *exciting* or *amazing* it was—you want to take another ride right away. However, *at the time of the experience*, the way you feel is so new to you that you are unable to classify it—not even as *exciting* or *amazing*. Nonetheless, you can certainly introspectively attend to your experience.

⁶ You may object that one who has an orgasm for the first time could at least classify it as a *pleasure* experience. However, first, it does not sound obvious to me that the first time one has an orgasm one would unhesitatingly classify it as a pleasurable experience (or as an unpleasant experience, for that matter). Secondly, if that does sound obvious to you, you can still imagine a person who has never had a pleasure experience before. Once they finally have one, arguably, there will be no previously encountered experience type under which they could classify it.

The second class features cases in which you *can* classify the experience (you do possess the right classificatory and recognitional abilities because you have already instantiated that experience type before), but you *do not* classify it because you do not want to—that is, because you *decide* not to classify the experience while introspecting. This occurs in some *meditation* practices. Some types of meditation require that you endeavor to ‘observe’ the goings-on within your mind without judging or describing them; leave any thought aside and observe what happens in your mind ‘as it is’, so to speak, and not as you would *judge* it to be upon classifying it under a familiar category. If you succeed, you engage in what I call ‘primitive introspection’. Something similar may have occurred to you even if you are not really a spiritual-meditative kind of person. You may have tried, sometimes, to introspect a certain experience of yours without forming any judgment about what that experience is—that is, by making an effort not to classify it as an instance of a familiar experience type. Perhaps, you may want to do this for *aesthetic* reasons: to contemplate and enjoy your experience better by attending to it in a non-classificatory way.

In the third class we have cases where the subject not only *can* classify the experience but also *does* classify the experience. You will (legitimately) wonder: what makes these cases of *primitive* introspection? By definition, they seem to be cases of *reflective* (i.e. classificatory) introspection. Indeed, these cases do involve reflective introspection. Therefore, they are *not* cases of *mere* primitive introspection, for mere primitive introspection entails the absence of reflective introspection—it requires that the subject deploys *no* classificatory or recognitional ability at all. However, primitive introspection can *co-occur* with some reflective introspection and indeed ordinary-life introspection of phenomenal states is often accompanied by some classification. Even in these cases two kinds of introspection of phenomenal states can still be distinguished. For one thing, they can be distinguished *conceptually*: for any state of reflective introspection there is a possible state of primitive introspection the subject would have been in had they been unable to recognize and classify the introspected phenomenal state. More importantly, two kinds of phenomenal-state introspection can be distinguished *phenomenologically*. Arguably, our introspective capacities outstrip our classificatory capacities. By reflective introspection one can only apprehend the phenomenal aspects which one can recognize and classify. But what one can apprehend by introspection exceeds what one can recognize and classify. So, it seems that even in cases in which one does classify the introspected experience one must *also* primitively introspect the experience.

By primitive introspection, one can grasp the details of the phenomenology that are inaccessible to reflective introspection.

For all these reasons, I believe that there is *prima facie* motivation to think that primitive introspection is a psychologically real phenomenon. In Chapter 2 I will develop a sustained argument for the existence of primitive introspection. To be sure, I do not claim that it is the *only* introspective process. Also, I do not claim, here, that it is a privileged or fundamental form of introspection.⁷ However, I maintain that primitive introspection is a full-fledged introspective phenomenon which can occur independently of other forms of introspection (e.g. reflective introspection).

Before continuing, a couple of clarifications may be helpful. First, primitive introspection is different from mere awareness or consciousness: *primitively introspecting* a phenomenal state is different from simply *having* a phenomenal state. The difference consists, roughly, in that primitive introspection involves *attending to* a phenomenal state, whereas the mere having of a phenomenal state does not involve such an attending. The role of attention in primitive introspection will be explained much more thoroughly in Chapters 3 and 4.

Second, primitive introspection does not reduce to mere attention directed to consciousness. As it will become clearer in Chapter 3, what distinguishes primitive introspection from mere attention to consciousness is that primitive introspection additionally involves *apprehension* of the relevant phenomenal state. Apprehending the phenomenal state implies *acquiring information* about its phenomenology. As I will argue in Chapter 6, such information acquisition is what fundamentally makes primitive introspection epistemically significant.

3. Primitive introspection between Cartesianism and skepticism

Not only primitive introspection is epistemically significant: its epistemic significance is distinctive and special. As I will argue in Chapter 6, knowledge by acquaintance possesses a distinctive and special epistemic property which I call *complete and perfect grasp* of the target phenomenal state. Roughly, the idea is that, in primitive introspection, there is no misrepresentation of the target phenomenal state. First, it can never be the case that one is in a state of primitive introspection but no phenomenal state is present (that is, there cannot be *hallucinatory* primitive introspection)—this is what I mean by

⁷ Although I am inclined to think so and would like to defend this idea in the future.

‘perfect grasp’. Second, by primitively introspecting a certain phenomenal state one acquires *all the information* available about the phenomenology of that phenomenal state—this is what I mean by ‘complete grasp’.

The details of primitive introspection’s epistemic distinctiveness will come out in due time. For now, it is enough to note that there is an important sense in which a traditional idea about the epistemic specialness of introspection is echoed by my claims about complete and perfect grasp in primitive introspection.

According to an important philosophical tradition going back at least to René Descartes, the beliefs we form on the basis of introspection (i.e. *introspective beliefs*) enjoy a privileged epistemic status: they are *infallible*—having them entails their truth. On this view, if on the basis of introspection you come to believe that you have a burning pain in your stomach, then it must be the case that you have a burning pain in your stomach. Similarly, if you come to believe on the basis of introspection that you currently have a visual experience as of a ripe tomato, then, necessarily, you do have such an experience. (Nota bene: the claim is *not* that your introspectively coming to believe that you have an experience as of a ripe tomato entails that *there is a ripe tomato* before you: what is entailed by your belief is that you have an *experience* as of a ripe tomato.) Or if you introspectively come to believe that you are (consciously) angry at your partner, then, necessarily, you are (consciously) angry at your partner. This feature of introspection—its necessarily outputting infallible beliefs—makes it *epistemically special* and *epistemically privileged*. It is *special* in the sense that it has an epistemic property that no other belief-forming or knowledge-acquisition method has (that is, infallibility-conduciveness). It is *privileged* in the sense that the epistemic property which makes introspection special also makes it *better* than any other belief-forming or knowledge-acquisition method.⁸

This philosophical tradition, however, has been harshly criticized (see e.g. Armstrong 1963, Dennett 1988, Kornblith 1998) and it is now generally agreed that introspective beliefs are *not* infallible. We are, at least sometimes, wrong about our experiences. Perhaps you are not really *angry* at your partner, but rather *regretful*, or *disappointed*. So, perhaps, your introspectively formed belief about your current emotion is mistaken. Perhaps

⁸ Not only infallibility, but also other special epistemic properties have been attributed to introspective beliefs within this philosophical tradition. Among them are *incorrigibility* (nobody can show that an introspectively formed belief of yours is false), *indubitability* (you cannot doubt that an introspectively formed belief of yours is true), and *self-warrant* (necessarily, introspectively formed beliefs are justified).

this sensation of yours, which you are now introspecting, is not really pain, but just itch. Perhaps you are undecided—you cannot wrap your mind around *what it is* that you are introspecting. If so, you cannot even *form* an introspective judgment in the first place. This kind of cases are typically put forward to show that introspection does *not* always culminate in infallible beliefs. Accordingly, introspection should not be considered as a *privileged* belief-forming or knowledge-acquisition method.

Not only introspection is no longer believed to possess any of the special epistemic features traditionally attributed to it, but even its *reliability* as a knowledge-acquisition method has been recently called into question (Schwitzgebel 2008). Several arguments have been put forward to the effect that we are *often* mistaken or uncertain about significant aspects of our inner life and that it might well be the case that our beliefs about the objects we see in our environment are much more epistemically secure than those about our own mental states. If so, introspection not only is *not a privileged* method to investigate the mind: other methods may even be better than introspection.

That introspection does not possess the special epistemic properties that have traditionally been attributed to it has now become orthodoxy. Introspection is *not* infallible. Perhaps, it is not even *reliable*. Yet, my claims about primitive introspection's epistemic specialness seem to outright neglect the thoroughly critical literature that has shaped that orthodoxy. Why focus on such an outdated, anachronistic idea?

The critical stance about introspection has its undeniable merits: the Cartesian tradition has undoubtedly overestimated the powers of introspection. On the other hand, though, skepticism has gone too far to the other end of the spectrum, or so it seems to me. Sustained criticism of the reliability of introspection has overshadowed what I take to be a grain of truth in the Cartesian idea. For although introspection sometimes, or even often, goes wrong, there may be *some* introspective process that does not. If Schwitzgebel's pluralist idea about the nature of introspection (i.e. that introspection consists of a plurality of processes, rather than reducing to a specific cognitive process) is sound, then, plausibly, the epistemic significance of introspection will be a function of the epistemic contributions of the different cognitive processes that are recruited in the introspective inquiry. So, it may be that only *some* of these processes are responsible of introspection's fallibility, or even of its unreliability. This would leave the possibility open for other processes to be reliable, perhaps *perfectly* reliable.

My suggestion is that, if we make the right distinctions and appropriately delimit the scope of the epistemic specialness claim, it may turn

out that there *is*, in fact, an introspective phenomenon or process that has, if not the properties attributed to introspection by the Cartesian tradition, at least some properties in their vicinity. Such an outcome would make justice to the Cartesian intuition, that has been so pervasive until about fifty years ago, and, at the same time, take to heart the skeptical warning. The present study of primitive introspection is a first pass at this enterprise: it is, in part, an attempt to vindicate the gist of the Cartesian idea while acknowledging its limits.

To be sure, I am not alone in this enterprise. The strategy outlined in the previous paragraph, that is, delimiting the Cartesian idea's scope of application, has been adopted by a number of contemporary philosophers. David Chalmers (2003), for example, restricts the infallibility thesis to what he calls 'direct phenomenal beliefs'.⁹ Horgan and Kriegel (2007) also argue for a restricted infallibility thesis, which only applies to what they call 'SPPB phenomenal beliefs': beliefs about the phenomenology of conscious experience (i.e. *phenomenal*) that are "singular, present, phenomenal in mode of presentation, and bracketed" (Horgan and Kriegel 2007: 128).¹⁰ In a similar vein, Brie Gertler (2012) argues that introspection yields epistemically privileged beliefs when those beliefs are formed upon 'pruning the epistemic appearances'.¹¹

⁹ Roughly, direct phenomenal beliefs are beliefs that involve the deployment of special concepts, what Chalmers calls 'direct phenomenal concepts', which are *partly constituted* by the very experience to which they refer.

¹⁰ A belief's being *phenomenal in mode of presentation* means that the phenomenal property which is attributed to the experience is picked out in a "phenomenal way": "deploying [a phenomenal] mode of presentation is a matter of undergoing oneself (or at least imagining undergoing) an experience that actually *instantiates* [the relevant phenomenal property]" (*ibid.*: 126, italics original). *Bracketed* phenomenal beliefs are beliefs such that their phenomenal mode of presentation is free from any presupposition about the relations the phenomenal property they pick out bears to other phenomenal properties:

[A] *bracketing* mode of presentation of phenomenal character [...] suspends any such presuppositions, so that their truth or falsity does not affect the content of the specific belief that employs such a mode of presentation. This is a mode of presentation that brackets out all relational information about the experience and its phenomenal character, including how experiences of this sort are classified by other subjects, how they are classified by oneself on other occasions, what their typical causes are, etc. It focuses (so to speak) on how the experience appears to the subject at that moment. (*ibid.*: 128, italics original.)

¹¹ This occurs when the subject adopts a "scrupulously cautious attitude":

Adopting a more cautious doxastic attitude has the effect of restricting the epistemic appearances, by raising the bar as to the strength (and kinds) of evidence regarding *p* that will dispose one to believe that *p*. [...] By adopting a scrupulously cautious doxastic attitude towards one's own experiences, one seeks to prune the epistemic appearances to the point where those that remain are exclusively determined by how things seem phenomenally—that is, by the phenomenal reality. (Gertler 2012: 109-110, footnote omitted)

The present study of primitive introspection shares the spirit of these views and adopts a similar strategy. However, it differs from them in an important respect. What all these views have in common is that they seek for the adequately restricted kind of *belief* to which the Cartesian idea applies. My study, instead, focuses on an introspective phenomenon which is *pre-doxastic* because *pre-conceptual*: it precedes the formation of any belief because it precedes the formation of any *concept*. As I will argue in Chapter 2, primitive introspection is a kind of introspection which is *non-conceptual* and thereby *non-propositional*. Its epistemic specialness, then, does not depend on the epistemic specialness of any introspective *belief*.

Conclusion

In this chapter, I have introduced the phenomenon of primitive introspection by contextualizing it in the larger debate about self-knowledge. I have explained that primitive introspection is a method through which one may acquire self-knowledge of a restricted set of aspects of oneself: it is a distinctively first-personal method through which one can non-descriptively acquire knowledge of the phenomenology of one's own current conscious experience. To be sure, first, I do *not* claim that primitive introspection is the *only* method to acquire self-knowledge, and not even the *only introspective* method; second, I do not claim that all kinds of piece of self-knowledge can be acquired via primitive introspection (knowledge about the contents of one's occurrent propositional attitudes, for instance, cannot be acquired just by primitive introspection: only knowledge of the phenomenology of one's current conscious experience can).

I have tried to elucidate what primitive introspection is by way of some examples and provided some preliminary motivation for its psychological reality.

Finally, I have explained how the apparently anachronistic quasi-Cartesian idea that primitive introspection is epistemically special interacts with recent criticism about introspection's reliability. My hope is that, by restricting the scope of epistemic privilege to primitive introspection, the gist of the Cartesian idea can be vindicated, despite the widely acknowledged limits of introspection as a belief-forming or knowledge-acquisition process.

Adequately 'pruned' beliefs, on Gertler's view, are not infallible, but still epistemically privileged: they are more strongly justified than any other empirical beliefs.

CHAPTER 2

THE PHENOMENAL-CONCEPT ACQUISITION ARGUMENT FOR PRIMITIVE INTROSPECTION

In Chapter 1 I have introduced primitive introspection and characterized it as non-classificatory introspection of phenomenal states. I have also provided ordinary-life examples of primitive introspection and some preliminary motivation for thinking that primitive introspection is a psychologically real phenomenon. In this chapter, I develop a sustained argument for the existence of primitive introspection, what I call the *argument from phenomenal-concept acquisition*.

First, I suggest that, with some caveats, the classificatory/non-classificatory distinction maps onto the conceptual/nonconceptual distinction (§1). Accordingly, primitive introspection can be characterized as a kind of phenomenal-state introspection that does not depend on the deployment of any phenomenal concepts. After some set-up considerations concerning phenomenal concepts (§2), the bulk of the chapter is devoted to the argument from phenomenal-concept acquisition for the existence of primitive introspection (§3). Roughly, the idea is that if all introspection were conceptual, the acquisition of most phenomenal concept would be mysterious. Denying the existence of primitive introspection entails a version of nativism about phenomenal concepts that would strike many people as highly implausible. I conclude the chapter with a few remarks concerning the extent of the argument from phenomenal-concept acquisition (§4).

*

1. Classification and concepts

As I have explained in Chapter 1, the most fundamental difference between primitive introspection and reflective introspection is that, whereas the latter involves classifying what is introspected as an instance of a certain experience type, the former does not involve any such classification. Importantly, classification requires *recognition*: it implies recognizing that the introspected experience is of a certain known, previously encountered kind.

Merely possessing the ability to so classify what is introspected is not sufficient for reflective introspection: to reflectively introspect, one must *exercise* this ability. In contrast, primitive introspection is possible in the absence of any classificatory ability. To be sure, one may primitively introspect also in cases where one *does* possess the relevant classificatory ability. What is essential to primitive introspection is that the subject does *not exercise* that ability.

Arguably, the capacity to classify what is introspected as an instance of a previously encountered experience type is due to possession of the right *concepts*; more specifically, of the right *phenomenal concepts*. (Compare: the ability to classify something as a cat depends on the possession of the concept CAT.) For this reason, in what follows, the classificatory/non-classificatory distinction will map onto the conceptual/non-conceptual distinction, where a mental state is conceptual *iff* it depends on the deployment of some concept(s), nonconceptual *iff* it does not depend on the deployment of any concept.¹

A caveat is in order. There may be some phenomenal concepts whose deployment does *not* imply classification. Let *pure phenomenal demonstratives* be phenomenal concepts with no descriptive component: phenomenal demonstratives like THIS and THAT. These can be used to refer to an experience but, unlike descriptive demonstratives (such as THIS PAIN, or THAT EXPERIENCE), are not accompanied by any description. Pure phenomenal demonstratives, by definition, do not involve any classification of the kind outlined above. Accordingly, if pure phenomenal demonstratives exist—an issue I do not intend to take a stand on here—then their deployment may be involved in primitive introspection.²

At the present stage, a further exception may be made for what David Chalmers (2003) calls *direct phenomenal concepts*. These are phenomenal concepts that are formed *upon* attending to the experience itself, which (according to Chalmers) partly constitutes their content. Therefore, they are not possessed *prior* to introspecting the relevant experience. By definition, then, direct phenomenal concepts do not involve classifying the experience they are associated with as an instance of a previously encountered experience type. Again, I will not take a stand on whether there really are such

¹ By ‘deploying’ a concept I mean exercising it. ‘Applying’ a concept is a way to deploy it. When you imagine a unicorn, you may deploy the concept UNICORN, even though you do not apply it to anything. When you see a table, and recognize it *as* a table, you not only deploy the concept TABLE, but also apply it to what you see.

² I will say a bit more about this in §4.

concepts. *If* there are, then primitive introspection may involve the deployment of direct phenomenal concepts too.³

What is fundamental to the distinction between primitive and reflective introspection, then, is that the former, unlike the latter, does not involve the deployment of any phenomenal concept *already* possessed by the subject *prior* to introspecting. Accordingly, reflective introspection is a conceptual introspective state, whereas primitive introspection is a nonconceptual introspective state, where an introspective state is conceptual *iff* it depends on the deployment of some phenomenal concept(s) already possessed by the subject prior to introspecting, nonconceptual *iff* it does not depend on the deployment of any phenomenal concept possessed by the subject prior to introspecting. On this characterization, the conceptual/nonconceptual distinction maps exactly onto the classificatory/non-classificatory distinction.

Note that primitive introspection being nonconceptual, it cannot be a propositional attitude.⁴ If propositions are ‘made up’ of concepts, the fact that primitive introspection is nonconceptual entails that it is non-propositional. Reflective introspection, in contrast, *may* be propositional, namely, if the concepts deployed while being in a state of reflective introspection are combined in the way required to form a structured proposition. You may introspect your headache experience *as* a headache experience (and thereby deploy the concept HEADACHE EXPERIENCE) without introspecting *that* this experience is a headache experience. In this case, you are in a non-propositional state of reflective introspection. However, from the descriptive phenomenal demonstrative THIS EXPERIENCE and the phenomenal concept HEADACHE EXPERIENCE you may also form the propositional content ‘This experience is a headache experience’. Your state of reflective introspection would then be conceptual *and* propositional. For the purpose of this dissertation, distinguishing between conceptual but non-propositional and

³ Although deployment of direct phenomenal concepts is *prima facie* compatible with my notion of primitive introspection, in Chapters 5 and 6 I will argue that there are independent reasons to reject the idea that primitively introspecting *depends on* deploying direct phenomenal concepts. I will briefly come back to this issue in §4.

⁴ I assume that propositions are structured and have concepts as their constituents. On other views of propositions (Stalnakerian or Russellian, for example), the claim that primitive introspection is non-propositional may be false. As noted, what is most important for my distinction is that primitive introspection does not involve deployment of previously possessed phenomenal concepts.

both conceptual and propositional states of reflective introspection is not of crucial importance.⁵

(Incidentally, primitive introspection is *uncommunicable*. If the concepts you deploy when you are in a state of reflective introspection are directly expressible by public language expressions, those expressions can be used to communicate to others what you reflectively introspect.⁶ For instance, when you reflectively introspect that your current experience is a headache experience, you are in a position to communicate to other people that you have a headache. By contrast, merely primitively introspecting your experience does not put you in the condition to communicate what you are attending to. For when you merely primitively introspect, there is no public language expression at your disposal which refers to *all and only* the instances of the kind of experience you are having.⁷)

One may object that there is at least one concept which is already possessed by the subject prior to introspecting and is deployed while introspecting at all times: the concept EXPERIENCE. When you introspect a certain sensation, even if you do not recognize it as a sensation of any particular sort (say a pain sensation or a coldness sensation), you must at least recognize it as an *experience*. The concept EXPERIENCE is an introspection-relative maximally generic concept: it is the maximally generic concept that can be applied by introspection. The objector's claim is thus that one must apply at least this concept at any time one introspects. Now, for one thing, even if subjects who are capable of introspection typically *possess* the concept EXPERIENCE, that they necessarily *deploy* it when they introspect is

⁵ Although this *will* be important for some developments of the work grounded in the distinction between primitive introspection and reflective introspection.

⁶ The content of reflective introspection, however, may be uncommunicable if, for instance, the subject does not know of any public language expression which could express the concepts they deploy in their reflective introspecting.

⁷ There might be other ways to express what you primitively introspect. First, even though you do not possess the suitable phenomenal concept to associate with the kind of experience you primitively introspect, you may form a new one offhand and invent a term to express it. Alternatively, you might simply express what you primitively introspect by saying "This experience is *thus* (i.e. as it phenomenally appears to me)". However, although in both cases it is true of you that you can somehow *express* the experience you attend to, in neither case can you *communicate* what you primitively introspect to others. For in the first case the expression you use is a *private* one, whereas to communicate anything one must rather use *public* language expressions. In the second case the expression is public, but it is too generic to convey any information about the relevant experience (any experience may be said to be 'thus'). Thanks to François Recanati for drawing my attention to this issue.

arguable. Plausibly, you may well introspect a certain sensation you have in this moment without introspecting *that* it is an experience.

Moreover, subtler issues arise when we try to specify how the content of the concept EXPERIENCE is fixed. It is often thought that the most effective way to fix the content of the concept EXPERIENCE is by *ostension*. Pinch your thigh vigorously: can you feel *that*? *That* is an experience. However, if this is the way we fix the content of EXPERIENCE, then it cannot be the case that every act of introspective attention toward an experience features the deployment of that concept, on pain of circularity. Still, one may propose that the content of the concept EXPERIENCE be fixed by description. The question, of course, is: what is the description which fixes that content? Three options come to my mind—I consider each of them in turn.

First, since experience is sometimes conceived as non-physical, the content may be fixed by the description ‘non-physical’. Accordingly, introspecting something as an experience would entail classifying it as non-physical. However, for one thing, many philosophers do *not* conceive experience as non-physical. More importantly, it is hardly the case that people classify what they introspect *as non-physical* every time they introspect. Therefore, introspecting something as an experience cannot amount to introspecting something as non-physical, especially if every instance of introspection involves classifying what is introspected as an experience.

Second, the content may be fixed by the description ‘simple, intrinsic, ineffable, private, and infallibly known’, since experience has often been conceived that way. However, again, not all philosophers reflecting about experience do conceive it that way: some of them deny that experience has any of the properties listed in the above description. Those theorists would deny that introspecting something as an experience entails representing something as simple, intrinsic, ineffable, private, and infallibly known. Moreover, and more importantly, most people outside philosophy virtually never represent something *as simple, intrinsic, ineffable, private, and infallibly known*: they simply do not typically have this kind of thoughts. Therefore, classifying something as simple, intrinsic, ineffable, private, and infallibly known cannot be involved at any instance of introspection.

A claim about conscious experience that, different from the previous two, would probably be accepted by many philosophers is that conscious experience is what gives us the impression that there is an explanatory gap between facts concerning consciousness and physical facts. Conscious experience may thus be characterized as what gives rise to the ‘mystery’ of consciousness, where the ‘mystery’ of consciousness consists in the fact that

consciousness does not seem to be explainable in physical terms.⁸ Accordingly, it may be suggested that the content of the concept EXPERIENCE is fixed by the description “the property which gives rise to the mystery of consciousness”, or something along these lines.⁹ However, this description is only available to philosophers of mind and perhaps a few particularly reflective subjects. People typically do not think about the mystery of consciousness. Surely, nobody has a representation as of the property which gives rise to the mystery of consciousness at every time they introspect.

Of course, other descriptions may be explored to fix the content of the concept EXPERIENCE. However, my prediction is that none of those will support the idea that the concept EXPERIENCE is deployed at any time one introspects. For it seems that even though those descriptions are *available* to one, it is very implausible that one mobilizes any of them at any time one introspects.

Finally, the argument from phenomenal-concept acquisition in §3 will suggest a further reason to doubt that the concept EXPERIENCE is necessarily applied at any act of introspection. It is unclear how the concept EXPERIENCE could be acquired if not by first having nonconceptual introspective representations of experiences. Alternatively, one would have to maintain that the concept EXPERIENCE is somehow innate. This might be tenable but sounds quite implausible: a zombie with unconscious perceptual representations might perhaps possess perceptual concepts, but it does not seem that it could possess the concept of EXPERIENCE.^{10;11}

*

To recapitulate, reflective introspection requires the deployment of (some of) the phenomenal concept(s) that are associated with the introspected experience and are already possessed by the subject prior to introspecting: it

⁸ Cf. Kriegel (2009: 3), who suggests to fix the reference of the term ‘phenomenal consciousness’ by the description ‘the property F, such that, in the actual world, F is responsible for the mystery of consciousness.’ One way he proposes to precisify the description is the following: “phenomenal consciousness is the property F, such that, in the actual world, F causally produces (in the suitably reflective subject, say) the sense that the facts of consciousness cannot be deduced from physical facts.” (Kriegel 2009: 3-4).

⁹ Those reported in footnote 8 are examples of alternative formulations of the description.

¹⁰ Thank to Uriah Kriegel for suggesting this example to me.

¹¹ Some (e.g. Roskies 2008) maintain that in order for a perceptual concept to be acquired on the basis of a certain perceptual experience, the subject must be aware of the content of the experience. If this is the case, then the zombie cannot acquire any perceptual concept either.

involves recognizing the introspected experience as an instance of a certain experience type. Therefore, in order for you to reflectively introspect that your current conscious experience is F, you must possess the phenomenal concept F, and deploy it while introspecting. For example, in order for you to reflectively introspect that your current experience is a headache experience, you must possess the concept HEADACHE EXPERIENCE and apply it to the experience you are introspectively attending to. Primitive introspection, in contrast, is nonconceptual. When you primitively introspect your current conscious experience, you need not possess any phenomenal concept and indeed *no* previously possessed phenomenal concept is applied when you *merely* primitively introspect (modulo the above discussion of putative pure phenomenal demonstratives and direct phenomenal concepts). So, to *merely* primitively introspect your headache experience, not only do you not *need* to possess the concept HEADACHE EXPERIENCE, even if you do possess that concept you must not deploy it. Of course, it is very likely that you know headache experiences well enough to be *unable* to introspect one without immediately recognizing it and thereby applying the concept HEADACHE EXPERIENCE to it. If so, then every time you primitively introspect your headache experience you *also* reflectively introspect that it is thus-and-so—you cannot *merely* primitively introspect it. (However, the two kinds of introspection may be distinguished even in cases in which one's experience is immediately recognized. For one thing, they may be discerned *conceptually*: for each state of reflective introspection, there is a state of primitive introspection that the subject could be in, had they not possessed the relevant phenomenal concepts.)

As noted in Chapter 1, there is another *prima facie* consideration that supports the distinction between reflective and primitive introspection. When we introspect, we often seem to be able to distinguish phenomenal properties that are more fine-grained than the concepts at our disposal. But what can be grasped by reflective introspection is restricted to the phenomenal properties for which the subject has a concept. Therefore, when one reflectively introspects, one will often also primitively introspect, and by primitive introspection will grasp the details of the phenomenology of one's current experience that outstrip one's phenomenal concepts.

In the bulk of the chapter I present a sustained argument for the existence of primitive introspection: I defend the thesis that some introspective states do not involve classifying what is introspected—that is, introspecting does not necessarily require recognizing what is introspected as an instance of a certain experience type. I call it the *argument from*

phenomenal-concept acquisition. Before that, some set-up considerations are in order.

2. Phenomenal concepts

Phenomenal concepts, as conceived here, are concepts which are associated with the phenomenology of experience. They are personal-level mental representations that enable the subject to (i) distinguish the experience they are currently introspecting from other current or past experiences and (ii) recognize it as an instance of a certain experience type. To each kind of experience corresponds a distinct phenomenal concept: to pain experience corresponds the phenomenal concept PAIN, to phenomenal-red experience corresponds the phenomenal concept PHENOMENAL RED, to papaya-taste experience corresponds the phenomenal concept PAPAYA-TASTE EXPERIENCE, and so on. Phenomenal concepts may be more or less determinate: PHENOMENAL SCARLET is a determinate of the determinable PHENOMENAL RED, which is a determinate of COLOR EXPERIENCE, which is a determinate of VISUAL EXPERIENCE.

Some phenomenal concepts are built up compositionally from other phenomenal concepts.¹² The phenomenal concept EXCRUCIATING PAIN, for instance, is built up by composition from EXCRUCIATING and PAIN. Of these, PAIN is associated with a *qualitative* phenomenal property, where qualitative phenomenal properties are phenomenal properties which define the qualitative aspect that characterizes each kind of phenomenal state: they are that which constitutes the difference between, say, phenomenal-red experience and phenomenal-blue experience, or between papaya-taste experience and stubbing-pain experience.¹³ On the other hand, EXCRUCIATING is associated with *intensity*, which is a *quantitative* phenomenal property.¹⁴ Thus, the phenomenal concept EXCRUCIATING PAIN can be acquired compositionally by putting together a phenomenal concept for a qualitative phenomenal property and a phenomenal concept of a quantitative phenomenal property.

¹² Provided that we assume a non-atomistic view of concepts. On a non-atomistic view, concepts can be built up compositionally from other concepts: BACHELOR, for instance, may be built up compositionally from UNMARRIED and MALE. On an atomistic view, instead, no concepts are composite.

¹³ Arguably, two qualitatively different phenomenal states (phenomenal states with different qualitative phenomenal properties) are phenomenal states of different kinds.

¹⁴ I will say more about qualitative and quantitative phenomenal properties in Chapter 4.

More generally, composite phenomenal concepts may be acquired compositionally from more component phenomenal concepts. Obviously, however, not all phenomenal concepts can be acquired in this way, on pain of infinite regress. There must be a layer of phenomenal concepts that are either not acquired at all or acquired non-compositionally.

Basic phenomenal concepts are phenomenal concepts which are *not* formed by composition from other concepts. They constitute the foundational layer out of which non-basic phenomenal concepts may be compositionally formed.¹⁵ Arguably, all phenomenal concepts associated with qualitative phenomenal properties are basic. Among these are concepts associated with perceptual experiences (PHENOMENAL RED, SOFTNESS EXPERIENCE, PAPAYA-TASTE EXPERIENCE, OLFATORY EXPERIENCE, etc.), algedonic experiences (PLEASURE, PAIN, SORENESS, THROBBING, etc.), bodily sensations (ITCH, TICKLE, ORGASM, etc.), interoceptive sensations (HUNGER EXPERIENCE, FAST-HEARTBEAT EXPERIENCE, etc.), emotions (ANGER, FEAR, etc.), moods (EUPHORIA, DEPRESSION, etc.). It is plausible to think that phenomenal concepts of this kind constitute the foundational layer, for it is hard to find more basic phenomenal out of which they could be built up. However, this plausible assumption is not quite necessary for the argument of this chapter. All that is necessary is that there be *some* foundational layer of non-compositionally acquired phenomenal concepts.

Acquiring a phenomenal concept means coming to possess it. One *possesses* a phenomenal concept iff one has it available to be deployed in cognition and action-guidance. *Deploying* a concept means using or exercising it, where exercising a concept is intended here as a *personal-level* process. Personal-level exercise of a concept involves active manipulation of the concept. Concept deployment is personal-level exercise of the concept in this sense. A phenomenal concept is *innate* just in case it is possessed but not acquired.

3. The argument from phenomenal-concept acquisition

Perhaps the most compelling reason to maintain that there must be a kind of introspective state which is nonconceptual (in the sense specified above) is that the acquisition of phenomenal concepts would be impossible otherwise. Roughly, the idea is the following. It is natural to think that we acquire phenomenal concepts by introspection. Now, if one's introspective state is conceptual, then one can only have it if one already possesses the concept(s)

¹⁵ On the atomistic view, *all* phenomenal concepts are basic.

associated with what is introspected. If *all* introspective states are conceptual, then the possession of all phenomenal concepts must be prior to the possibility of introspecting. But if this is the case, then phenomenal-concept acquisition is mysterious. The full argument is much more complex, however, and must be developed more slowly. This is the purpose of the present section.

The argument just sketched partly retraces the argument from *perceptual*-concept acquisition for nonconceptualism about perceptual experience (Heck 2000; Peacocke 2001). A variant of the latter has been put forward by Adina Roskies (2008), who argues that conceptualism about perceptual content must be rejected because it entails an unacceptable version of nativism about perceptual concepts (the thesis that perceptual concepts are innate, i.e. are possessed without being acquired). Although it differs from hers in many respects, the argument I propose is partly inspired by Roskies'.

3.1. *The argument: first step*

Assuming that non-basic phenomenal concepts are acquired by composition from more basic phenomenal concepts, the focus of the present argument is the acquisition of *basic* phenomenal concepts, i.e. the phenomenal concepts which constitute the foundational layer out of which composite phenomenal concepts are formed. The general structure of the argument from phenomenal-concept acquisition is as follows:

- (P1) (Almost) all basic phenomenal concepts are acquired.
- (P2) For most basic phenomenal concepts, if they are acquired, they are acquired by introspection.
- (P3) If all introspective states are conceptual, then it is not the case that most basic phenomenal concepts are acquired by introspection.
- (C) Not all introspective states are conceptual.

A note on the quantifiers. What do I mean by '(almost) all'? Consider this toy model: if 'most' is defined as 'more than 50%', then '(almost) all' should be read as 'equal or very close to 100%'. That would allow for a few phenomenal concepts being innate consistently with (P1) being true. I use these unusual quantifiers to protect the argument from certain eventualities inessential to its core idea.

Thus, denying (P1), as framed here, would entail a version of nativism about phenomenal concepts that would strike many of us as highly implausible. For if, as it seems to be, the foundational layer of basic phenomenal concepts includes very many concepts associated with

qualitative phenomenal properties, such a foundational layer is quite large—it includes a great number of phenomenal concepts. Denying (P1) would thus entail that a great number of phenomenal concepts are innate. To be sure, plausible versions of nativism are defended in the philosophical and cognitive-science literature (e.g. Susan Carey’s (2009) core cognition theory). However, none of them implies that *a great number* of phenomenal concepts are possessed from birth. Indeed, all they imply is that *a few* conceptual representations are innate.¹⁶ That a newborn possesses a great number of phenomenal concepts would mean that they possess representations that enable them to (i) distinguish a great number of experiences from one another and (ii) recognize each of them as an instance of a certain experience type. I find this very implausible. Those who agree with me on this will accept that, even if a few phenomenal concepts may be innate, almost all of them must be acquired. At the very least, we may think of the central thesis of this chapter as the following conditional: if radical nativism about phenomenal concepts (i.e. the theory according to which most basic phenomenal concepts are innate) is false, then not all introspective states are conceptual.

The next two subsections are aimed at defending (P2) and (P3). In §3.2 I argue that for most basic phenomenal concepts, if C (i.e. the phenomenal concept associated with the phenomenal property C) is an *acquired* phenomenal concept, then C is acquired by introspection of C-experiences. In §3.3 I argue that if C is to be acquired by introspection of C-experiences, then the relevant introspection of C-experiences must be nonconceptual.

3.2. Defending (P2)

In this subsection I argue that, even if some basic phenomenal concepts C could be acquired on the basis of something other than introspection of C-experiences, there are reasons to think that it cannot be the case that *most* of them are. Presumably, if a basic phenomenal concept C is acquired, but not through the introspection of C-experiences, then C is acquired either (i) through the introspection of *other* kinds of experience, or (ii) on the basis of something other than introspection. In the remainder of this subsection, I argue that even if (i) and (ii) were live options for *some* basic phenomenal

¹⁶ The innate conceptual representations posited by core cognition are abstract in kind and limited in number. They belong to three domains: objects (including representations of causal and spatial relations), numbers, and agents (including representations of goals). Even though the latter domain may include some folk psychological concepts (goal, attentional state, etc.), core concepts surely do not include a great number of phenomenal concepts. Indeed, they only include a few (plausibly not more than 1%).

concepts, they could not be for *most* of them. First, I argue that there are reasons to discard (i); at the very least, if (i) could ever be an option, it would only be for a very few phenomenal concepts (§3.2.1). Then, I consider two potential ways to defend option (ii). The first consists in adopting a sort of ‘transparency’ account of phenomenal-concept acquisition. I argue that this option, if viable at all, could only account for the acquisition of a restricted number of phenomenal concepts: the acquisition of the majority of our phenomenal concepts would still be in need of explanation (§3.2.2). The second one is grounded in the idea that phenomenal concepts are acquired not by introspecting an experience, but by the mere having of an experience. I argue that the mere having of an experience is not yet sufficient for one to form a phenomenal concept (§3.2.3).

3.2.1. Extrapolation. It may seem that (i) can be straightforwardly ruled out. If *c* is acquired through introspection, then it is much more plausible that it is acquired through introspection of *C*-experiences than that it is acquired through introspection of non-*C*-experiences. Plausibly, if the concept PAIN EXPERIENCE is acquired through introspection, it is acquired through introspection of pain experiences, rather than, say, coldness experiences. However, one might resist this line of argument by appeal to a Hume-inspired ‘missing shade of phenomenal blue’ objection.¹⁷ It might be argued that, even though one may acquire the phenomenal concept PHENOMENAL-BLUE₇ by introspecting bluish₇ experiences, one may also acquire it by introspecting bluish₆ experiences and bluish₈ experiences plus extrapolation.

My reply to this objection is twofold. *First*, it is not straightforward that PHENOMENAL-BLUE₇ can be acquired by extrapolation. To be sure, what *can* be acquired by extrapolation is the description <the phenomenal color between phenomenal-blue₆ and phenomenal-blue₈ in the phenomenal-color spectrum>. However, that this description amounts to the phenomenal concept PHENOMENAL-BLUE₇ is far from evident. For one thing, I doubt that such a description *alone* (that is, independently of the subject’s having ever introspectively attended to a bluish₇ experience) would enable one to distinguish and recognize a bluish₇ experience upon encounter. Moreover, once we consider other phenomenologies (phenomenal properties other than phenomenal colors), the extrapolation option sounds even more implausible. Take bodily sensations, such as itches or orgasms. An account such that the phenomenal concepts associated with those experiences are acquired via introspection of other experiences plus extrapolation seems hardly defensible.

¹⁷ Thanks to Andrew Lee and Emile Thalabard for drawing my attention to this objection.

From introspection of what experiences could, say, the concept ITCH (or ITCH EXPERIENCE) be acquired by extrapolation?

Second, even if phenomenal-concept acquisition by extrapolation were possible, it surely could not be the case that all or even most basic phenomenal concepts are acquired by extrapolation.¹⁸ Even if PHENOMENAL-BLUE₇ could be acquired by extrapolation from PHENOMENAL-BLUE₆ and PHENOMENAL-BLUE₈, the latter would have to be acquired by introspection of bluish₆ and bluish₈ experiences respectively (if they were acquired by introspection at all). So even if *some* phenomenal concepts could be acquired via extrapolation, only a very *few* of them could be acquired that way—*most* phenomenal concepts surely cannot.

3.2.2. *Transparency.* (ii) says that if a basic phenomenal concept C is acquired, it is acquired on the basis of something other than introspection. It seems to me that the most plausible way to defend (ii) is to adopt a sort of *transparency* theory of phenomenal-concept acquisition and maintain that basic phenomenal concepts are acquired on the basis of *perceptual* experience. On this view, basic phenomenal concepts are acquired not by attending to one's experience at all (thus not on the basis of introspection), but rather by attending to external objects and their properties. Phenomenal-concept acquisition therefore does not require introspective attention but merely depends on some sort of *perceptual* process: it merely depends on the subject's perceptually representing (and perceptually attending to) their environment and their body. Plausibly, a transparency theorist of this sort would argue that by having a certain perceptual experience, and by attending to the object or property the experience represents, one acquires a certain perceptual concept and it is on the basis of the latter that the relevant basic phenomenal concept is formed. So, for instance, the phenomenal concept PHENOMENAL RED is acquired on the basis of the perceptual concept RED, which in turn is acquired by the subject's having perceptual experiences as of (or perceptually attending to) red objects.

As long as phenomenal concepts associated with color experiences are concerned, the transparency account just outlined may sound, at least *prima facie*, plausible. However, even for the acquisition of phenomenal-color concepts a transparency account might be resisted. For the order of acquisition between, say, PHENOMENAL RED and RED is a traditionally highly

¹⁸ Nota bene: extrapolation is different from composition. That PHENOMENAL-BLUE₇ is acquired by extrapolation from PHENOMENAL-BLUE₆ and PHENOMENAL-BLUE₈ does not imply that PHENOMENAL-BLUE₇ is built up by composition from PHENOMENAL-BLUE₆ and PHENOMENAL-BLUE₈.

controversial issue. The notion that PHENOMENAL RED is derived from RED, rather than the other way around, is not totally obvious. According to a long-standing tradition in philosophy, going back to Galileo, Descartes and Locke, colors are *secondary qualities* of objects. Unlike qualities such as shape, size, motion, and number (i.e. *primary qualities*), which are objective and defined independently of anyone's experience, *secondary qualities* are defined in terms of subjective responses and do not feature in the scientific explanation of reality. There are at least two contemporary versions of this traditional view. According to *dispositionalism*, colors are powers or dispositions to elicit in perceivers visual experiences of a certain type (see e.g. Johnston 1992, Peacocke 1984, and Levin 2000).¹⁹ A poppy is red in virtue of its having the power to cause reddish experiences in normal subjects under normal conditions. According to *projectivism*, colors are not properties of external objects at all: strictly speaking, they are properties of visual fields. In color experience, properties of the visual field are erroneously represented as belonging to external objects—they are 'projected' to external objects (see Boghossian and Velleman 1989). On this view, a poppy is said to be red in virtue of the fact that when a subject sees the poppy, they erroneously attribute to it a property (the property of being red) which in fact belongs to a region of their visual field. On both versions of the secondary quality view, color properties are analyzed in terms of phenomenal-color properties. Color experience being more basic than color on this view, it would be natural to accompany it with an account of color concepts according to which color concepts are grounded in phenomenal-color concepts. Thus, if RED refers to the property of being disposed to elicit phenomenally red experiences under normal conditions, then, arguably, one will need to possess the concept PHENOMENAL RED prior to acquiring RED: one could not grasp the concept that refers to the external object's property without grasping the concept that refers to the experience caused by that property. Now, it is not my intention here to *prove* that phenomenal-color concepts are prior to color concepts. I am merely trying to highlight one substantial challenge our objector faces.

Regardless of how the particular issue concerning phenomenal-color concepts is to be settled, there is a more general and more serious challenge the transparency theorist of phenomenal-concept acquisition faces. For even

¹⁹ Accordingly, color experience consists in representing objects as having the disposition to cause a certain type of experience (under standard conditions). For instance, an experience as of a red poppy consists in representing the poppy as having the disposition to cause a reddish experience (under standard conditions). Clearly, a *prima facie* worry for dispositionalist accounts of color experience is *circularity* (cf. Boghossian and Velleman 1989). See Peacocke (1984) for a reply to this objection.

if the transparency account is viable for the acquisition of basic phenomenal concepts associated with color experiences (and perhaps also of basic phenomenal concepts associated with some other perceptual experiences), it seems deeply problematic when it comes to other kinds of experiences, such as algedonic experiences (pain and pleasure), bodily sensations (such as tickles and itches), emotions, moods, and imagination. Indeed, the transparency theorist of phenomenal-concept acquisition is committed to four increasingly controversial claims concerning those experiences:

- [1] All the experiences just listed are intentional.
- [2] For each of these experiences, its phenomenal character is exhausted by its intentional content.
- [3] Introspection of any of these experiences involves attending not to the experience itself, but to its intentional object.
- [4] The phenomenal concepts associated with these experiences are acquired by attending not to the experience itself, but by attending to its intentional object.

Claim [1] expresses a commitment to *intentionalism* about conscious experience, that is, the view that every conscious experience is *directed to* an object. As, when I think about David Bowie, my thought is intentionally directed to David Bowie, any conscious experience, according to the intentionalist, is directed to some object. Since intentional states may be intentionally directed toward objects that do not exist (like the child's belief that Santa Claus will bring her a bicycle this year), intentionalists introduce the notion of intentional *content*, which is often spelled out in terms of *representation*. In the case of belief, which is usually considered as a paradigmatic intentional state, the proposition believed is its intentional content. If I believe that David Bowie is a great artist, the intentional content of my belief is the proposition 'David Bowie is a great artist'. Such a content is representational: it represents the world as being a certain way, and has accuracy conditions (my belief may be true or false depending, for example, on whether David Bowie is really a great artist). So, for instance, your visual experience as of a yellow banana has a yellow banana as its intentional content: it represents a yellow banana and is accurate just in case there is a yellow banana before you (and your visual experience is appropriately related to the banana). According to the intentionalist, *all conscious experiences* are intentional in this sense: they have an intentional content.

Claim [2] expresses a particular version of intentionalism: what is usually called *strong representationalism*. Strong representationalism is the view that the phenomenal character reduces to, or is exhausted by, the intentional content of the experience. Consider again your visual experience

as of a yellow banana. According to the strong representationalist, the phenomenology of your visual experience is fully determined by its representational content of, i.e. the yellow banana. The main tenet of strong representationalism is that the phenomenology of *any* experience is fully determined by its intentional content.

Claim [3] is a direct consequence of what is usually called the *Transparency of Experience* thesis. The idea is that conscious experiences are, so to speak, ‘transparent’: when you try to focus your attention on your experience, all you can do is focusing (or keeping focusing) your attention on what the experience is *about*, what it is directed to (i.e. its intentional content). The transparency thesis was famously defended by Gilbert Harman (1990) and Michael Tye (1995b):

Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree, including relational features of the tree “from here”. (Harman 1990: 39)

Focus your attention on a square that has been painted blue. Intuitively, you are directly aware of blueness and squareness as out there in the world away from you, as features of an external surface. Now shift your gaze inward and try to become aware of your experience itself, inside you, apart from its objects. [...] The task seems impossible: one’s awareness seems always to slip through the experience to blueness and squareness, as instantiated together in an external object. (Tye 1995b: 30)

Perceptual, and especially visual experiences seem to make the transparency idea particularly vivid and indeed the transparency thesis may seem at least *prima facie* plausible as far as these experiences are concerned. Transparency theorists, however, typically argue that the transparency thesis is true not only of perceptual experiences, but of *all* conscious experiences: for any conscious experience, one can never attend to the experience itself; all one can attend to is what the experience is *about* (i.e. what the experience is an experience *of*). Accordingly, introspecting any conscious experience can only involve attending to the intentional content of the experience—it can never involve attending to the experience itself.

Finally, claim [4] is just a restatement of what I have called, above, the ‘transparency account of phenomenal-concept acquisition’: for any phenomenal concept *c*, associated with experience *C*, *c* is acquired not by attending to *C*, but by attending to *C*’s intentional object. To defend [4], the transparency theorist of phenomenal-concept acquisition is committed to [1], [2], and [3].

Claims [1]-[3], however, have all met with fierce objections in the extant literature. The anti-transparency arguments are copious and most of them very convincing. Considering them in detail exceeds the scope of this dissertation, so I will limit myself to summarizing some of those arguments.

As for [1], although intentionalism seems to be intuitively true of *some* experiences (e.g. *perceptual* experiences), that it is true of *all* experiences is not straightforward.²⁰ Consider algedonic experiences first. Some intentionalists (or representationalists) have argued that what *pain* experiences represent is bodily damage (Dretske 1995, Tye 1995b, Tye 1997): when you feel pain in your right knee, you have an experience that represents bodily damage in your right knee. Even granting that pain experiences are intentionally directed toward bodily damage in this way, one may wonder what the intentional content of *pleasure* might be, since surely it does not seem to be some kind of *bodily flourishing* (Massin 2013). Bodily sensations such as tickles, itches, and orgasms are also often cited as counterexamples to [1] (Block 1995a and 1995b, for example, argues that orgasms are not intentional): there does not seem to be a straightforward way to determine what the intentional content of these states could be. On the representationalist side it has been argued that, for example, while having an orgasm “[w]hat one experiences, in part, is that something very pleasing is happening down *there*” (Tye 1995a: 269), and that that is the intentional content of an orgasm experience. On a plausible interpretation of the view, ‘pleasing’ means ‘causing a pleasure experience’. However, first, although orgasm experiences clearly have pleasing features, they do not strike me as having the propositional content ‘something causing a pleasure experience occurs down *there*’; particularly, orgasms do not seem to have the metarepresentational feature that such an account would involve.²¹ Second, arguably, defining the intentional content of orgasms in terms of pleasure seems to require a prior definition of the intentional content of pleasure experiences which, as just noted, is not straightforward. Furthermore, defining an intentional content for *moods* has posed a particularly difficult challenge to defenders of [1] (Searle 1983). Even a hard-core representationalist like Dretske ends up admitting that the extent of the representationalist thesis may be somewhat limited by some “experiences—a

²⁰ Incidentally, there is growing literature raising doubts even about the plausibility of intentionalism about perceptual experiences. Travis (2004), for instance, argues that perceptual experiences are not representational. See Brogaard (2014) for further discussion.

²¹ For an alternative objection, see Block (1996: 33): “I can have an experience whose representational content is that my *partner* is having a very pleasing experience down *there* that changes in intensity, and although that may be pleasurable for me, it is not pleasurable in the phenomenally impressive way that graces my own orgasms.”

general feeling of depression, for example—about which I do not know what to say” (1995: xv).

On the representationalist side, efforts have been made to overcome the issues mentioned in the previous paragraph and defend the idea that all experiences have intentional content. Moods, for instance, are sometimes said to be intentional states directed to the entire world, or to any object one may happen to experience (see e.g. Solomon 1976). However, even if suitable intentional contents can be found for all the types of experience listed above, that their phenomenal character is *exhausted* by such contents, as required by [2], is even more controversial.

For one thing, there are some aspects of algedonic and emotional experience that do not seem to fit smoothly that sort of reductive representationalist account, namely their affective or evaluative component (Aydede and Fulkerson 2014). Consider pain experience again. We can distinguish at least two components in pain experiences: a sensory component (the qualitative aspect of pain, which may differ depending on the kind of pain experience—that which makes the difference between, say, a burning pain and a stubbing pain) and an affective or evaluative component (roughly, the negative feeling associated with the pain experience).²² Both aspects are *phenomenal*, so the strong representationalist should reduce both of them to some representational content. As mentioned above, the sensory aspect of the phenomenology is reduced to the representation of bodily damage. The affective aspect, instead, is reduced to representing such bodily damage as *harmful*. This representational content fixes the accuracy conditions of a pain experience: at the very least, one has a pain experience only if the bodily damage (located in the relevant body part) is harmful. But imagine that your pain in your knee is due to the “removal of scar tissue due to a serious burn by a doctor to prevent further pathological complications” (Aydede and Fulkerson 2014: 186). In this case, the bodily damage is *not* harmful: on the contrary, it *prevents* future harm. Strong representationalism predicts that, in this case, what you have is not real pain, but merely *illusory* pain. Yet, this seems to be simply false, at least on a widespread, commonsense (as well as scientific: see Aydede 2009) notion of pain. Your pain experience caused by the surgical operation is *not* illusory.

Moods especially seem to escape the representationalist reduction entailed by [2]. A case against representationalism about moods is put forward

²² There is a pathological condition, called *pain asymbolia*, that seems to support this two-component analysis of pain experience. In people affected by pain asymbolia, the pain sensation is felt but no affective component is associated with it.

by Kind (2014), who argues that although moods may have intentional content, their phenomenology is too rich to be fully captured by their intentional content. Moreover, she claims, the phenomenal character of moods and their intentional content can vary independently of one another (Kind 2014: 124-129). Arguably, the phenomenology of a depressive mood may change without the world seeming any different: the depressive feeling may, for instance, increase in intensity without the world looking any 'blacker'. Moreover, our perceptual representation of the world may change consistently without this being correlated with any change in mood phenomenology: the phenomenology of your depressive mood is not affected by the fact that the objects represented by your perceptual experience and the properties you represent them as having change.

A more general worry about strong representationalism is expressed by Charles Siewert (2004). Recall, strong representationalism is the view that the phenomenal character reduces to the representational content. As Siewert points out, a consequence of the view is that "phenomenal features can be *explained by identifying them with* a species of representation" (2004: 27, emphasis original). The representational content, to which the phenomenal character reduces, must be spelled out *entirely* in terms of worldly objects and their properties. Crucially, it must be spelled out *independently* of any reference to the phenomenology, that is, independently of *how* objects and their properties *appear* (look, sound, smell, etc.) to one (in Siewert's terms, it must be "specified in terms other than merely as *ways of seeming*"), on pain of circularity. Since strong transparency theorists maintain that their claim is introspectively supported, we should expect that such a non-circular specification of representational contents be introspectively available. However, on a closer look, this seems rarely, if ever, feasible. For although mere reference to external objects and their properties may help us give a partial or coarse-grained specification of the representational content, an *exact* specification will require reference to *how* the relevant object *appears* in experience. Consider the following example from smells by Siewert:

It seems that, when I think of precisely *what* kind of odor it is that I am smelling, I am sometimes quite unable to think specifically of just *that* odor, relying on introspective resources, in any other way than simply as: *the odor of what smells this way to me*. I may of course classify it in some more informative way [...]. But I can't always or even commonly use such characterizations to distinguish all the ways of smelling to me that I can distinguish in thought. In any case, the effort to give such a characterization is guided by some sort of attention that precedes it, which focuses the question "What is *that* odor?" on something's smelling this way to me [...]. So, to think,

of an odor in this manner, I attend to its smelling to me in a certain way.
(2004: 28, emphasis original)

Siewert argues that further evidence against the idea that representational content is fully specifiable, introspectively, independently of any reference to phenomenal character, comes from perspectival vision and gestalt switch. In both cases, when one tries to specify the content of one's experience on the basis of introspection, one needs to appeal and thereby attend to how things *look* to one, thus to the phenomenal character of one's experience.

As for [3], a vast critical literature has been produced against it (Block 1996, Kind 2003, Loar 2003, Siewert 2004, Smith 2008). Indeed, various arguments in defense of the claim that introspection involves attending to one's experience have been provided (Goldman 2006, Petitmengin 2006, Hill 2009, Siewert 2012). As an example, consider moods again. The transparency of experience thesis predicts that, when you are, say, in an elated mood, and you introspect your elation, your attention is directed not to the experience you have, but rather to some external object, or perhaps to some bodily states. Yet, this does not seem to be the case. As Kind points out, “[i]n attending to my experience of elation, it seems most natural to describe what I’m attending to as the elation itself. I’m not simply attending to some feature of a changing series of things, or even to some unbound feature. Rather, I focus directly on what it feels like to be elated” (2014: 130). Similar considerations apply to other kinds of experience. Attending to a pain sensation, to an itch, or to an orgasm really seems to involve attending to the experience itself, to its phenomenal character, rather than to any object such an experience may represent (even granting that we can find a plausible intentional object for itches and orgasms). Some philosophers argue that the Transparency of Experience thesis fails even to apply to some instances of the kind of experience that seems to make the thesis initially plausible, namely, visual experience. By articulating a consideration already present in Boghossian and Velleman (1989), Smith (2008) argues that the phenomenon of blurry vision is a counterexample to the transparency thesis. Seeing things blurrily does not consist in representing objects *as being blurry*: blurriness is a feature of the visual experience, not of seen objects. Accordingly, when one is aware of blurriness one is aware of a feature of one's experience and not of a property of object's in one's environment.²³

²³ Smith strengthens his point by insisting on the fact that the phenomenology of blurry vision is utterly different from the phenomenology of seeing *fuzzy* objects, that is objects that have undefined borders. In the latter case, fuzziness *is* represented in experience as a feature of the seen objects.

If we consider again Siewert's argument above, we see that there is also a more general reason to reject claim [3]. For, if, in order to specify, on the basis of introspection, the content of one's experience one needs to appeal to the way things appear to one, one must be able, in introspection, to *attend to the experience itself*:

[W]hen, for example, there is some figure that looks blue and square to me, its *looking* to me that way "falls within the scope of my attention", just as much as, and together with, the figure itself and its blueness and squareness. The figure, its properties, and its appearing to me, all come together as a package, as far as this act of attention is concerned. (Siewert 2004: 20)

The idea is that, in introspection, attention is always, at least partly, directed toward the experience itself. Importantly, directing one's attention to the experience does *not* entail *diverting* it, or turning it *away*, from what the experience is about (that is, from its content). As Siewert makes clear, attention to the experience *comes together* with attention to the object. This makes justice to what strikes me as right in the transparency intuition, namely that, at least as far as some experiences are concerned (surely visual experience, probably also some other perceptual experiences, and perhaps some non-perceptual experiences too), drawing one's attention to the relevant experience seems to partly involve drawing it to (or keeping it focused on) what the experience is about.

It is also important to notice that the denial of the transparency thesis expressed by [3] is distinct from and does not entail any of the following controversial theses. First, it does not entail non-intentionalism, that is, it does not imply rejecting the thesis that all conscious experiences have intentional content. This is important because, notwithstanding the difficulties that emerged during our discussion of [1], intentionalism is widespread among philosophers of mind and many of them would regard with skepticism any view that is inconsistent with it. As argued by Tim Crane (2000), the denial of transparency does not even entail the denial of a *strong* form of intentionalism: it is compatible with a view such that the phenomenal character can be reduced to the intentional features of experience. Crane invites us to consider that the intentional structure of experience is constituted not only by its intentional content, but also by the intentional mode (or attitude), that is, the relationship the subject bears to the intentional content while having the relevant experience (Searle 1983). What anti-transparency entails is that the phenomenal character is not reducible to the intentional *content*. This leaves the possibility open for it to reduce to the conjunction of *intentional content and intentional mode*. Relatedly, the denial of the transparency thesis does not entail commitment to the existence of

qualia, that is, to features of experience that are non-intentional and intrinsic to it (Crane 2000; Siewert 2004). Finally, anti-transparency does not entail anything like an inner sense theory of introspection, that is, a theory according to which introspectively attending to an experience is analogous to perceptually attending to a worldly object, as if the experience could be attended to in isolation, independently of its being an experience of a certain subject (Siewert 2004 and 2012).

Once we adequately qualify and analyze [3] and thereby recognize, on the one hand, its introspective and phenomenological inadequacy and, on the other hand, how little theoretical commitments its denial involves, we come to have very good reasons to reject [3] and, with it, the whole idea that introspective attention cannot be directed toward the experience itself.

To my knowledge, there are no direct arguments against (or for) [4] in the extant literature. However, if credence in either of [1]-[3] is weakened by the arguments mentioned above, so will be that in [4].

Moreover, there seem to be independent reasons to reject [4]. Except perhaps perceptual experiences, it seems much more plausible that the concepts we use to think about our experiences are formed by attending to properties of the experience itself, rather than to what the experience is about. Aydede (2009) develops an argument from pain reports which may provide this intuition with some support. He compares pain reports (“I feel a jabbing pain in the back of my hand”) with perceptual reports (“I see a dark discoloration on the back of my hand”) and argues that, notwithstanding their superficial similarity (they both seem to express a perceptual relation between the subject and a physical property of a body part of the subject—tissue damage in the former report and discoloration in the latter report), they express, in fact, fundamentally different thoughts. For, upon closer semantic analysis, different from perceptual reports, pain reports do *not* express the attribution of a physical property to a physical object (body part) plus the presence of a perceptual relation between the subject and the relevant property. Rather, they express the self-attribution of an *experience*, namely a pain experience (which may itself represent presence of bodily damage). This suggests that, different from perceptual concepts such as DARK DISCOLORATION, the concept PAIN is normally used by people to refer to a certain type of experience, rather than to any worldly object or property of worldly object:

[I]n pain, contrary to first appearances, our immediate and spontaneous interest (epistemic or practical) is in the experiences themselves in the first place, rather than in what objects or conditions these experiences represent. (2009: 541)

If the concept PAIN is normally used to refer to pain *experiences*, rather than to whatever these experiences are *about*, then it is much more plausible that such a concept is acquired by attending to pain experiences themselves, rather than to their intentional content.

Arguably, this line of thought generalizes to other kinds of conscious experience. Take emotions. Thinking introspectively (or making an introspective judgment) about the anger I feel is different both from thinking (making a judgment) about the person who causes my anger experience and from thinking (making a judgment) about the bodily states that accompany my anger.²⁴ When I think that my anger experience is intense, my thought makes no reference to the person who causes my anger or to any of my current bodily states. Accordingly, it seems much more plausible that the concept ANGER EXPERIENCE is acquired by attending to the anger phenomenology itself, rather than to the cause of the anger or some bodily states: it would be weird if we thought about the properties of our experiences through concepts that were originally formed by attending to something other than phenomenal properties.

Similar observations apply to moods like euphoria or depression. Even accepting that they are intentional states directed at the whole world, “casting happy glows or somber shadows on every object and incident of our experience”, as Solomon (1976: 173) writes, it seems that thinking about one’s depressive mood is different from thinking about the objects in one’s environment as shrouded in ‘somber shadows’. When I think about my depressive mood, it is the bad way I feel I think about, not the things around me. It may be that when I feel depressed I *also* have characteristic thoughts about my surroundings (that this object is pointless, that that person does not care about me, etc.). But, although they may be *related* to my depressive mood (perhaps partly *caused* by it), those thoughts are not *about* it. Accordingly, it is much more plausible that the phenomenal concept associated with depressive mood is acquired by attending to the depressive experience itself, rather than to objects in one’s environment.

Moreover, the transparency account of phenomenal-concept acquisition seems to imply that the phenomenal concept DEPRESSIVE EXPERIENCE be acquired by merely attending to any object one may encounter when one is depressed, arguably in virtue of something that all the objects one encounters when one is depressed appear to have in common (e.g. being shrouded in ‘somber shadows’). However, I find this highly implausible. For

²⁴ These correspond to the two main representationalist accounts of emotional phenomenology. See e.g. Prinz (2004).

one thing, it is arguable that, for each mood, there is some property that *all* the objects one encounters when one is in that mood appear to have in common. First, when one is depressed, the negative aspects of things around one may become more salient, or one may attribute some negative properties to the objects around one, but it does not seem to me that exactly *the same* property is attributed or becomes salient. Secondly, it is arguable that *all* the objects one encounters when one is depressed appear as having such negative properties. Moreover, it seems to me that one could acquire the concept DEPRESSIVE EXPERIENCE from inside a dark room, having perceptual experience of no object at all, simply by attending to the way one currently feels.

Therefore, the defender of the transparency account of phenomenal-concept acquisition faces considerable challenges. To defend such an account, not only should it be established that [1], [2], and [3] are true of all phenomenally conscious experiences. It should also be shown that all phenomenal concepts are acquired by attending not to experience itself, but to its intentional objects. For the reasons just adduced, such challenges are hard to overcome. Therefore, even if a transparency account of *some* basic phenomenal concepts (those associated with some visual experiences and perhaps those associated with some other perceptual experiences) could be developed, a transparency account of the acquisition of *all* basic phenomenal concepts does not seem to be plausible.

3.2.3. *Mere consciousness.* Besides the radical transparency view just discussed, there is another option my objector may consider in defense of (ii)—i.e. the claim that if a basic phenomenal concept *C* is acquired, it is acquired on the basis of something other than introspection. It might be argued that, to acquire basic phenomenal concepts, *introspecting* one's experience is not required: all one needs is *having* the relevant experience. So, for instance, to acquire the phenomenal concept PAPAYA-TASTE EXPERIENCE, all one needs is to have a papaya-taste experience—one does not need to introspect the experience. Accordingly, basic phenomenal concepts are (typically) acquired not via introspection but via the having of the experience.²⁵

The difference between *having* and *introspecting* an experience consists, at the very least, in the fact that the latter but not the former involves *attending to* the experience. Introspection is something we actively do, a mental act we perform. Therefore, it does not occur automatically but requires an effortful and voluntary act on the part of the subject. Such an act partly

²⁵ I owe this objection to an anonymous referee for *Erkenntnis*.

consists in the subject's drawing their attention to the experience. Merely having the experience, in contrast, does not require any effort or voluntary act and can occur independently of the subject's attending to the experience.²⁶

As noted, acquiring a phenomenal concept *C* entails (a) coming to have a representation of the experience *C* which must (b) enable the subject to distinguish *C* from other experiences and recognize it as an instance of the *C*-experience type and (c) be available to be deployed by the subject in cognition and action-guidance. Now, for one thing, that merely having a conscious experience implies forming a representation of it is controversial. On first-order representationalism about consciousness (Dretske 1995, Tye 1995b), for example, having a conscious experience entails representing things in one's environment, but it does not entail representing the experience itself. On this view, then, having an experience is not sufficient for having a representation of it—therefore it is not sufficient for acquiring a phenomenal concept of the experience.

There are, however, other theories of consciousness—e.g., higher-order representationalism (Rosenthal 1997, Lycan 1996) and self-representationalism (Kriegel 2009)—that *do* imply that having a conscious experience entails having a representation of it. These theories—different from first-order representationalism—thereby fulfill (a). However, whether they fulfill (b) and (c) is debatable. To acquire the phenomenal concept *C*, a subject must form a representation *of the right kind*: the representation must enable the subject to distinguish and recognize *C*-experiences. Yet, it seems that merely having an experience is not enough for the subject to subsequently be able to distinguish and recognize an experience of the same kind.²⁷ Imagine you are an aspiring sommelier. Arguably, part of your aspiration is to acquire a great number of fine-grained wine-taste-associated phenomenal concepts,

²⁶ Unless we adopt a view, such as Prinz's (2011), on which an experience can only be conscious if it is attended to. On views of this kind, the notion of having an experience collapses into the notion of introspecting an experience (as conceived here).

²⁷ Defenders of the HOT theory of consciousness (Rosenthal 1997) might disagree. The higher-order representation being a thought, on this view, it is propositional and conceptual and thereby allows for recognition. However, arguably, the HOT theory faces a similar challenge to that faced by conceptualism about introspection, i.e. that of explaining phenomenal-concept acquisition. Plausibly, to acquire a phenomenal concept *C* one must, at the very least, *have* a conscious *C*-experience. But if having a conscious *C*-experience requires possessing *C* (i.e. the phenomenal concept that partly constitutes the unconscious thought that makes a *C*-first-order representation conscious), it is mysterious how we come to possess phenomenal concepts at all. I thus would not recommend my objector to take the HOT theory as a basis for the claim that phenomenal concepts can be acquired via the mere having of an experience.

i.e., phenomenal concepts that enable you to distinguish and recognize taste experiences associated with a great number of different types of wine. To achieve this result, for any wine-taste experience you want to acquire a concept of, you need to gain and store information about its phenomenology. Quite obviously, merely drinking a large variety of wines (and thereby having taste experiences associated with many different types of wines) is not sufficient to accomplish this task. (If the task were so easy to accomplish, I myself would be an expert sommelier by now.) The information-gaining-and-storing process required for wine-taste phenomenal-concept acquisition is an effortful and deliberate activity, which involves your directing and focusing your attention to each wine-taste experience. More generally, for you to acquire a representation that enables you to distinguish and recognize C-experiences, you must not only *have* a C-experience, but also *attend to*, and thereby *introspect* it.

For similar reasons, the mere having of an experience is not enough to satisfy (c). As noted, possessing a concept requires it being available for personal-level deployment in cognition and action-guidance. Accordingly, a subject possesses a phenomenal concept C only if they can actively and deliberately retrieve it and manipulate it to form thoughts about the experience C and take decisions for action based on those thoughts. But merely having an experience is not enough to gain such personal-level cognitive access.²⁸ Gaining personal-level cognitive access requires the subject's actively collecting information about the phenomenology of the relevant experience, which in turn, as noted, requires attending to and thereby introspecting the experience.

In sum, the mere having of an experience C is not enough to acquire the phenomenal concept C. To acquire a mental representation of C which enables the subject to distinguish and recognize C-experiences and which is available to be deployed in cognition and action-guidance, the subject must *introspect C*.

*

For most basic phenomenal concepts, then, no acquisition method alternative to introspection is available. *Extrapolation* is not a promising phenomenal-concept acquisition method because, for one thing, it does not seem to be

²⁸ My reflection on this point has benefitted from conversation with Luca Gasparri.

sufficient for the subject to acquire the ability to distinguish a certain experience and recognize it upon encounter; for another thing, even if it did, only a very few phenomenal concepts could be acquired by extrapolation. The *transparency* method is no better off: at most, it is only available for the acquisition of a small number of phenomenal concepts (namely those associated with some perceptual experiences). Finally, the mere *having* of an experience is never sufficient to acquire a phenomenal concept. Therefore, for most phenomenal concepts, if they are acquired, they must be acquired by introspection.

3.3. *Defending (P3)*

In the previous subsection, I have argued that, except perhaps for phenomenal concepts associated with some perceptual experiences (for which a transparency account might perhaps be developed), it is plausible to suppose that if a basic phenomenal concept *c* is acquired, then it is acquired by introspecting *C*-experiences. In this subsection, I argue that if all introspective states were conceptual, then no phenomenal concept *c* could be acquired by introspecting *C*-experiences. Conceptualism about introspective states therefore entails an implausible version of nativism according to which most basic phenomenal concepts are innate.

An introspective state's being conceptual implies that one cannot have it without deploying some (already possessed) concept(s). If introspecting a *C*-experience depends on the deployment of some phenomenal concept(s), then it depends either on the deployment of *c* (i.e. the phenomenal concept associated with the phenomenal property *C*) or on the deployment of some phenomenal concept other than *c* (*tertium non datur*). Accordingly, if the content of all introspective states is conceptual, for any phenomenal concept *c*, if *c* is acquired by introspection, then either (ia) *c* is acquired by introspecting *C*-experiences, where introspecting a *C*-experience depends on the subject's deploying the concept *c*, or (ib) *c* is acquired by introspecting *C*-experiences, but where introspecting a *C*-experience depends on the subject's deploying some concept other than *c*.

Now, there is no *c* for which (ia) can be an option. For if introspecting *C*-experiences depends on the subject's deploying *c*, then *c* must already be possessed by the subject in order for them to introspect the experience. The point is that if the ability to have an introspective state depends on deployment of a certain concept, then that concept must be possessed by the subject *prior to* having the relevant introspective state. Accordingly, it cannot be acquired by way of that very introspective state.

What about (ib)? If introspecting C-experiences requires the deployment of some phenomenal concepts, but does not depend on the subject's deploying C, then it depends on the subject's deploying some concepts F_1, \dots, F_n different from C. Accordingly, acquiring C by introspecting C-experiences would require that the content of C be built up by composition from other introspective concepts F_1, \dots, F_n . If, for instance, the concept PAIN were acquired by introspecting pain experiences and pain experiences depended on the subject's possessing some concepts other than PAIN, then the acquisition of PAIN would depend on composition of those other concepts.

However, although it may be an option for some phenomenal concepts, (ib) is not an option for the phenomenal concepts under consideration here, namely *basic* phenomenal concepts (for basic phenomenal concepts, by definition, are *not* built up by composition from other concepts).

Moreover, a conceptualist theory according to which introspecting a C-experience depends on the subject's deploying some concept different from C strikes me as pointless. The whole point of maintaining that introspecting an itch experience depends on the subject's deploying some concept is, plausibly, that it depends on the subject's deploying the concept ITCH. The idea is that one cannot introspect an itch experience unless one can distinguish it from non-itch experiences and recognize it *as* an itch experience. Why should one maintain that introspecting an itch experience depends on the subject's recognizing their experience as something *other* than an itch?

So, even if some phenomenal concepts can be acquired by composition, the foundational layer of non-compositionally acquired phenomenal concepts still has to include a great number of introspectively-acquired concepts, whose acquisition cannot be accounted for in a conceptualist framework. For all I am arguing here, there may well be many phenomenal concepts acquired by composition of more basic phenomenal concepts—as long as the foundational phenomenal concepts, from which the others are compositionally acquired, are themselves acquired through a nonconceptual form of introspection (what I have called primitive introspection). The only alternative is that most of the numerous foundational-layer phenomenal concepts are innate. But as I have already argued, such a radical form of nativism is not very promising.

*

I conclude that, if radical nativism about phenomenal concepts is to be avoided, then for many introspective concepts C , C must be acquired by introspecting C -experiences. On a conceptualist framework, there is no viable way in which such concepts may be acquired. If we are to deny that they are innate, we must accept that they are acquired with the aid of a nonconceptual form of introspection—primitive introspection.

4. Pure demonstratives, direct phenomenal concepts, and primitive introspection

In this section I want to briefly clarify the extent of the argument from phenomenal-concept acquisition developed in §3, by making some considerations about the relationship between primitive introspection and the special kinds of concept mentioned in §1, that is, pure demonstratives and direct phenomenal concepts.

At least three possible views are compatible with the outcome of the phenomenal-concept acquisition argument:

- [a] There is a kind of introspection that is wholly nonconceptual (it does not involve the deployment of *any* concept at all);
- [b] There is a kind of introspection that only involves the deployment of pure demonstratives;
- [c] There is a kind of introspection that only involves the deployment of direct phenomenal concepts.

All these views are consistent with the idea that there is a kind of introspection that does not depend on the deployment of any phenomenal concept *already* possessed by the subject *prior to* introspecting. Accordingly, the view that there is primitive introspection may be articulated in at least three different ways:

- [a] Primitive introspection is a kind of introspection that is wholly nonconceptual (it does not involve the deployment of *any* concept at all);
- [b] Primitive introspection is a kind of introspection that only depends on the deployment of pure demonstratives;
- [c] Primitive introspection is a kind of introspection that only depends on the deployment of direct phenomenal concepts.

Now, the view *I* favor is the one expressed by [a]. However, what I will say in the remaining of this dissertation will also be compatible with [b]. Indeed, it is not wholly clear that [a] and [b] are distinct views. For, if mental demonstration is possible at all (*cf.* Millikan 2012), what does deploying a *pure* phenomenal demonstrative (i.e. a demonstrative with *no* descriptive component) involve if not simply drawing one's attention toward a certain phenomenal state? But that drawing one's attention toward a phenomenal state is necessary to primitive introspection is a claim that I am ready to make (indeed, it is a claim I am going to defend in Chapter 3). So, at a closer examination, [b] seems to collapse into [a].

[c], instead, *is* distinct from [a] and, as we will see, it is inconsistent with some claims I am going to make in the remaining of this dissertation. On my view, primitive introspection does not depend on the deployment of direct phenomenal concepts. I am going to argue for this in Chapters 5 and 6, but the essential idea is that, if there are direct phenomenal concepts, their formation depends on primitive introspection. Therefore, it cannot be the case that primitively introspecting requires deploying direct phenomenal concepts (on pain of circularity).

Conclusion

In this chapter I have proposed an argument from phenomenal-concept acquisition for the existence of primitive introspection. I have argued that denying the existence of primitive introspection entails an implausible, radical version of nativism about phenomenal concepts. For if the majority of basic phenomenal concepts is to be acquired on the basis of introspection, then they must be acquired on the basis of a nonconceptual form of introspection. There are only three alternatives: extrapolation, transparency, and the mere having of an experience. However, as I have shown, if any of these alternative methods is viable at all, all it can allow for is the acquisition of a very small number of phenomenal concepts. Therefore, if phenomenal concepts are acquired at all, their acquisition must be grounded in primitive introspection.

PART 2: NATURE

CHAPTER 3

METAPHYSICAL STRUCTURE: PRELIMINARIES

In the previous chapter I argued for the *existence* of primitive introspection. In this chapter and the next I am going to investigate more deeply its *nature*. In the next chapter, I am going to propose an account of the metaphysical structure of primitive introspective states. By clarifying the main features of primitive introspection, the present chapter prepares the ground for that.

At a first approximation, primitive introspection is a process by which the subject can get information about the phenomenology of their current conscious experience without classifying or recognizing it as an instance of a certain experience type (that is, without deploying any phenomenal concept—bearing in mind the caveats explained in Chapter 2). It therefore involves a mental act, or state, which is directed toward one’s current conscious experience. To better understand what primitive introspection is and how it enables phenomenal information acquisition, a more extensive explanation of what it is and what it involves is in order. First, an *act* of primitive introspection is not the same as a *state* of primitive introspection. Therefore, an account of each and an explanation of their relationship are due. Second, to get information about the phenomenology of one’s current experience, one needs to *attend to* that experience. Therefore, primitive introspection involves attention, and what is attended to is the phenomenology of one’s conscious experience. An account of both the role of attention in primitive introspection and the sense in which the phenomenology of one’s conscious experience is attended to is in order. Moreover, attention alone is insufficient to explain the epistemic role of primitive introspection. A condition must be added to explain the fact that primitive introspection enables the subject to form judgments about, and possibly know, their own current experience. Finally, an elucidation of the target of the state of primitive introspection is needed to better understand that toward which a primitive introspective act is directed. The present chapter’s purpose is to make a first pass at an understanding of primitive introspection by attempting a treatment of these issues, some of which will be addressed more deeply in the next chapter.

First, I make some preliminary remarks about what is involved when one primitively introspects (§1). I outline the features of primitive introspection and make explicit my main assumptions about them. Second, I draw a distinction between *act*, *target* and *state* of primitive introspection (§2).

This will be particularly helpful to understand the main features of primitive introspection and pursue a clearer analysis of its nature. Third, I outline what I take to be the main features of the primitive introspective act, namely *attention* and *apprehension* (§3). Finally, I focus on the *target* of primitive introspection, namely that to which the primitive introspective act is directed (§4 and §5).

*

1. Four conditions for primitive introspection

Primitive introspection is a process through which one can acquire information about the phenomenology of one's current conscious experience without classifying it as an instance of a certain experience type. There are at least four necessary conditions for primitive introspection: (i) the introspected experience must be *conscious*; (ii) the introspected experience must be *present*; (iii) the subject must *attend* to the phenomenology of the relevant experience; (iv) the subject must concomitantly *apprehend* the experience.

As regards (i), as explained in Chapter 1, primitive introspection is a kind of *introspection of phenomenal states*. Recall, introspection of phenomenal states is the distinctively first-personal method through which one can get knowledge of the phenomenology of one's current *conscious* experience. Therefore, *ex hypothesis*, primitive introspection only targets phenomenally conscious states.¹ To get a better grasp of why this is so, consider the following example. Although Lucie trusts her partner more than anyone else, and she is certain that her partner will never be unfaithful to her, her behavior is sometimes very akin to that of a jealous person. On one plausible interpretation of Lucie's behavior, she is jealous of her partner. However, she has never consciously *felt* jealous: jealousy is not among Lucie's conscious states. By consequence, and quite straightforwardly, Lucie cannot *introspect* her jealousy. She might become aware of her behavior, observe that it is very akin to that of a jealous person, and thereby come to realize that she has an unconscious jealousy toward her partner. In such a case, she eventually comes to believe that she is jealous; however, it is not by introspection, but rather by behavior observation and inference, that she

¹ Some philosophers (Rosenthal 2005; Carruthers 2005) think that there can be unconscious phenomenal states. If there are, they are *not* potential targets of *introspection of phenomenal states* as I characterize it here (in such case what I call 'introspection of phenomenal states' should be rather labeled 'introspection of *conscious* phenomenal states').

comes to believe that. Her jealousy being unconscious, introspection cannot be directed toward it. Since primitive introspection is a kind of introspection, it requires that what is primitively introspected be a conscious experience.

As also noted in Chapter 1, introspection of phenomenal states only targets one's *current* conscious states. Therefore, a phenomenally conscious state can be primitively introspected only if it is currently present—as per (ii): past experiences (or future experiences, for that matter) cannot be primitively introspected.

An experience's being conscious is not sufficient for one to be *introspectively* aware of it. In order to introspect, one must also *attend* to the experience—as per (iii).² By introspection, one acquires some information about the phenomenology of one's experience which one could not acquire by merely *having* the experience. The acquisition of such additional information is partly due to the fact that when one introspects one's experience, one attends to its phenomenology.³ Imagine you have a quite intense stomachache, but you keep working notwithstanding it. Arguably, although your stomachache experience is conscious (the pain does not disappear), you can divert your attention from it while focusing on your work.⁴ When you decide to go and see a doctor, and she asks you to describe what the stomachache feels like, you cannot answer her question unless your attention is directed to the stomachache experience: you need to focus your attention on the experience to get the relevant information about its phenomenology and report it to the doctor. Such information was not available the same way to you when most of your attentional resources were absorbed by your work.⁵

² On the basis of transparency considerations already encountered in Chapter 2, strong representationalists such as Tye (1995b) or Dretske (1995) would probably argue that experiences cannot be attended to—only their intentional object can. I hope to have provided, in Chapter 2, sufficient motivation for being suspicious about this kind of reasoning. The reader should also bear in mind that rejecting radical transparency and strong representationalism à la Tye is consistent with intentionalism (even with a strong version of intentionalism), as well as with a *weak transparency* view such that, by attending to one's experience, one *also* attends to the experience's intentional object.

³ I will say something more about what is involved in such an attending in §3.1 of this chapter. The relationship between primitive introspection, attention and information acquisition will be dealt with extensively in Chapter 6.

⁴ Some (e.g. Prinz 2011) would object that the stomachache experience cannot be conscious unless you attend to it—attention is necessary for consciousness. Although I disagree (arguments against the necessity of attention for consciousness are put forward by Mole 2008 and Smithies 2011), my claim that attention is necessary for introspection can be accepted regardless of whether attention is also necessary for consciousness.

⁵ My discussion of the relationship between attention and availability of phenomenological information is vague at purpose here. At the present stage of the discussion I am just drawing

Therefore, in order to be introspected, an experience must not only be conscious, but also attended to. The same applies to the more specific case of primitive introspection: it being a kind of introspection, primitive introspection requires drawing one's attention toward one's experience, beside such an experience's being conscious.

Merely drawing one's attention to a certain phenomenal state might not yet be sufficient to introspect it.⁶ To introspect a phenomenal state, one must apprehend it—as per (iv). Apprehending a phenomenal state implies, first, *referring* to it and, second, *acquiring information* about its phenomenology. It may be that, typically, attending to a phenomenal state implies apprehending it. However, there may be cases in which a phenomenal state is attended to but not apprehended. If, for example, you undergo a sudden intense pain in your stomach, your attention is probably grabbed by it. The way I use the term here, to *apprehend* the phenomenology of your pain state that sudden exogenous switch of attention is not yet sufficient: you must also engage in a voluntary and effortful activity, by which you keep attending to the pain state and acquire information about its phenomenology (about 'how it feels'). Such an activity does not necessarily follow the exogenous switch of attention: right after your attention is grabbed by the pain, you might decide not to engage in sustained attention to your pain state and switch your attention back to work instead.⁷

The notion of apprehension may sound somewhat mysterious. I am going to come back to it in due course—in Chapter 6 apprehension will be explained in terms of information acquisition. For now, I encourage the reader to rely on the examples provided above, as well as his or her own intuitions

an impressionistic picture of the main features of primitive introspection. The epistemological role of attention in introspection will be treated more extensively in Chapter 6. I will say a bit more on this already later in this chapter, in §3.1.

⁶ Thanks to David Chalmers for bringing my attention to this point.

⁷ The conceptual separation between attention and apprehension may also be established by showing that one can apprehend an experience without attending to it, i.e. when one is *peripherally aware* of the relevant experience (see e.g. Kriegel 2009, Ch. 5). I am sympathetic to a view along these lines, though developing it would require argumentative work which outstrips the scope of the present paper (for one thing, it relies on a controversial assumption, i.e., that for a mental state to be conscious the subject must be conscious of it). *If* there is non-attentive apprehension in peripheral awareness though, you may wonder how this is supposed to square with the idea, suggested in the main text, that apprehension requires voluntary sustaining of attention (thank to Uriah Kriegel for pointing this out to me). I cannot develop a full reply to this worry here (though this is planned to be object of further work), but, roughly, the idea would be that there are two kinds of apprehension, *attentive* apprehension (featuring in primitive introspection) and *non-attentive* apprehension (featuring in peripheral inner awareness). I will elaborate a bit further on this in §3.2.

on the matter, to get a preliminary grasp on the difference between mere attention and attentive apprehension.

Arguably, attention to a phenomenal state is typically accompanied by apprehension of the state's phenomenology. However, since the two are separable, they should be treated as distinct necessary conditions for primitive introspection.⁸ Primitive introspection could be therefore characterized, at least, as *attentive apprehension of the phenomenology of one's current conscious experience*. I take the above to be *necessary* conditions for primitive introspection, although I do not claim that they are also collectively *sufficient*.

2. Act, target, and state of primitive introspection

As characterized above, primitive introspection looks more like a mental *activity* than a mental *state*: both drawing one's attention to the phenomenology of one's experience and apprehending it, require an activity on the part of the subject. For one thing, both are *effortful*: they cannot be achieved if the subject takes a passive attitude toward their experience. Moreover, both are *voluntary*: they can only be achieved if the subject wants to achieve them.

It might be objected that, quite often, drawing one's attention toward a certain aspect of one's experience does not require any effort or act of will on the part of the subject. If one suddenly undergoes a very intense experience, for example, one's attention is automatically attracted by that experience. In such a case, one's will plays no role in the direction of one's attention. Even though the direction of one's attention does not necessarily require that one perform an action, *keeping* one's attention focused on the experience, as well as *apprehending* its phenomenology, *is* effortful and voluntary. Indeed, the mere shift of one's attention is not sufficient for having a primitive introspective act: without apprehension, it cannot be primitive introspective and without a voluntary and effortful focus of one's attention, it cannot be an act.

⁸ Note that, however, my claims concerning primitive introspection do not rely on the assumption that attention and apprehension are separable. If it turns out that attention entails apprehension, this will simply reduce to three the number of necessary conditions for primitive introspection, namely (i) consciousness, (ii) presence, and (ii) attention (entailing apprehension).

A primitive introspective act is directed toward (an aspect of) one's current conscious experience.⁹ It involves drawing one's attention to (and/or keeping it focused on) the experience and apprehending at least some of its phenomenology. What is primitively introspected (the relevant aspect of one's current conscious experience) is what I call the *target* of the primitive introspective act (or the *target of primitive introspection*).

Primitive introspection, however, is not just a mental act. By primitively introspecting, the subject enters a new mental *state*, namely a *primitive introspective state*. Thus, to every primitive introspective act corresponds a primitive introspective state. At a first approximation, a primitive introspective state is the *result* of a primitive introspective act.¹⁰ It is the conscious mental state the subject is in in virtue of engaging in the activity of primitive introspection. The state of primitive introspection is a conscious state with phenomenology. Therefore, *it is a phenomenal state* (it is important to bear this in mind because it will come up again in Chapter 4, where I develop my account of the metaphysics of the state of primitive introspection). When you primitively introspect the pain sensation in your stomach, the introspective state you are in has phenomenology, one aspect of which is painfulness. Therefore, the state of primitive introspection is a phenomenal state. An account of how the state of primitive introspection relates to the target of primitive introspection is an account of the *metaphysics* of primitive introspection. This will be the topic of Chapter 4.

A note on the relationship between the primitive introspective act and the primitive introspective state. I said that a primitive introspective state is the *result* of a primitive introspective act. This is misleading though, for it seems to imply that once the primitive introspective act yields its outcome (i.e. the primitive introspective state), it ceases to be performed. This is an inaccurate description of what really happens. To be sure, since a shift of attention is needed to enter a state of primitive introspection, it could legitimately be said that the act of primitive introspection (or, better, the *attention* component of it) *causes* the state of primitive introspection. Therefore, there is a sense in which the primitive introspective act often temporally precedes the primitive introspective state: typically, to enter a primitive introspective state, one first needs to direct one's attention toward the target experience. However, arguably, only the attention-shift aspect of the primitive introspective act temporally precedes the primitive introspective state. When one apprehends the phenomenology of one's experience, one is

⁹ The "one aspect of" specification will become clearer in §5.

¹⁰ Although the term "result" is not fully appropriate (more on this in a moment).

already in a primitive introspective state. Moreover, to continue to be in a primitive introspective state, one must keep engaging in the activity of primitive introspection: one must keep one's attention focused on the target experience and apprehend its phenomenology, both of which are effortful and voluntary activities carried out by subject. Therefore, the act and the state of primitive introspection occur concurrently. Even though they cannot be distinguished through a temporal criterion, they are distinct aspects of the phenomenon of primitive introspection.

Consider, by analogy, a person who is running.¹¹ On the one hand, running is an *activity*: it involves a series of actions on the part of the subject (putting one foot forward, while pushing the back foot on the ground to get the leap, compensating by moving the opposite arm forward to keep balance, and so on and so forth). On the other hand, there is a peculiar physical *state* the subject is in while s/he is running: fast heartbeat, deep respiration, muscle contraction, and so on. Although there is a sense in which the act of running temporally precedes the physical state of running (the subject could not achieve the state of running did not s/he perform the actions that initiate the movement of running), the act of running does not reduce to such an initial trigger. To continue to be in the state of running, the subject must keep performing a series of actions that are characteristic of the act of running. Although the act and the state of running co-occur, they are distinct aspects of the running phenomenon. Analogously, the act of primitive introspection initiates a state of primitive introspection but does not reduce to the triggering of the introspective state. To continue to be in a primitive introspective state one must keep performing some mental actions, namely attending to and apprehending the target phenomenal state.

So, an account of the act of primitive introspection is an account of the cognitive activity the subject must engage in in order to introspect. An account of the primitive introspective state is an account of the metaphysical structure of the mental state the subject is in while s/he is engaged in the activity of primitively introspecting.

The metaphysical structure of a primitive introspective state is more complex than that of a merely conscious experience: it involves not only the aspect of the phenomenology of one's conscious experience that is primitively introspected, but also a certain amount of attentional resources. Moreover, in virtue of its metaphysical structure, the primitive introspective state can be epistemically significant: by being in a primitive introspective state one is in a position *know* one's current experience, as I will argue in Chapters 5 and 6.

¹¹ This analogy was suggested to me by Raffaello Antonutti.

An analysis of the metaphysical structure of primitive introspective states will be provided in Chapter 4. Hopefully, the present paragraph's assertions, which may now sound rather sibylline, will become clearer and fairly motivated.

If all this is right, inquiry on the nature of primitive introspection should rely on the distinction between the act, the target, and the state of primitive introspection. The primitive introspective *act* is an act of attention and apprehension. A more detailed characterization of it is provided in §3. What the primitive introspective act is directed to is the *target* of primitive introspection. Roughly, the latter is the subject's primitively introspected experience. In §4 I make some preliminary considerations concerning the extent and temporality of the target of primitive introspective acts. A more detailed characterization of the target of primitive introspection requires specifying its ontological status. In §5 I consider a number of options and point out the challenges they face. A primitive introspective *state* is the conscious state one is in in virtue of engaging in the activity of primitive introspection. It is a state of attentive apprehension. Even though it differs from the primitively introspected target experience, the primitive introspective state bears an intimate relation with it. As noted, this will be discussed in Chapter 4.

3. The act of primitive introspection

As noted in the previous section, the act of primitive introspection involves (1) a shift of attention, or a redistribution of one's attentional resources, and (2) apprehension of the target experience. In this section, I say a bit more about the notions of attention and apprehension at work here.

3.1. Attention

There is an intuitive notion of attention, according to which attention is “the selective or contrastive aspect of the mind: when you are attending to something you are contrasting what you pick out with what remains in the background.” (Watzl 2011: 843). From a pre-theoretical point of view, attention seems to refer to a unified and pretty clearly defined phenomenon. However, research on attention in psychology and cognitive sciences has called this pre-theoretic intuition into doubt. Attention has turned out so complex a notion, that many have doubted it really refers to one unified phenomenon. For one thing, it is appealed to for explaining a considerable number of cognitive processes (selection of information, feature binding,

simultaneous cognitive processes' competition, coordination of movement etc.). Reductionist accounts of attention often seek to explain attention in terms of one of those processes. However, there is no single underlying mechanism that all those processes seem to have in common. Moreover, it has been shown that the mechanisms which underlie many of those processes can occur even in the absence of attention.¹² Secondly, there are different forms or notions of attention, partly depending on the following distinctions (Watzl 2011b): *all-or-nothing* (on/off) vs. *gradable* attention (the former implies that either attention is directed to a certain object, or it is not, the latter implies that some but not all one's attentional resources are directed toward a certain object); *focal* vs. *global* attention (the former is directed toward a specific object, whereas the latter is distributed over a whole scene); *voluntary* vs. *involuntary* attention (as noted in the previous section, attention can be intentionally drawn to something by the subject or it can be captured by something independently of the subject's will); *endogenous* vs. *exogenous* (the former is internally controlled, whereas the latter is controlled by the external stimulus).

Many have taken the heterogeneity of allegedly attention-related phenomena to show that there is no such thing as what we intuitively call "attention".¹³ Some take an eliminativist stand and argue that "attention" does not refer to any psychologically real phenomenon. Others prefer a disjunctive account, according to which "attention" really refers to a bunch of fundamentally different phenomena. Those who attempt to defend the idea that attention is a psychologically real phenomenon and picks out a natural kind rather than a fragmented collection of disparate phenomena, often appeal to personal-level marks of attention. By "personal-level" I mean "involving the person as a whole", rather than merely one or more of its subsystems (sub-personal states and processes). This may be spelled out in

¹² Watzl (2011: 847) calls the first *The Disunity Problem* and the second *The Overgeneralization Problem*. (Mole 2011b: 63) calls the overall problem the "predicament of explanatory overburdening": "The problem is that the phenomena that attention is expected to explain are too many and too various for any one theory of attention to do justice to them all."

¹³ The idea is well summarized by Allport (1993: 203): "[M]ost contemporary theories of information processing in general, and selective attention in particular, view attention as some sort of causal mechanism. However, even a brief survey of the heterogeneity and functional separability of different components of spatial and nonspatial attentional control prompts the conclusion that, qua causal mechanism, *there can be no such thing as attention*. There is no *one* uniform computational function, or mental operation (in general, no *one* causal mechanism), to which all so-called attentional phenomena can be attributed. On the contrary, there is a rich diversity of neuropsychological control mechanisms of many different kinds (and no doubt many yet to be discovered), from whose cooperative and competitive interactions emerge the behavioral manifestations of attention."

different ways, each of which has its pros and cons (Kriegel 2012). What matters for the present purpose is, roughly, that (i) personal-level phenomena are *conscious* phenomena (unconscious states and processes do not qualify as personal-level phenomena) and (ii) personal-level phenomena are phenomena that are *available* to the subject for deliberate cognizing, reasoning, and action guiding. Accordingly, the notion of attention at play here is *conscious* attention, which is characterized by its phenomenological manifestations and its personal-level functional role. In what follows I present what strike me as the most promising accounts of the phenomenology, the functional role and the metaphysics of attention. I will assume that conscious attention involved in primitive introspection is a personal-level phenomenon and has the features specified by these accounts, regardless of whether *all* attention is to be characterized this way.

Attention participates in the overall phenomenology of one's experience. It shapes the structure of one's overall conscious experience by making some of its aspects central, others peripheral (Kriegel 2009; Watzl 2011a, 2017). The way this structure-shaping affects the phenomenology may be better understood by an example. Consider the auditory experience you have when you listen to a jazz concert.¹⁴ Sometimes you may enjoy your musical experience in its totality; sometimes you may focus on the sound of a particular instrument. The saxophone often occupies the center of your attention, but you may decide to focus first on the bass, then on the piano, then on the drums, and so on. Such a shifting of the focus of your attention makes one sound come to the foreground, while the others retrocede to the background. When you shift the focus from the piano line to the sax line, for example, the overall experience acquires a new configuration such that the sax line becomes the prominent or salient aspect of it, whereas the piano line becomes less prominent, as all the other aspects of the experience. Moreover, it may not be the case that all the aspects other than the one made prominent by the focus of attention lie on one single flat background level. As noted by Watzl (2011a: 156), "the background does not appear to be completely unstructured. In many cases, when you are attending to the saxophone, the sound of the piano is experienced as relevant for or close to the experience of the melody played by the saxophone. By contrast, your pain might be experienced as further out in the periphery or at the fringe of consciousness; you are only marginally aware of your pain." The idea is thus that attention does not simply select an aspect of the experience by making it stand out against a uniform background; it rather organizes the center/periphery structure of the experience in such a way that different aspects of the overall

¹⁴ Both Watzl (2011a: 146) and (Kriegel 2009: 173) use a similar example.

phenomenology may acquire different levels of prominence.¹⁵ To use an evocative metaphor, attention is what provides your experience with “phenomenal depth”: it fixes the focal aspect of the field of consciousness that constitutes the most prominent point around which the rest of the experience is organized at various levels of salience.

Therefore, attention makes a non-negligible contribution to the overall phenomenology of one’s experience.¹⁶ Importantly, it does so even if it does not add any *new element* in the phenomenology. When the (previously silent) sax comes in at the end of the piano solo, a new element is introduced in the overall phenomenology of the listener’s experience. When instead both instruments are playing and the focus of attention is shifted from the piano to the sax, the overall phenomenology changes even though no new phenomenal element is introduced. Attention is thus a structural feature of the phenomenology: it contributes to the phenomenology of the overall experience (it is an *aspect* of it) without being an *item* in it (Kriegel 2009: 172).

The foregoing phenomenological considerations may be called in support of what strikes me as the most promising metaphysical approach to the nature of attention, namely the one adopted by Christopher Mole (2011a, 2011b) and, before him, by F.H Bradley (1886) and Alan White (1964).¹⁷ As noted, one reason why the very existence of attention as a natural kind has been called into doubt is that there seems to be no single cognitive process type which is common to all the allegedly attention-related phenomena and that can thereby be identified as *the* attentional process. Mole argues that both the reductionist and eliminativist project are misled because they are subject to the same category mistake. A theory of the nature of attention should not look for a specific process-type to be identified with attention: as noted, there does not seem to be such a thing. Rather, attention is a *way*, or *manner*, in which mental activities and cognitive processes are carried out. In other words, attention is an *adverbial* phenomenon:¹⁸ “what’s essential to attention is not the facts about which processes are taking place but is, instead, the facts about how the things that happen happen.” (Mole 2011b: 71). Slowness, hastiness, and carefulness are not processes or activities

¹⁵ Arguably, this implies that attention is a gradable, rather than an on/off phenomenon.

¹⁶ These considerations concerning the effects of attention on the phenomenology will be crucial for our discussion of the metaphysics of primitive introspection in Chapter 4.

¹⁷ Bradley and White are cited by Mole as former defenders of his own approach in (2011b: 65-66).

¹⁸ Adverbial phenomena are state- or process-modifiers, rather than being themselves states or processes: they are *ways* states occur or processes are carried out. Compare: preparing breakfast vs. preparing breakfast hastily. Preparing breakfast is a process, hastiness is a process-modifier: it is the way the breakfast-preparing process is carried out.

themselves. Rather, they are ways in which processes and activities can be carried out: one may run *slowly*, pack *hastily*, and ride *carefully*. Similarly, one may carry out a certain cognitive process *attentively*.

Mole's adverbial theory is that attention is a *way* sub-personal cognitive processes occur: there is attention iff a set of cognitive mechanisms occur in unison.¹⁹ However, there is no reason not to apply the adverbial account to *conscious* attention (i.e., attention as a personal-level phenomenon).²⁰ On such an account, conscious attention is not itself a conscious state or process but rather a way in which conscious states and processes occur. This seems to fit the phenomenology of attention pretty well. As noted, even though it does not feature in the phenomenology as an item in it, attention modifies the phenomenology of the overall experience. These phenomenological considerations may suggest that conscious attention is a *phenomenology modifier*: it is the way the center/periphery structure of one's overall phenomenological field is structured.

I am not going to provide an argument for the adverbial theory of the nature of conscious attention. Although it presupposes that attention structures one's overall field of consciousness in the way just described, the account of the nature of primitive introspective states that I am going to provide in Chapter 4 does not rely on the assumption that such a structuring *is* conscious attention (that is, that the nature of conscious attention reduces to the structuring of the field of consciousness). Rather, it is compatible with the structuring being intimately related to conscious attention without being identical to it (e.g. a product of conscious attention or an activity involving conscious attention).²¹

Attention being characterized as a personal-level phenomenon, its functional role cannot be reduced to sub-personal information processing. Declan Smithies (2011) defends the view that the functional role of attention is to make information accessible to the subject for the rational control of action, reasoning, and verbal report. On the assumption that beliefs and actions based on unconscious information are *not* available for such a *rational* control, Smithies' view entails that attention is a personal-level phenomenon (since, recall, in the present framework unconscious states and processes are sub-personal). Attention is thereby characterized as a *mode of consciousness*

¹⁹ "A person is paying attention if and only if, among the set of resources that that person could bring to bear in performing whatever task he is engaged with, there are no resources that are doing anything other than serving that particular task" (Mole 2011b: 67).

²⁰ And indeed this is the kind of account Bradley (1886) and White (1964) put forward.

²¹ For a defense of the view that attention *is* the mental activity of structuring consciousness see Watzl (2017).

(again, on the assumption that personal-level phenomena are conscious phenomena). Smithies is committed to the twofold thesis that attention is sufficient and not necessary for consciousness. There are well-known objections to this thesis, coming from experimental psychology. Against sufficiency, it has been argued that blindsight patients can attend to objects without being phenomenally aware of them (Kentridge, Heywood, and Weiskrantz 1999). Against non-necessity, cases of inattentive blindness are taken to show that absence of attention entails absence of consciousness (Mack and Rock 1998; Simons and Chabris 1999). Smithies argues that all these empirical studies show, if anything, is that some *causal*, non-rational influence on action, reasoning, and verbal report may be necessary and not sufficient for consciousness. Such a causal access being unconscious, it does not yet amount to attention, whose role is to make information *rationally* (rather than merely *causally*) accessible for action, reasoning, and verbal report control.

One worry with Smithies' reasoning is that it seems to characterize attention's functional role in *normative* terms (the "rationally accessible" aspect of it appeals to the notion of *rationality*, which is a normative notion), whereas, arguably, there should be a non-normative way to characterize a phenomenon's functional role—at least on some central ways of thinking about functional role.²² I believe that this challenge may vanish upon a better articulation of Smithies' account. Instead of characterizing the role of conscious attention as that of making information *rationally accessible*, we should rather think of it as that of making information *available to be used* by the subject in carrying out a specific *set of cognitive processes* (i.e. personal-level processes such as conscious cognition, action guidance and verbal report) in a specific *way*, that is, deliberately and with conscious control.²³ From this angle, rationality is not definitional of the functional role of consciousness. Having information available to be used deliberately and consciously, though, implies its being available as an (epistemic) reason for belief or as a (practical) reason for action. The latter bit makes justice to Smithies' idea that information acquired through conscious attention is available to the subject for the *rational* control of action, reasoning and verbal report.

I want my account of primitive introspection to be as neutral as possible as on the nature of attention. Therefore, I am not going to commit myself to Smithies' thesis that *all* attention is conscious attention. Moreover, I do not

²² I owe this objection to Uriah Kriegel.

²³ Recall, the idea that conscious attention makes information available in the way described here came up in our discussion of phenomenal concept acquisition in Chapter 2.

intend to commit to the idea that making information accessible in the way described above is the *only* functional role of conscious attention. More cautiously, I am going to assume that, regardless of whether there is unconscious attention and of what its functional role is, *one of* the functional roles of *conscious* attention is to make information available to the subject for the deliberate and conscious control of personal-level processes such as conscious cognition, action guidance and verbal report.

To sum up, on the notion of attention at work in my account of primitive introspection, attention is a personal-level phenomenon. It is phenomenologically manifest in that it organizes the center/periphery structure of one's overall experience, even though it does not add any new phenomenal element to it. As for its nature, it is an adverbial phenomenon: it is a modifier of conscious states and processes, rather than a conscious state or process itself. Finally, one of its functional roles is to make information available to the subject to be used in deliberate and conscious control of action, reasoning, and verbal report (*personal-level accessible* for short).

Therefore, when you attend to your stomachache experience, in virtue of your attending, the structure of your overall phenomenal field is organized in such a way that the stomachache experience becomes prominent, whereas all the other aspects of your current conscious experience retrocede to the background, possibly at different levels of prominence. Accordingly, introspective attention (i.e. the kind of attention involved in primitive introspection) is an adverbial phenomenon: it modifies your overall phenomenology without being an element in it. By attending to your stomachache experience, some information about its phenomenology, unavailable to you before primitively introspecting, becomes personal-level accessible to you.

3.2. *Apprehension*

In virtue of the act of attention, information about the phenomenology of the target experience becomes personal-level *accessible* (i.e. *available* to be used at the personal level). But in order for the subject to *gain access* to that information (i.e. in order for it to be actually *used*), a mere act of attention is not yet sufficient. To acquire the relevant information, an act of *apprehension* is needed. The latter is a voluntary and effortful act by which the subject (i) refers to the primitively introspected experience and thereby (ii) acquires information about its phenomenology.

In primitive introspection, one cannot apprehend the phenomenology of one's experience without attending to it: apprehension requires attention. It seems, however, that at least in some cases one can apprehend the phenomenology of one's experience even if one does not attend to it: simply having a conscious experience is sufficient to acquire at least some information about its phenomenology. Even though you are concentrated on your work, and the stomachache experience lies at the periphery of your phenomenal field, you might nonetheless have access to some aspects of its phenomenology. For example, you may be aware of the continued presence of that experience (the stomachache is not gone) or notice an increase or decrease in its intensity. Far from being mainstream among philosophers of mind, the idea that merely having a conscious experience is sufficient for apprehending its phenomenology should nevertheless be taken in serious consideration. To be sure, I want my account of the nature of primitive introspection to be neutral about this issue. Those who reject the claim that merely having a conscious experience is sufficient for apprehending its phenomenology might want to stick to the idea that apprehension requires attention. For those who feel attracted by the view that mere conscious experience may involve apprehension, my proposal is to distinguish two distinct notions of apprehension: attentive apprehension and non-attentive apprehension. Both are kinds of *apprehension*: both involve referring to the relevant experience and acquiring information about its phenomenology. However, (i) they do not provide the subject with the same amount of information (attentive apprehension provides you with more and more detailed information than non-attentive apprehension) and (ii) information acquired via attentive apprehension is accessible and available to be deployed by the subject in cognition and action guidance in a way which information acquired via non-attentive apprehension is not. So, even if some information about the phenomenology of the experience could be accessible just in virtue of the subject's merely having the experience (thus independently of the subject's attending to it), there still would be two important differences between the case in which the relevant experience is attended to and that in which it is not. First, when the experience is attended to, more detailed information about the phenomenology is available. This is the reason why, to reply to the doctor's specific questions about what the stomachache feels like, you need to attend to your stomachache experience. Second, when the experience is attended to, the information about its phenomenology is available for the deliberate and conscious control of thought, action, and verbal report, which is not the case when attention is directed elsewhere. Before attending to your stomachache experience the information about its phenomenology is not available for you to report it to the doctor.

The kind of apprehension featuring in primitive introspection is *attentive* apprehension: it involves the subject's attending to the target experience (i.e. the experience which is referred to and information about whose phenomenology is acquired). Accordingly, the act of primitive introspection is an act of *attentive apprehension*. It being an act involving attention, it structures the phenomenal field in such a way that the target experience becomes its most prominent aspect. In virtue of such a structuring the subject apprehends the phenomenology of the target experience: they can refer to it and acquire information about its phenomenology.

4. The target of primitive introspective acts: preliminary considerations

4.1. Restriction on target

Typically, the overall state of consciousness of a subject at a time is extremely rich and complex: at any one time, a considerable number of aspects or goings-on can be distinguished within one's overall conscious experience. Different conscious states, such as perceptual, proprioceptive and algedonic states, as well as states of imagination, emotions, and thoughts may occur simultaneously in one's inner life. In this very moment, for example, among other things, I see the laptop screen before me, I hear cars passing outside, I feel my legs crossed, I feel slightly nervous, and I am thinking about what the best account of the target of a primitive introspective act might be. Some philosophers (e.g. Dainton 2000) provide an atomistic account of simultaneous conscious states of a subject and maintain that these are distinct conscious experiences, which may occur independently of one another. Others (e.g. Tye 2003) argue for a holistic individuation of conscious states, according to which they are rather inseparable parts or aspects of one single overall experience. Although I tend to prefer the holistic individuation of conscious states (Giustina 2017), my account of primitive introspection is compatible with both the holistic and the atomistic view.

Plausibly, all the conscious states simultaneously had by a subject are experienced by them as occurring together.²⁴ However, that all of them can be *introspected* at once is implausible, or, at least, very infrequent.²⁵ For, as noted, introspection is attentive apprehension. And, at least in ordinary cases,

²⁴ This claim is usually part of a stronger thesis—the unity of consciousness thesis—according to which all the phenomenally conscious states of a subject at a time are unified. See Bayne (2010) for a sustained defense of this thesis.

²⁵ Something like introspection of all or most of the conscious states simultaneously had by a subject may occur in some cases of deep meditation. Thanks to Luca Gasparri for pointing this out to me.

one can only attend to a selected aspect of one's current conscious experience. To see this, consider your current overall conscious experience and try to attend to it in its entirety at once. I bet you will not accomplish the task very easily. What I guess you *can* easily do is focus on one particular aspect of your experience (say the tactile sensation of the shoe under your right foot), then on another (say your visual experience of black signs on a white background), then on still another one (say your current emotional experience, whatever it is), and so on. Although this exercise of your attention might make you think that you can introspect your whole conscious experience at once, this is probably an illusion. What you have just done is more likely an introspective 'travel' through different aspects of a conscious experience that is *extended in time*, or alternatively, through specific aspects of different conscious experiences which occur at *different times*.²⁶ In any case, what you *did not* do is introspect all those aspects at once (at the *same time*). Therefore, primitive introspection is, typically, attentive apprehension of an *aspect* of one's current conscious experience.²⁷

It may be that the target of primitive introspection sometimes coincides with one's current overall conscious experience. However, normally, only a part or aspect of the overall experience can be primitively introspected. In any case, the 'size' and complexity of the target of primitive introspection depends on the subject's attentional capacities: it depends on how many aspects of one's current overall experience can be simultaneously at the center of one's phenomenal field (i.e., roughly equally and maximally prominent).

4.2. *Temporality*

As anticipated in the previous subsection, one preliminary question concerns whether what is introspected is (a) an 'aspect' of a temporal part of a temporally extended experience or rather (b) an aspect of an instantaneous experience at a certain instant *t*. This question is closely related to a broader issue, concerning the nature of temporal consciousness: is conscious experience instantaneous or temporally extended? For on the one hand it seems that we can only be aware of what is present. On the other hand, we can be aware of change, movement and succession and when we do, it does not seem to us as though discrete instantaneous experiences occurred one

²⁶ I am going to address this issue in the next subsection.

²⁷ One important question about the nature of primitive introspection is: what are those 'aspects' of conscious experience which primitive introspective attention is directed at? An answer to this question partly depends on one's preferred ontology of conscious experience. I am going to address it in §5.

after the other. Rather, our experience appears to us continuous and temporally extended. A further question concerns the *possibility* of instantaneous experience: is temporally unextended experience metaphysically possible? That is, if a creature were to exist only at instant t (if it were created and immediately destroyed), could it have a conscious experience?²⁸ A negative answer to the latter question would imply that philosophers' talk of 'conscious experience at t ' is nothing but an abstraction, a theoretical device, which does not correspond to any psychologically real phenomenon. If so, then the further question arises about whether such an abstraction can nonetheless successfully explain the phenomena of consciousness.

Those issues are both interesting and important. However, an account of the target of primitive introspection does not really require settling them. You may think that the temporal structure of primitive introspection depends on the temporal structure of its target. Depending on one's preferred account of the temporal structure of conscious experience, you may argue, primitive introspection can or cannot be instantaneous and it can or cannot be temporally extended. For example, on the view that there are only temporally unextended experiences occurring one after the other, primitive introspection must be temporally unextended too. However, this does not seem to be necessarily the case: the temporality of conscious experience and that of primitive introspection may come apart. For even if one maintains that conscious experience is always instantaneous, one can accept that a series of subsequent instantaneous experiences can be introspected through one single introspective state. On the other hand, one who maintains that conscious experience must be temporally extended, may well accept that introspective states are temporally *unextended*. For instance, one may maintain that what is introspected are instantaneous 'aspects' of a conscious experience that by its nature is temporally extended. Therefore, it seems that the two questions (about the temporal structure of introspection and the target conscious experience respectively) are orthogonal.

In what follows, my analysis will be focused on introspective states or events that occur at a certain time t . There are at least two possible understandings of a mental state occurring at a time t . On one interpretation, a mental state occurring at t is an instantaneous state (i.e. a state occurring at instant t). On another interpretation, a mental state can be said to occur at t even though it is slightly extended in time, say one or two seconds (what

²⁸ Thanks to Enrico Terrone for drawing my attention to this question.

is often called the *specious present*),²⁹ or at least the minimal amount of time a mental state must last in order for it to be conscious. The latter interpretation seems to me much more plausible. For if the object of our analysis is *conscious* experience, then, *a fortiori*, we must allow conscious experience (and, by consequence, introspective states) to have at least the duration corresponding to the minimal amount of time a mental state must last in order for it to be conscious. Therefore, in what follows I will assume *t* to be a short undetermined temporal duration, possibly lasting up to two one or two seconds.

5. The target of primitive introspective acts: ontology

Up to now, I have left the ontology of the target of primitive introspection unspecified. To refer to the target, I have used the expressions ‘conscious state’, ‘phenomenal state’, and ‘experience’ interchangeably. In some places, to emphasize that, typically, what is primitively introspected is not the overall experience one has at a certain time, I have used the expression ‘aspect of one’s overall experience’. In this section, I am going to clarify some points about the ontology of the target of primitive introspection, in order to talk about it with a higher level of precision. Particularly, I will consider five possible candidates for constituting the target of primitive introspection: phenomenal-property instances, particulars (bearers of phenomenal properties), phenomenal events, phenomenal states, and phenomenal parts. I will analyze each candidate and specify under what conditions they correspond to genuinely alternative accounts of the target of primitive introspection.

I am not going to argue for the superiority of one account over the others. Rather, I am going to explore different options, and point out their main weaknesses. Rather than defending a specific account of the target of primitive introspection, the purpose of this section is to get a sense of the kind of thing a primitive introspective act is directed at and get an idea of the main challenges faced by different accounts of the target of primitive introspection.

²⁹ The widespread use of the term ‘specious present’ in philosophy and psychology is due to William James (1890) who nevertheless credits E.R. Clay with both the introduction of the term ‘specious present’ and the idea that the ‘sensible present’ has a duration (Dainton 2014).

5.1. *The Phenomenal Property Instance View*

The first option I am going to discuss is what I call the ‘Phenomenal Property Instance View’. On this view, what is primitively introspected is a certain *instance of phenomenal property* of one’s current conscious experience. Phenomenal properties are the subjective or qualitative properties of experience. It is common practice to introduce the notion of phenomenal character of experience by way of Thomas Nagel’s expression ‘what it is like’ (Nagel 1974). What it is like to see the blue sky is different from what it is like to see a yellow banana, and both are different from what it is like to taste papaya. Each conscious state has a certain kind of ‘subjective feel’ which contributes to the overall phenomenal character of one’s experience.

The term ‘phenomenal property’ is often considered as a synonymous of ‘quale’. I prefer not to use the latter expression, though, because of the theoretical baggage often associated with it. Qualia are sometimes characterized as intrinsic, non-representational, non-physical, private, infallibly known, and/or ineffable properties. Although I believe that some of the features usually associated with qualia do apply to at least some phenomenal properties too (e.g. their being intrinsic and not reducible to representational properties), I do not want, at this stage, to *rely on the assumption* that all phenomenal properties have such features. To accept the *Phenomenal Property Instance View*, one can adopt a maximally neutral stand with respect to phenomenal properties, according to which they simply are the phenomenal character of experience.

Eliminativists about phenomenal properties maintain that there is no entity to which the term ‘phenomenal property’ refers. Although eliminativism about phenomenal properties may be defensible on a more committal characterization of phenomenal properties (such as the one mentioned above, according to which phenomenal properties are intrinsic, non-representational, non-physical, ineffable, etc.), it strikes me as inappropriate if a more neutral characterization of phenomenal properties is adopted. For it seems undeniable that conscious experience *does have* a phenomenal character: there *is* something it is like to have a certain experience. This is what I refer to when I use the expression ‘phenomenal property’.³⁰

³⁰ To be sure, some philosophers—so-called ‘illusionists’—deny that there is anything it is like to have a certain experience (see e.g. Dennett 1988, Frankish 2016, and Kammerer 2016). My view and theirs lie on incompatible grounds. So, arguably, accepting the existence of primitive introspection implies rejecting illusionism—a cost I am ready to pay. Besides, I believe no definitive argument can be provided either for or against illusionism. I will briefly come back to this in Chapter 4.

So, on the *Phenomenal Property Instance View*, what is primitively introspected is an instance of phenomenal property. When you contemplate the blue sky while tasting papaya and feeling the wind's brushing your cheeks, (you or)³¹ your experience exemplifies three distinct phenomenal properties. In general, at any one time, a number of phenomenal properties are instantiated (e.g. phenomenal blueness, phenomenal papaya-taste, phenomenal wind-cheek-brushing). On this view, when one primitively introspects, one attends to one of those phenomenal property instances. In the example just described, suppose you focus your introspective attention on the phenomenal blueness which features in your current experience and, in doing so, you do not classify or recognize it as an instance of any experience type. Such an instance of phenomenal blueness is the target of your primitive introspective act.

An issue about phenomenal properties concerns their relationship to experiences and subjects of experiences. Are they properties *of* experiences? Or, rather, are experiences *nothing but* instances of phenomenal properties? If so, what are phenomenal properties properties of? Several options are available. I limit myself to listing some of them. One is that, at any one time, a subject has one overall experience which possesses a number of phenomenal properties. Another option is that at any one time a subject has a number of experiences, or conscious mental states, each of which possesses one (or more) phenomenal properties. On both options, experiences are particulars and phenomenal properties are universals that can be exemplified by different experiences. (The metaphysical ingredients of this view are roughly the same as what I will call the 'Particular View' in the next subsection). To adopt this view, one would need to specify what kind of particulars experiences are, which may turn out not to be an easy task.

To avoid this difficulty, one may maintain that experiences *are* phenomenal-property instantiations, or phenomenal *tropes* (they are abstract particulars). On one version of this view, at any one time a subject has a number of experiences, each of which is a phenomenal-property instantiation. On another version, at any one time a subject has one single experience which is constituted by a bundle of phenomenal-property instantiations.³² For both views, the main challenge consists in explaining in virtue of what an ensemble

³¹ I will come back to the issue whether it is you or your experience that exemplifies phenomenal properties in the next paragraph.

³² Arguably, there is no substantial difference between the two views: they seem to only disagree about what should be called 'experience', the phenomenal-property instantiation or the bundle of phenomenal-property instantiations.

of phenomenal-property instantiations is unified and belongs to one and the same subject.

A related issue, which I just mention here, concerns the relationship between experiences and subjects. Does a subject *have* an experience? Is the subject the substratum that instantiates phenomenal properties? Or is the subject just a bundle of phenomenal property instantiations? Unfortunately, I cannot discuss these issues here. However, it is important to bear in mind that the account of primitive introspection I am going to provide is compatible with different answers to these questions and that different versions of the view can be developed depending on one's preferred metaphysics of subjects and experiences.

5.2. *The Particular View*

According to what I call the 'Particular View', experiences are particulars, bearers of phenomenal properties. The metaphysics underlying the *Particular View* is similar to that underlying the first version of the *Phenomenal Property Instance View* outlined above, i.e., the view according to which experiences (intended as particulars) are what exemplifies phenomenal properties. On this view, primitively introspecting your papaya-taste experience implies performing a primitive introspective act directed toward an experience (a particular), which exemplifies (among others) the papaya-taste phenomenal property.

One may spell out the *Particular View* in at least two different ways. The first option is that at any one time a subject has one single experience which instantiates a number of phenomenal properties (we may call this the 'One-Particular View'). The second option is that at any one time the subject has many experiences, each of which exemplifies one or more phenomenal properties (call this the 'Many-Particulars View').

As for the first version of the *Phenomenal Property Instance View*, the main challenge for the *Particular View* is to spell out what kind of particulars experiences are. Moreover, each the two specifications of the view meets a challenge.

The *One-Particular View* has to explain how it is possible for one's overall experience to be the target of primitive introspection. As noted in §4.1, what can be primitively introspected is typically restricted to an aspect or portion of one's overall experience—one typically does not introspect one's overall experience at once. On the *One-Particular View*, however, at any one time the subject has one single experience (which, arguably, coincides with

one's *overall* experience) and *that* is the target of primitive introspection. The challenge for the *One-Particular View* is thus to solve this (apparent) tension.

The *Many-Particulars View*, on the other hand, needs be implemented with an explanation of the unity of consciousness. As noted in §4.1, at any one time one's conscious experience is phenomenally unified. Yet, if at any one time a subject has many distinct conscious experiences, what explains the fact that such experiences are phenomenally unified? On the *One-Particular View*, this unity is explained by the fact that at any one time a subject has *one single* overall experience (which exemplifies a number of phenomenal properties). By contrast, on the *Many-Particulars View*, such a straightforward explanation is not available. Unity is to be explained otherwise, perhaps by appeal to a certain relationship among simultaneous phenomenal events, which ties them together in a special way.³³ The challenge the *Many-Particulars View* faces is that of providing a satisfying and sufficiently explanatory story about such a relationship.

5.3. *The Phenomenal Event View*

Another option is to take the target of primitive introspective acts to be *phenomenal events*. On this view (call it the 'Phenomenal Event View'), one's stream of consciousness is constituted of a number of phenomenal events, like seeing the blue sky at eight o'clock, or feeling hungry at noon. Those phenomenal events which occur simultaneously constitute one's current overall conscious experience. Accordingly, primitively introspecting involves attending to one of these phenomenal events. Such a view might take at least two different forms, depending on one's metaphysical account of mental *events*. According to Jaegwon Kim, an event is an exemplification of a property at a time (Kim 1993). On the Kimean model, thus, events are *structured*: they are constituted by an object, a property and a time. An event $[x, P, t]$ exists iff the object x exemplifies the property P at time t . Therefore, on this view, events are reduced to more fundamental ontological categories—properties and objects. According to Donald Davidson, instead, events are unrepeatable individuals, which are not reducible to more fundamental constituents. They are thus *sui generis* particulars which constitute a fundamental ontological category (Davidson 1969). Davidson's first proposal was to individuate events by their unique position in the chain of causes and effects: two events are identical iff they have the same causes and effects. However, acknowledging

³³ Dainton (2000) has a theory of the unity of consciousness along these lines, according to which the different experiences a subject has at t are unified in virtue of their being co-conscious (where co-consciousness is a special relation they bear to each other).

the circularity of such an account (an event's causes and effects are themselves events), he modified his identification criterion for events and maintained that two events are identical iff they have the same spatiotemporal location (Davidson 1985).³⁴

Arguably, a Kimean *phenomenal* event is an exemplification of a *phenomenal* property at a time. A Kimean phenomenal event $[x, P, t]$ exists iff the mental particular x exemplifies the phenomenal property P at time t . It is an open question what exactly the mental particular which exemplifies such a property is. It is often taken to be the *subject*, but it might also be the subject's current overall experience, one of one's current conscious states, or perhaps something else. In any case, on a 'Kimean Event View' ('Kimean View' for short) of the target of primitive introspective acts, what is primitively introspected is a phenomenal property exemplification at a time. Now, you may have noticed that this version of the *Phenomenal Event View* sounds quite similar to the *Phenomenal Property Instance View*. Indeed, as stated above, the latter says that what is primitively introspected is an instance of phenomenal property. As noted, one framing of the *Phenomenal Property Instance View* is that at any one time an experience (or a mental state, or a subject) exemplifies one or several phenomenal properties. Kimean events do not constitute a fundamental ontological category, but are reducible to more fundamental entities (properties and particulars). It thus seems that the ontology underlying the *Kimean View* is the same as that underlying this version of the *Phenomenal Property Instance View*. If so, then the two views involve exactly the same metaphysical ingredients, though make slightly different predictions about the exact target of a primitive introspective act: on the *Phenomenal Property Instance View* the target is a phenomenal-property instance; on the *Kimean Event View* the target is the instantiation of a phenomenal property at a time. As before, the main challenge for the *Kimean View* will be to specify what exactly the substrate that exemplifies the phenomenal properties is.

The *Davidsonian Event View* ('Davidsonian View' for short), on the other hand, does sound like a genuine alternative to the *Phenomenal Property Instance View*. On this view, phenomenal events which constitute the target of primitive introspective acts are irreducible and unrepeatable mental particulars. Two are the main challenges to the *Davidsonian View*.

The first hinges on the idea, mentioned above, that one's conscious experience is somehow unified at any one time. On the *Phenomenal Property*

³⁴ Both the Kimean and the Davidsonian view have pros and cons, which I am not going to discuss here. For a discussion of them see Schneider (2016).

Instance View, this unity can be explained by the fact that at any one time a subject has one single overall experience (which exemplifies a number of phenomenal properties). As we saw in the previous subsection, the same solution is available to the *One-Particular View*. Similarly, the *Kimean View* can accommodate phenomenal unity by maintaining that an ensemble of phenomenal events is unified in virtue of there being one single particular (e.g. one's overall experience, or the subject) which exemplifies different phenomenal properties at the same time. By contrast, the *Davidsonian View* (as the *Many-Particulars View*) does not have an equally straightforward explanation of phenomenal unity. The *Davidsonian View* thus owes us an explanation of how several different Davidsonian phenomenal events stich together and form one single unified state of consciousness.³⁵

The second challenge has to do with the *individuation* of Davidsonian phenomenal events. As mentioned before, Davidson's original proposal that events are individuated in terms of their causes and effects is threatened by circularity. What about his later proposal? The latter is that events are individuated by their position in the spatiotemporal framework. A *prima facie* worry concerns which exactly is the spatial location of phenomenal events. Here is Davidson's suggestion:

I think we do assign a location to such an event when we identify the person who remembered, decided or solved: the event took place where the person was. [...] Mental events (by which I mean events described in the mental vocabulary, whatever exactly that may be) are like many other sorts of events, and like material objects, in that we give their locations with no more accuracy than easy individuation (within the relevant vocabulary) demands. Aside from a few dubious cases, like pains, itches, pricks and twitches, we have no reason to locate mental events more precisely than by identifying a person, for more

³⁵ I am wittingly leaving aside, here, the issue about whether, on the *Phenomenal Property Instance View*, phenomenal properties are instantiated by the subject's experience or by the subject themselves, and simply assuming the former. This is not because I consider the issue unimportant (on the contrary), but because bringing it in would lead us too far from the focus of the present section. Taking a stance about the nature of the self surely does affect one's view of the unity of phenomenal consciousness and thorough reflection should be devoted to the ways both issues interlace. At any rate, whether one takes experiences or selves as fundamental, any ontology of consciousness implying a number of distinct mental particulars (experiences, events or states) at *t* will face the problem of explaining how those relate. Such a problem does not necessarily affect the Phenomenal Property Instance View or the Kimean View, for one may take phenomenal properties to be instantiated either by one single unified experience or by one single unified self. (Of course, if one takes phenomenal properties to be instantiated by different simultaneous mental particulars, one's view does face the relevant problem.)

than this would normally be irrelevant to individuation. (Davidson 1969: 304)

Davidson's view is that the location of phenomenal events is wherever their subject is. However, this view seems to face a further difficulty. By defending his previous causal account of the individuation of events, Davidson himself suggests the following objection against the spatiotemporal account:

[I]t seems natural to say that two different changes can come over the whole of a substance at the same time. For example, if a metal ball becomes warmer during a certain minute, and during the same minute rotates through 35 degrees, must we say these are the same event? (Davidson 1969: 306)

Intuitively, this does not seem to be the case. The objection appears even more serious when it comes to phenomenal events. For if an event is individuated by its spatiotemporal location and a phenomenal event's spatiotemporal location coincides with that of its subject, then all the phenomenal events of a subject at a time are in fact the same event. But this does not seem to make justice to the fact that at any one time we can *distinguish* and *individuate* different goings-on in our overall phenomenal field. A Davidsonian could argue that for every subject there really is one single phenomenal event at a time (this is what accounts for synchronic unity of consciousness), and that what accounts for the fact that different goings-on can be distinguished is that phenomenal events have a complex internal structure.³⁶ Of course, the main challenge for this view will be to show how a phenomenal event is internally structured and how this can account for distinguishability.

Another option is to give up the idea that a phenomenal event's spatiotemporal location coincides with that of the subject and opt for a more fine-grained story about phenomenal events, according to which their location coincides with that of their neural correlates: a certain phenomenal event occurs wherever its neural correlate is located. A consequence of this view is that, being individuated by exactly the same spatiotemporal location, the phenomenal event and the neural event are in fact one and the same event. This is good news for physicalists. One potential worry for this view, however, may be that the same phenomenal-event type could have been realized by a different neural-event type, and different neural event types may have different spatiotemporal locations. Therefore, the spatiotemporal location of a phenomenal event is contingent to it, so it seems to fail to provide the event's individuation condition.

³⁶ Something similar to what Tye (2003) argues about experiences.

5.4. *The Phenomenal State View*

A somewhat related account of the target of primitive introspective acts is the ‘Phenomenal State View’, according to which the target of primitive introspective acts is a *phenomenally conscious state* like seeing a tree, tasting chocolate, and feeling pain in one’s knee. On this view, a subject has a number of distinct phenomenal states at any one time and primitively introspecting involves attending to one of them. Whether the *Phenomenal State View* differs from the *Phenomenal Event View* depends on one’s conception of a *state*. If one takes states to be a particular kind of event, then the two views are identical. If, instead, one resists the idea that states are events (perhaps on the basis of the intuition that whereas events involve change, states do not, and change plays a substantial role in the definition of one’s ontological categories), the two views are distinct.

Even in the latter case, though, some analogies can be identified between the two views. Particularly, a distinction can be drawn between two versions of the *Phenomenal State View*, which retraces the above distinction between the Kimean and the Davidsonian Event View. According to the Kimean version of the *Phenomenal State View*, phenomenal states are instantiations of phenomenal properties. The difference with the Kimean Event View is that whereas an event consists in an object’s *instantiating* a property (which involves change), a state consists in an object’s *having* a property (which does not involve change). Similarly to the *Kimean Event View*, this *Kimean State View* seems to collapse into a version of the *Phenomenal Property Instance View*. On a Davidsonian version of the *Phenomenal State View*, by contrast, phenomenal states are particulars. Both the Kimean and the Davidsonian version of the *Phenomenal State View* face similar challenges to their phenomenal-event counterparts.

5.5. *The Mereological View*

Finally, one might take a *mereological* approach, and maintain that what is introspected is a certain *part* of one’s current conscious experience (‘Mereological View’). On this view, at any one time one has one overall experience in which a certain number of parts can be distinguished. Seeing the screen before me and hearing the cars passing outside, for instance, are phenomenal parts of my current overall experience.³⁷ One challenge for the

³⁷ My current experience may have also other kinds of parts, for example physical parts (i.e. the parts of its neural correlates), or intentional parts (i.e. parts of its intentional content).

Mereological View is to specify how the phenomenal parts that are primitively introspected are individuated: exactly what parts are primitively introspected? My seeing the whole scene before me (which includes my laptop's screen at the center, a bottle and an apple on the right, a pencil, a sheet of paper and a lamp on the left), my feeling the sole of the shoe under my right feet, and my hearing cars outside are phenomenal parts of my current experience. But so are my seeing the screen before me, my seeing an apple on the right, and my seeing a lamp on the left, which are parts of my seeing the whole scene before me; also my feeling the sole under my toes and my feeling the sole under my heel are parts of my current experience, they being parts of my feeling the sole under my feet; and so is my feeling the sole under my big toe; and so on and so forth. Now, which of these parts do I introspect when I primitively introspect? When I draw my introspective attention to my right foot, is it feeling the sole under the foot that I primitively introspect, or rather my feeling the sole under my toes, or maybe under my big toe?

As noted in the previous section, introspective attention normally cannot be directed at one's whole overall experience. Now, it seems that, if what is introspected are experiential parts, such parts should be sufficiently 'small' for the subject's attention to be directed at them. For instance, it seems that one's overall visual experience, or one's overall proprioceptive experience usually cannot be attended to by way of one single act of attention (perhaps they can by way of a number of subsequent acts of attention). So, the target of primitive introspection should be a subpart of them. But, one might ask, exactly *how* 'small'? Requiring the *Mereological View* to determine the *exact* 'size' an experiential part must be in order for it to be primitively introspected is perhaps too demanding. A reasonable answer may be: the target of a primitive introspective act is the part of one's current experience toward which one can direct one's attention, however 'big' this part might be. The 'size' of the primitively introspected experiential part, therefore, depends on the power and limits of one's introspective attention. It is just because our introspective attention usually cannot encompass our current overall visual experience, nor our overall proprioceptive experience, that those typically cannot be the target of a primitive introspective act. What experiential parts are primitively introspected is thus a contingent matter, which depends on one's introspective attention capacities and may vary from subject to subject, as well as for the same subject across time.

The candidate for the target of primitive introspective acts under consideration here is *phenomenal parts*, i.e. parts of the experience's phenomenology.

A non-negligible task for the defender of the *Mereological View* will be to define the relationship between the overall experience and its parts. Some of the issues they will need to consider are the following. Are the parts separable from the whole? That is, can they exist independently of the overall experience they are part of? Or does their existence depend on that of the overall experience? Is the phenomenology of the overall experience obtained by composition of that of its parts? Or is it something over and above (thus not reducible to) the phenomenology of the parts? If the latter, how does this interact with the fact that primitive introspection is usually directed to a phenomenal part of the experience, rather than to the overall experience as a whole? In other words, how does the phenomenology of the overall experience affect the phenomenology of the part which is primitively introspected? And how does it affect the phenomenology of the primitive introspective state? These and probably many other questions should be answered for the *Mereological View* to be a complete account of the target of primitive introspective acts.

*

I have considered different possible accounts of the target of primitive introspective acts and pointed out the main challenges they face. As noted, the purpose of this section was not to settle the question about the ontology of the target of primitive introspection. Rather, by exploring different candidates, my purpose was to help get a sense of what a primitive introspective act can be directed at. My theory of primitive introspection is compatible with all the mentioned accounts. In what follows, I will more often use the expression ‘phenomenal state’ to refer to the target of primitive introspection. This is just for sake of simplicity and should not be read as a commitment to a particular account of the target of primitive introspection. The reader should always bear in mind that what I call ‘phenomenal state’ may be a phenomenal property instance, as per the *Phenomenal Property Instance View*, a mental particular (a bearer of phenomenal properties), as per the *Particular View*, the instantiation of a phenomenal property by a subject (or by an experience) at a time, as per the *Kimean Event View*, a Davidsonian phenomenal event, as per the *Davidsonian Event View*, a Davidsonian phenomenal state, as per the *Davidsonian State View*, or a phenomenal part, as per the *Mereological View*. In Chapter 4 I will briefly come back to this and suggest how my account of the metaphysics of primitive introspection could

fit these different views of the ontological status of the target of primitive introspection.

Conclusion

In this chapter, I have prepared the ground for an account of the metaphysics of primitive introspection. I have specified the necessary conditions of primitive introspection, namely *consciousness*, *presence*, *attention*, and *apprehension* and suggested that primitive introspection is best seen as a *process* involving at least three elements: the *act*, the *state* and the *target* of primitive introspection.

I have characterized the act of primitive introspection as an act of *attentive apprehension*. Conscious attention has been characterized as a personal-level phenomenon, which is phenomenologically manifest in that it organizes the center/periphery structure of one's overall experience. As for its nature, it is an adverbial phenomenon: it is a modifier of conscious states and processes, rather than a conscious state or process itself. Conscious attention makes information available to the subject to be used in deliberate and conscious control of action, reasoning, and verbal report. Accordingly, it enables the subject to *apprehend* the target phenomenal state, that is, to refer to it and acquire information about its phenomenology.

Finally, I have explored different ways the target of primitive introspection may be characterized. The account of the metaphysics of primitive introspection that I will develop in the next chapter is compatible with all the aforementioned views of the specific ontological status of its target.

CHAPTER 4

THE INTEGRATION ACCOUNT OF THE METAPHYSICS OF PRIMITIVE INTROSPECTION

In Chapter 3 I made some preliminary remarks about the nature of primitive introspection. In this chapter, I am going to dig into its metaphysics. My principal aim is to develop a version of the acquaintance account to explain the metaphysical structure of the mental state the subject is in when they primitively introspect.

An account of the metaphysics of primitive introspection is an account of the relationship between the act, the target, and the state of primitive introspection. Most fundamentally, it describes what the activity of primitively introspecting consists in and how the state of primitive introspection relates to the target phenomenal state. In Chapter 3 I explained what the act of primitive introspection involves. Here I focus on the relationship between the target and the state of primitive introspection. I argue that a satisfactory account of the metaphysics of primitive introspection must fulfill two desiderata (§1): what I call the *no introspective hallucination desideratum* (there can be no hallucination in primitive introspection) and the *phenomenal modification desideratum* (primitive introspecting a phenomenal state modifies some aspects of its phenomenology). I show that none of the best available theories satisfies both desiderata. The *inner sense* theory of introspection may have the resources to explain *phenomenal modification*, but it cannot satisfy the *no introspective hallucination* desideratum (§2). The *acquaintance* theory does fit the *no introspective hallucination* desideratum. However, extant versions of the acquaintance theory do not explain *phenomenal modification* (§3). I propose my own version of the acquaintance view, what I call the *integration account of primitive introspection*, as an improvement on extant acquaintance accounts (§4). I show that the integration account satisfies both desiderata. I take this to be a compelling reason to prefer the integration account over its competitors as an account of the nature of primitive introspection.

To be sure, I do not take the following to show that the integration account is the *sole* satisfactory account of the metaphysics of primitive introspection. Other equally satisfactory accounts may be developed, provided that they satisfy both desiderata. Here I limit myself to argue that, given that

it satisfies both desiderata, the integration account is one such satisfactory account.

*

1. Two desiderata for an account of the metaphysics of primitive introspection

An account of the metaphysics of primitive introspection should explain the relationship between the state and the target of primitive introspection. A satisfactory account of the metaphysics of primitive introspection must fulfill two desiderata: what I call the *no introspective hallucination desideratum* and the *phenomenal modification desideratum*. Let us consider each of them in turn.

1.1. The no introspective hallucination desideratum

In perception, there is the possibility of hallucination. One has a perceptual hallucination when it perceptually seems to one as if there is an object before one, whereas no object is actually there. If, for instance, by effect of direct brain stimulation, you have a visual experience as of a pink rat although nothing is before you, you are undergoing a visual hallucination as of a pink rat. My claim is that introspection is fundamentally different from perception in this respect. In introspection, it is never the case that it introspectively seems to one as if one has a certain phenomenal state whereas no phenomenal state is actually there.¹

The no-introspective-hallucination claim is distinct from and does not entail the no-introspective-*illusion* claim, i.e. that introspection can never *mischaracterize* an introspected phenomenal state. One undergoes an introspective illusion when it introspectively seems to one as if phenomenal state ϕ has the phenomenal property P , whereas ϕ does not have P , or when it introspectively seems to one as if ϕ lacks Q , whereas ϕ does have Q . Thus, in introspective illusion, introspection misattributes properties to a phenomenal state that *is* present (which makes introspective illusion crucially

¹ Arguably, the no-introspective-hallucination claim has been part of the orthodoxy in philosophy, at least since Descartes. Recently the claim has been called into question (e.g. Dennett 1988), but it still has a substantial number of defenders (among which Hill 1991, Pitt 2004, Horgan and Kriegel 2007, Gertler 2011).

different from introspective hallucination). Different from introspective hallucination, introspective illusion *does* sometimes occur (one may misattribute a certain phenomenal property to one's current experience).

The no-introspective-hallucination claim is a metaphysical-necessity claim: it is *metaphysically impossible* for one to have an introspective state as of a certain phenomenal state but no phenomenal state be actually there.² For a theory of primitive introspection to satisfy the no-introspective-hallucination desideratum, then, it is not enough that it entail that, in the actual world, there cannot be introspective hallucination (say, due to how human cognitive system is built). It must entail that, in all worlds in which there is primitive introspection, it is never the case that one has a state of primitive introspection as of a certain phenomenal state but no phenomenal state be actually there.³

Providing a definitive argument in favor of the no-introspective-hallucination claim exceeds my present purpose. (Indeed, the main purpose of this chapter is *not* to provide a sustained argument for the two desiderata but rather, given the desiderata, develop an account of the metaphysics of primitive introspection which satisfies them.) I will nonetheless motivate it before moving on. On my view, the main motivation for endorsing the claim comes from strong intuitions about introspection. Denying the no-introspective-hallucination claim would imply that it is possible that it introspectively seems to one that one has an experience although one is having *no* experience at all. But this seems simply inconceivable. I may well be wrong about *which phenomenal properties* this experience (which I am introspecting) has. But how can I be wrong about there being an experience at all? To see why the impossibility of hallucination in primitive introspection is a *metaphysical* impossibility, consider the following reasoning. As noted in §2 of Chapter 3, the state of primitive introspection is a conscious state with phenomenology. But being a conscious state with phenomenology entails being a phenomenal state. Therefore, in every possible world where there is primitive introspection, being in a state of primitive introspection entails being in a phenomenal state.

Incidentally, I believe that *no* definitive argument for (or against) the no introspective hallucination claim can be provided. Consider illusionism about phenomenal consciousness, i.e. the theory that, although it introspectively seems to us that we are phenomenally conscious, we are not (see e.g. Dennett

² In other words, it is metaphysically impossible that it introspectively seems to one as if there is a certain phenomenal state but no phenomenal state is actually there.

³ This specification will be relevant to the evaluation of the inner sense account in §2.2.

1988, Frankish 2016, and Kammerer 2016). According to illusionists, not only introspective hallucination is possible—we undergo introspective hallucinations *all the time*: it is always the case that it introspectively seems to us as if there is a phenomenal state although no phenomenal state is there. (A terminological note: in the context of the conceptual distinction between introspective *illusion* and introspective *hallucination* just drawn, the label ‘hallucinationism’ would probably be more appropriate than ‘illusionism’. I stick to ‘illusionism’, though, to be consistent with the terminology most commonly used in the literature.) Now, the primary reason for the realist about phenomenal consciousness to reject illusionism is that it is not only hard to believe, but simply *inconceivable* that phenomenal consciousness is just a hallucination—that although it seems to me that phenomenal consciousness is present it really is not. Phenomenal consciousness is *right here*, in *this* experience: how could one seriously claim it does not exist? Illusionists, on the other hand, maintain that not only phenomenal consciousness being an illusion is conceivable—it is the way things actually are. The widespread resistance to the idea that phenomenal consciousness is an illusion, the reasoning goes, is due to the extraordinary powerfulness of the illusion itself. On the illusionist’s view, the realist’s inconceivability claim does not constitute a threat to illusionism—rather, it is predicted by the theory (Kammerer 2016). On the realist’s view, by contrast, the best explanation of such inconceivability is that there can be no introspective hallucination.

Realists and illusionists therefore ground their theories on incompatible fundamental claims: that there cannot be introspective hallucination on the one hand and that there can be on the other. I suspect that neither claim can be defended in a non-question-begging manner—although I wish I would be wrong on this point. For there does not seem to be any antecedently plausible truth—a more fundamental truth than ‘I am conscious’ or ‘I am not conscious’—on the basis of which the dispute could be settled. Hence my skeptical position about the possibility of a definitive argument for the no-introspective-hallucination claim.

Even though a definitive argument for the no-introspective-hallucination claim cannot be provided, I still think that there are no good reasons to resist it. A consideration which is sometimes adduced against the no-introspective-hallucination claim is that there are at least some cases in which we are wrong about our current experience: a certain experience appears a certain way to introspection, but we then realize that the experience was actually different from how it appeared. If we can be wrong about some

properties of experience, why could we not be also wrong about whether there are phenomenal states at all?

An oft-cited case is the following.⁴ A college student is being initiated to a fraternity. He is told that, as part of his initiation, he will be blindfolded and a spot on his throat will be cut with a razor. However, after he is blindfolded, no cut is drawn on his throat; an ice cube is put on it instead. As soon as he feels a sensation, the student cries out, convinced that he is having a pain sensation. However, after a few moments, he realizes he was wrong: no pain sensation has ever been there, for all he has had all along is a coldness sensation. Therefore, the reasoning goes, throughout the first split second, the student has an introspective state as of a certain experience (i.e. pain sensation) which is actually not there. However, what the fraternity initiation case shows, if anything, is that the student may misattribute a certain phenomenal property to his experience—i.e. that he may be wrong about *which* sensation he is introspecting. It does not show that the student has an introspective state as of a certain sensation (a pain sensation) although no sensation at all is actually there, for there still is a sensation the student introspects, namely the coldness sensation.

So, even if cases like the fraternity initiation show that one could be wrong about *which* sensation one is introspecting, they do not show that one cannot introspect a sensation without *any* sensation being felt.

1.2. *The phenomenal modification desideratum*

As noted in Chapter 3, primitive introspection requires conscious attention. We are going to see throughout the course of this subsection that conscious attention affects the phenomenology of one's experience. Therefore, when it is introspected, a phenomenal state undergoes a change in its phenomenology—a change which is due to the subject's drawing their introspective attention toward it.

As already mentioned in Chapter 3, a study of the effects of conscious attention on the phenomenology of experience has been put forward by Uriah Kriegel (2009, Ch. 5) and, more recently and more thoroughly by Sebastian Watzl (2017, Ch. 8 and Ch. 9). Consider the auditory experience you have when you listen to a jazz concert.⁵ You may first enjoy your musical experience in its totality, in which case your attention is roughly equally

⁴ The case is attributed to Roger Albritton by Christopher Hill (1991: 128).

⁵ The music example is put forward by both Kriegel and Watzl.

distributed across all the instruments. Then, you may decide to focus your attention on the sound of, say, the saxophone. Later, you may shift the focus of your attention toward the sound of the piano. Intuitively, such a shifting of the focus of your attention makes a difference in the phenomenology of your auditory experience. The way it is like to you to listen to the music changes partly depending on what aspect of the music you attend to. Or consider a meditation session in which you are asked to focus on different parts of your body in turn. You attend to your feet, then to your ankles, your knees, your pelvis, and so on and so forth through all the parts of your body. Again, each shift of attention implies a phenomenological change.

That attention makes a difference to the phenomenology of one's experience seems intuitively true.⁶ Indeed, this was an important issue for early experimental psychologists. E.B. Titchener (1912), for example, while spelling out the introspective method distinguishes between "free consciousness" and "controlled consciousness":⁷

A conscious state or process is free when it is neither evoked nor influenced by the intent to observe; it is controlled when it arises under the influence of an introspective intent and as the object of a consequent attention especially directed upon it. (Titchener 1912: 493)

Drawing one's attention to a certain conscious state or process affects its phenomenology. For this reason, the introspectionist psychologist must bear in mind that attended and unattended conscious states are to be treated differently and that "we have no right to generalise a priori from the controlled to the free" (*ibid.*).

So, that attention somehow affects the phenomenology is widely accepted. However, *how* the effect of attention on the phenomenology is to be spelled out is not straightforward. For one thing, as Watzl (2017: 161) points out, there is empirical evidence that attention has different kinds of effect on *visual* appearances (Carrasco, Ling, and Read 2004). As Carrasco and her colleagues have shown, attention affects (at least) apparent contrast, apparent spatial frequency, apparent size, apparent color saturation, apparent spatial relations, and apparent duration in visual experience.⁸ Visually attended objects appear to have higher contrast than unattended objects; they also look bigger, more saturated, closer, and faster.

⁶ Watzl (2017:155) claims that that attention affects the phenomenology of experience is even "uncontroversial". I think that his claim is a bit too strong though.

⁷ He borrows the distinction from Müller (1911).

⁸ I refer to Watzl's (2017: 161) discussion here.

Some philosophers and early experimental psychologists have pointed at a more general effect of attention on appearances (an effect that is not restricted to *visual* appearances): attended things appear clearer, more defined. One way to spell this out is in terms of determinable vs. determinate properties (Watzl 2017: 168-69).^{9;10} The idea is that attention affects *how determinate* things look. By attending to a certain feature, one discerns it in a finer-grained way. When you attend to a leaf on a tree, its features (color, shape, size) are experienced by you more determinately than those of the unattended leaves—this, arguably not just in virtue of the relevant leaf occupying the foveal area of your visual field: attending to objects in the periphery of the visual field makes them look more determinate too (although probably not as determinate as if they were both attended to and occupying the foveal area). When you attend to the saxophone sound at the jazz concert, the sound appears more determinately to you than when you do not attend to it: among other things, you can distinguish its pitch and timber in a finer-grained way. Similarly, when you attend to a certain part of your body, you proprioceptively experience some of its features (shape, weight) more determinately than those of unattended body parts.

Another effect attention sometimes has on the phenomenology is modification of the *intensity* of a certain experience. At the jazz concert, one of the effects of shifting your attention to the sound of the saxophone may be that the sound becomes a bit louder. At the meditation session, when you draw your attention to your right foot, the foot may appear to you a little bit heavier. Sometimes, when you attend to a pain sensation of yours, say a headache, it may occur that by attending to it the pain becomes slightly more intense—a little bit more painful. Or if you try to attend to the phenomenology of a certain emotional experience of yours, say, boiling anger, it may occur that drawing your attention to it somewhat decreases its intensity—the anger cools down a bit.¹¹

As both Kriegel (2009) and Watzl (2017) observe, the effect of attention on the phenomenology is not limited to modifying the way things appear to

⁹ Watzl refers to the work of Nanay (2010) and Stazicker (2011) for an account along these lines. These philosophers take experience of more determinate properties to be what *constitutes* the phenomenology of attention. However, we need not accept such a strong view here: for the present purpose it is sufficient to point out that experience of more determinate properties is an effect that attention sometimes has on the phenomenology of experience.

¹⁰ As for the determinable-determinate relationship consider the following: red is a determinate of the determinable color; crimson is a determinate of the determinable red.

¹¹ Cf. Brentano (1874: 30): “If someone is in a state in which he wants to observe his own anger raging within him, the anger must already be somewhat diminished, and so his original object of observation would have disappeared.”

the subject—it is not limited to modifying the qualitative aspects of one’s experience in the ways specified above. Perhaps the most fundamental effect of attention on conscious experience is that it shapes the experience’s *structure*: it organizes one’s overall conscious experience in such a way that some of its aspects become more central, other more peripheral. Becoming more central or more peripheral has a direct effect on *salience*: the more central, the more salient. Recall what we saw in Chapter 3 considering the jazz example. When you shift the focus from the sound of the saxophone to that of the piano, the overall experience acquires a new configuration such that the piano occupies the center of your phenomenal field and thereby becomes the most prominent or salient aspect of it, whereas the sax becomes less prominent, as all the other aspects of the experience. The other, less salient aspects of your experience do not lie on one single flat background level: the background is itself structured. The sound of the sax, as that of the other instruments, is experienced as more salient than, say, the pain in your knee, for they are closer, in terms of relevance, to the sound of the piano. Attention thus organizes the center/periphery structure of the experience in such a way that different aspects of the overall phenomenology may acquire different levels of prominence. To repeat the metaphor I used in Chapter 3, attention provides your experience with “phenomenal depth”: it fixes the focal aspect of the field of consciousness that constitutes the most prominent point around which the rest of the experience is organized at various levels of salience.

Whereas at least some of the abovementioned phenomenological effects of attention on how things appear to the subject are contingent—they do not necessarily occur at each attention focusing or shifting,¹² shaping the center-periphery structure of conscious experience does occur at every instance of attention focusing or shifting: it is a (nomologically) necessary effect of attention on the phenomenology. Therefore, every time one attends to a certain object or to a certain aspect of one’s experience, such object or aspect becomes more salient, whereas other aspects become less salient.

¹² As mentioned above, people like Nanay and Stazicker think that the phenomenology of attention is exhausted by increased clarity (understood as experience of more determinate properties). If they are right, then change in clarity is a (nomologically) necessary effect of attention on the phenomenology. Some may also argue that any instance of attention focusing or shifting entails a (however slight) intensity modification. Watzl (2017: 171-181) develops two arguments against the view that the phenomenological effect of attention reduces to change in appearance properties. As I point out in the next paragraph, I remain neutral on whether the change in phenomenology associated with salience reduces to change in some appearance properties (clarity, determinacy, intensity or else).

In the specific case of *introspective* attention, every time one introspectively attends to a certain phenomenal state, such phenomenal state becomes *more salient*—at least in some cases, it may become the *most* salient aspect of one’s phenomenal field. Moreover, at least sometimes, introspective attention also modifies some appearance properties (i.e. properties associated with how things appear to the subject in experience) of the target phenomenal state. It may modify its *intensity*, as when introspecting a pain sensation makes it more painful. It may modify its *clarity*: introspectively attending to a phenomenal state may render its phenomenology more determined and defined. Since any attention shift entails a change in the center-periphery structure and thereby a change in salience, any act of primitive introspection entails a change in the target phenomenal state’s salience. Whether any change in salience is associated with a change in intensity or clarity is not obvious. One appealing view, for instance, is that the phenomenal change associated with increased salience is to be (at least partly) spelled out in terms of increased phenomenal clarity. On such a view, the more salient an element in the phenomenal field is, the more phenomenally defined, clear, and determined it is. However, there may be reasons to resist the idea that any phenomenal change associated with salience entails a change in the way things appear to us (Watzl 2017: 171-81). If so, the phenomenology of salience will be *sui generis* (i.e. not reducible to other types of phenomenology). I remain neutral about this point. What matters to my present purpose is that any act of primitive introspection entails a phenomenal modification, i.e. the phenomenal modification which is associated with salience increasing.

Therefore, at any time it is introspected, a phenomenal state undergoes a phenomenal change. The phenomenology of the target phenomenal state is not the same before and while being primitively introspected: some aspects of it (i.e. those associated with salience) *change* upon the relevant phenomenal state’s being introspected. Such a phenomenal change being due to the act of primitive introspection, it depends on the presence of a state of primitive introspection.¹³ Accordingly, the target phenomenal state, *post-change* in phenomenal features associated with salience, cannot exist independently of the state of primitive introspection. This is because certain phenomenal features of the target phenomenal state can only be present when the relevant phenomenal state is primitively introspected. Therefore, the target phenomenal state, *qua* modified by primitive introspection (i.e. with all the

¹³ For, recall from Chapter 3, engaging in the activity of primitively introspecting entails being in a state of primitive introspection.

phenomenal features it possesses while being primitively introspected), cannot be present independently of the state of primitive introspection.

A satisfactory account of the nature of primitive introspection must not only accommodate the possibility of phenomenal modification. It must also (1) account for the fact that such phenomenal modification occurs at every instance of primitive introspection and (2) explain how one can primitively introspect notwithstanding phenomenal modification.¹⁴

*

As noted, an account of the metaphysics of primitive introspection explains the relationship between the target and the state of primitive introspection. Arguably, the best available candidates for an account of the metaphysics of primitive introspection are the inner sense view and the acquaintance view, which spell out the relationship between the introspected state and the introspective state in causal terms and in constitutive terms respectively.¹⁵ In the next two sections I consider them in turn.

2. The inadequacy of the inner sense account

2.1. The inner sense view

According to the inner sense view, introspection is in important respects similar to perception. As the name suggests, on this view, introspection involves the activity of an internal ‘sense’, which is in some respects analogous to the human senses like vision or audition but differs from them in that it is directed to one’s internal states, rather than to external objects.¹⁶ A view along these lines was famously put forward by Locke. Contemporary versions of the inner sense view have been defended by Armstrong (1968) and Lycan (1996),

¹⁴ Thanks to Farid Masrouf for making point (2) clear to my mind.

¹⁵ I am *not* claiming that the inner sense view and the acquaintance view are the best possible accounts of introspection *tout court*—I only claim they are the best available options as far as *primitive* introspection is concerned. When it comes to knowledge of the contents of propositional attitudes like belief and desire, for example, theories such as the rationalist transparency view (Moran 2001; Byrne 2005) or the inferentialist view (Cassam 2015) are excellent candidates (probably better candidates than the inner sense view and the acquaintance view).

¹⁶ Nota bene: inner sense theorists do not maintain that introspection literally involves the deployment of an ‘inner eye’, or of any specific introspective sense *organ*, as instead vision, audition, and the other senses do.

who characterize the inner sense as a sort of self-scanning process, or a self-detection mechanism. On this view, introspection consists in directing one's internal detection mechanism or self-scanner toward a conscious state. This yields a higher-order introspective state that represents the target conscious state.

The relationship between the higher-order introspective state and the target conscious state is analogous to the relationship between the perceived object and the perceptual state that represents it. In perception, the perceptual system *detects* a certain object and its properties. The object and its properties are the perceptual process' *input*. The *output* of the perceptual process is a perceptual *representation* of the relevant object and its properties.¹⁷ The relationship between the object and the representational state is *causal*: the former causes the latter. Analogously, on the inner sense view, in introspection, the internal scanner detects a certain conscious state. The latter is the introspective process' input. The output of the introspective process is a higher-order representation of the relevant conscious state. As before, the relationship between the target (first-order) conscious state and the higher-order introspective representation is causal: the former causes the latter. The inner sense view may therefore be characterized thus: one has an introspective state only if one has a higher-order mental state which represents the target first-order conscious state.¹⁸

Arguably, an *inner sense account of primitive introspection* explains the relationship between the act, the target, and the state of primitive introspection as follows. The act of primitive introspection involves directing one's internal scanner toward the target phenomenal state. The internal scanner can collect information about the phenomenology of the target state. This act produces a higher-order state which represents the target phenomenal state in a non-classificatory way. Such a higher-order non-classificatory representation is the state of primitive introspection.

In perception, the perceptual state and the object seen are distinct: they can exist independently of each another—the object can exist without being seen, and the perceptual experience can exist even if no object is there to be perceived (this is what happens in cases of perceptual hallucination). Analogously, on the inner sense view, the state of primitive introspection and the target phenomenal state are distinct: each can exist without the other. Accordingly, the target phenomenal state can exist without being primitively

¹⁷ I am assuming here a representationalist account of perceptual experience.

¹⁸ See Gertler (2011, chapter 5) for a full explanation of the inner sense view.

introspected and the state of primitive introspection can exist independently of there being a phenomenal state to be introspected.

As noted, in perception, the perceptual state is related to the object seen by a causal process. Similarly (on the inner sense view), in primitive introspection, the higher-order introspective state stands in a causal relation to the target phenomenal state: the target phenomenal state causes the state of primitive introspection. Therefore, the relationship between the state of primitive introspection and the target phenomenal state is contingent, as is the relationship between the perceptual representation and the perceived object. Given the contingency of the relationship between them, the state of primitive introspection may misrepresent the target phenomenal state, as a perceptual state may misrepresent the target object.

The inner sense account might be claimed to have the resources to satisfy the phenomenal modification desideratum. Plausibly, the inner sense theorist could argue that the phenomenal change associated with increased salience is due to the way the higher-order introspective state represents the target phenomenal state. Accordingly, the state of primitive introspection represents the target phenomenal state *as being more salient* (more intense, more defined, clearer, and so on).¹⁹

However, the inner sense account cannot satisfy the no-introspective-hallucination desideratum. As noted, on the inner sense view, the state of primitive introspection and the target phenomenal state can exist independently of each other. This entails that there can be a state of primitive introspection even if no phenomenal state is there to be introspected. Just as a visual representation can be formed even in the absence of any object before one, a state of primitive introspection can be formed even in the absence of any phenomenal state.

2.2. Rejoinder attempt n. 1: mere nomological necessity?

Inner sense theorists may object that, although their account entails that introspective hallucination is *metaphysically* possible, it is compatible with its being *nomologically impossible*. Although in principle there can be creatures who are in a state of primitive introspection even in the absence of *any* phenomenal state, this cannot happen in the actual world because of how our introspective mechanisms are built. The inner sense theorist might thus

¹⁹ The view could be developed in such a way that this does not entail *recognizing* that the phenomenal state is more salient—if, for instance, the higher-order introspective state involves a *nonconceptual* (mis)representation of the target state.

develop an account of human cognitive system to the effect that, in the actual world, one can only be in a state of primitive introspection if there is a separate phenomenal state to which it is directed.

However, this is unsatisfying. For one thing, it is not clear that the nomological impossibility claim can convincingly be defended from an inner sense perspective. Every known mechanism in the natural world sometimes goes off-track. Why should the introspective mechanism be an exception? In the inner-sense framework, the nomological-necessity move smells a bit *ad hoc*: at the very least, the inner sense theorist owes us an articulated explanation of the nomological impossibility of hallucination in primitive introspection. Arguably, the best and simplest explanation of the nomological impossibility of introspective hallucination is its *metaphysical* impossibility.²⁰

Moreover, nomological impossibility is not yet enough to satisfy the no-introspective-hallucination desideratum. As motivated in §1.1, the impossibility of hallucination in primitive introspection is a *metaphysical* impossibility. The state of primitive introspection is itself a conscious state with phenomenology: necessarily, when you primitively introspect the pain sensation in your stomach, you are in a conscious state with phenomenology (arguably, a phenomenology which is characteristic of pain sensations). Accordingly, the primitive introspective state is itself a phenomenal state. Therefore, being in a state of primitive introspection entails being in a phenomenal state. Arguably, the phenomenal properties of the primitive introspective state are somehow inherited from its target: plausibly, the painful character of your pain-sensation directed primitive introspective state somehow depends on the phenomenal character of the introspected pain sensation. As we will see, the acquaintance account accommodates this intuition. The inner sense view, however, allows for cases in which there is primitive introspection but no phenomenal state. It seems that, in such cases, the state of primitive introspection could not have phenomenology, because there would not be any phenomenal state whose phenomenal properties the primitive introspective state could inherit. Therefore, in those cases there would be a state of primitive introspection but no phenomenal state, which is inconsistent with the no-introspective-hallucination desideratum.

²⁰ I owe this argument against the inner sense theorist's nomological-necessity move to Uriah Kriegel.

2.3. Rejoinder attempt n. 2: can inner sense deliver metaphysical necessity?

Inner sense theorists may reject the reasoning at the end of §2.2 and argue that the inner sense account can be compatible with the metaphysical impossibility of introspective hallucination after all. The idea is the following.²¹ We distinguish the good cases from the bad cases of primitive introspection. In the good cases, one has two distinct mental states, the higher-order introspective state and the target phenomenal state, causally related the way explained above. In the bad cases, there is only one mental state—the introspective state. However, in the bad cases, the introspective state itself has phenomenology—it is a phenomenal state. Therefore, both in the good and in the bad cases, being in a state of primitive introspection entails there being a phenomenal state.

I have two remarks in reply to this objection. First, the inner sense theorist distinguishes between good and bad cases. However, it is unclear what makes bad cases *bad* on the view sketched above. Arguably, the inner sense theorist thinks that those cases are bad because the state of primitive introspection does not track a phenomenal state that is actually there. But this is incorrect, since, by hypothesis, as soon as you primitively introspect you thereby *have* the relevant phenomenal state—the introspective state does track an existing phenomenal state, which happens to coincide with the introspective state itself. Therefore, if ‘good’ and ‘bad’ are to be interpreted in referential terms (where the bad case is a case of reference failure), then the ‘bad’ case here is not really bad (it is not a case of reference failure). Moreover, it is also unclear how ‘bad’ cases are supposed to be bad if we interpret the good/bad distinction in epistemic terms. For there seems to be no reason to doubt that, in the so-called ‘bad’ case, the epistemic status of the introspective state will be, at the very least, equally good as that which the introspective state has in the case in which it tracks a separate (i.e. metaphysically independent) phenomenal state. In fact, since in this case the introspected state is not independent of the introspective state (they coincide), it is plausible that the epistemic status of the introspective state is even *better*, here, than in the case in which there are two distinct mental states, since the absence of an intermediary causal process may reduce the possibilities of error. It would then seem that, in fact, the ‘bad’ case is at least just as good as the ‘good’ case (perhaps even *better* than the ‘good’ case, or even *maximally* good).

Secondly, every state of primitive introspection aims at a phenomenal state: to be introspective, a mental state must be *about* a phenomenal state.

²¹ This objection was raised to me by François Kammerer.

Therefore, at any instance of primitive introspection, we must conceptually distinguish between the introspective state and the target state it aims at. In the ‘good’ case, this is straightforwardly accounted for by the fact that the introspective state and the target phenomenal state are separate states. As for the ‘bad’ case, the inner sense theorist will plausibly have to maintain that, although the introspective state and the target phenomenal state are not distinct (there is only one mental state), we can still conceptually tell them apart mereologically: the introspective state is (partly) constituted by the target phenomenal state. This move seems to drift considerably apart from the spirit of the inner sense view and get closer to that of the acquaintance view, which spells out the relationship between the introspective state and the target phenomenal state in terms of constitution.

The resulting view is a *disjunctive* account, which explains ‘good’ cases in terms of inner sense and ‘bad’ cases in terms of acquaintance. Now, disjunctive accounts, of course, are not necessarily theoretically bad. However, arguably, a disjunctive account should be discarded if (i) a unified account with equal explanatory power is available or (ii) the disjuncts do not track a theoretically interesting distinction. As noted, so-called ‘bad’ cases are neither referentially bad nor epistemically bad and indeed it is unclear what the good/bad distinction is supposed to track.²² Moreover, I am going to argue that an alternative unified account *is* available. Therefore, there seem to be no reason to endorse the disjunctive version of the inner sense view.

*

To sum up, although it may accommodate the phenomenal modification desideratum, the inner sense account could only accommodate the no introspective hallucination desideratum by committing itself to a disjunctive theory of dubious theoretical value. This seems to be a good reason for exploring an alternative account.

²² The inner sense theorist might not want to say that it tracks the distinction between cases in which there are two distinct states (the introspective state and the target phenomenal state) and cases in which there is just one state (the introspective state), for this would clearly be just a question-begging *ad hoc* move.

3. The inadequacy of extant acquaintance accounts

On the acquaintance view, the metaphysical relationship between the introspective state and the introspected state is much more intimate than on the inner sense view. As noted, on the inner sense view the introspective state and the introspected state are distinct (they can exist independently of one another) and the relationship between them is metaphysically mediated or indirect—it involves a causal process. On the acquaintance view, by contrast, the relationship between the introspected state and the introspective state is metaphysically *immediate* (direct): no state or process mediates between them (Gertler 2011: 90-91).

The metaphysically immediate relationship of acquaintance is typically explained in terms of *constitution* (BonJour 2000; Gertler 2001, 2012; Chalmers 2003; Pitt 2004; Horgan and Kriegel 2007; Balog 2012): the introspective state is (partly) constituted by the introspected state.²³ This may be understood in mereological terms: the introspected state is a (proper) part of the introspective state. Therefore, on the acquaintance view, when one introspects, there are no two distinct states, the introspective and the introspected, but rather one single (introspective) state (which is partly constituted by the introspected state). Different from the inner sense view, then, on the acquaintance view the introspective and the introspected state are not independent of one another. Although the target phenomenal state can exist independently of the introspective state, at least the introspective state's existence depends on that of the target phenomenal state. Brie Gertler (2001), who provides a particularly rigorous and detailed analysis of the metaphysics of acquaintance, spells out the constitutive relationship partly in terms of what she calls *embedding*: the introspective state *embeds* the introspected state, where a mental state *a* is *embedded* in a mental state *b* iff (i) if *b* is present, *a* must be present too, and (ii) *a* can be present even if *b* is not present.²⁴

The idea then is that, when one primitively introspects, one's state of primitive introspection embeds the target phenomenal state: (i) if the introspective state is present, the target phenomenal state must be present too and (ii) the target phenomenal state can be present even if the

²³ The acquaintance relationship is usually spelled out in terms of constitution, but whether the constitution is full or partial is often left unspecified.

²⁴ Embedding is not the only condition that must be satisfied in order to have introspection of phenomenal states on Gertler's account. It must also be the case that (a) the introspective state refers to the introspected state and (b) it is in virtue of the embedding relationship that the introspective state refers to the introspected state.

introspective state is not present.²⁵ When you primitively introspect a pain sensation in your right knee, the pain sensation (partly) constitutes your state of primitive introspection. Arguably, this entails that the *phenomenology* of your introspective state is constituted by the phenomenology of the pain sensation: your introspective state has a pain phenomenology and this is due to the fact that the pain sensation (partly) constitutes your introspective state. If no pain sensation were there, you could not have been in this state of primitive introspection—as per (i). However, the existence of your pain sensation is independent of whether the sensation is introspected or not—as per (ii). The latter captures a straightforward intuition about introspection. Intuitively, the pain sensation you are now introspecting could have existed (and indeed *did* exist) before your introspecting it and, plausibly, it will continue to exist after you cease focusing your introspective attention on it. More generally, phenomenal states do not need to be introspected in order to exist.

The acquaintance account satisfies the no-introspective-hallucination desideratum. As per (i), the state of primitive introspection cannot be present without the target phenomenal state being present too. The constitutive relationship entails that the existence of the introspective state depends on that of the introspected phenomenal state. Therefore, there cannot be state of primitive introspection if no phenomenal state is present.

Incidentally, it is perhaps worth pointing out that not only the acquaintance account satisfies the no-introspective-hallucination desideratum by implying that there cannot be a state of primitive introspection without *a* phenomenal state being there. It entails that there cannot be a state of primitive introspection without *its target* being there. Therefore, the metaphysical necessity of there being a phenomenal state at any instance of primitive introspection depends on the fact that any state of primitive introspection entails the presence of its target. This makes justice to the intuition that the primitive introspective state's phenomenology is somehow inherited from its target and that *this* (rather than the primitive introspective state's merely *having* phenomenology) is what grounds the impossibility of introspective hallucination.

However, as it stands, the acquaintance account does not satisfy the phenomenal modification desideratum because it does not explain the phenomenological change the target introspective state undergoes upon being

²⁵ I am applying Gertler's account to primitive introspection here, although she does not explicitly propose this as an account of *primitive* introspection—she proposes this as an account of introspection of phenomenal states more generally.

primitively introspected. As noted in §1.2, due to primitive introspection's entailing an act of attention, the phenomenology of the target phenomenal state is not the same before and while being primitively introspected: the phenomenal aspects of it which are associated with salience change upon the relevant phenomenal state's being primitively introspected. Importantly, such a phenomenological change *depends* on the state of primitive introspection: it is in virtue of its being primitively introspected that the target phenomenal state undergoes the relevant phenomenological change. Therefore, the target phenomenal state, *post-change* in phenomenal features associated with salience, depends on the state of primitive introspection.²⁶ The introspected phenomenal state, *qua* modified by primitive introspection, could not have existed independently of the state of primitive introspection. To be sure, the claim here is *not* that a phenomenal state *of the same kind* could not have existed independently of being primitively introspected. There could be a creature who undergoes an instance of the same kind of phenomenal state (with the same salience, intensity, clarity, and so on) without the relevant phenomenal state being primitively introspected. Nor is the claim that, for *any* phenomenal state ϕ with a certain (*sufficiently high*, or *maximal*) degree of salience, ϕ can only exist if primitively introspected. Rather, the claim is that, when a token phenomenal state is primitively introspected, *that token* phenomenal state which is actually primitively introspected, and which has thereby undergone a phenomenological change associated with salience, could not have existed independently of the relevant state of primitive introspection.

It therefore seems that the state and the target of primitive introspection are *interdependent*. Not only the introspective state depends on the introspected state: there is an important sense in which the introspected state, once introspected, depends on the introspective state. Such interdependence is neglected by extant acquaintance accounts. As noted, extant acquaintance accounts imply that the introspected state can exist independently of the introspective state. Although this captures an important intuition about introspection—namely that a phenomenal state may exist before and continue to exist after (and therefore independent of) being introspected, it fails to account for the idea that the target phenomenal state, *post-change* in phenomenal features associated with salience, depends on the state of primitive introspection.

²⁶ The dependence here is constitutive, rather than causal. Perhaps causal dependency may feature in an inner-sense picture, where the target and the state of primitive introspection are distinct and causally related. In the acquaintance framework, however, the dependence is constitutive.

To be sure, the acquaintance view is not *incompatible* with phenomenal modification. Recall, however, that to satisfy the phenomenal modification desideratum, a theory of the metaphysics of primitive introspection must not only be *compatible* with phenomenal modification. It must also (1) account for the fact that such phenomenal modification occurs at every instance of primitive introspection and (2) explain how one can primitively introspect notwithstanding phenomenal modification. What extant acquaintance accounts lack is an *explanation* of why phenomenal modification occurs at every instance of primitive introspection and of how primitive introspection is possible although the target phenomenal state undergoes a phenomenal change at any time it is introspected.

I suggest that the acquaintance account should be implemented so that it can explain phenomenal modification. A more refined version of the acquaintance account needs to be developed which accounts for the fact that the introspected phenomenal state, *qua* modified by the act of primitive introspection, cannot exist independently of the state of primitive introspection. On the other hand, the account must capture the intuition (as the other versions of the acquaintance account do) that phenomenal states can exist independently of being primitively introspected. Although there is a *prima facie* tension between these two requirements (the introspected state is claimed to be both dependent on and independent of the introspective state), an account can be developed on which the two requirements are both satisfied—there is no *ultima facie* tension between them. The next section aims to develop such an account.

4. A proposed solution: the integration account of primitive introspection

In this section I develop an account of the metaphysics of primitive introspection which satisfies both the no-introspective-hallucination desideratum and the phenomenal modification desideratum. The account I propose—what I call the ‘integration account’ of primitive introspection—is a version of the acquaintance view. While setting up the grounds of the theory, I will make some substantive assumptions about the metaphysics of the mind. I argue that if the assumptions I make are accepted, the integration account is a good account of the metaphysics of primitive introspection. However, I do *not* argue that the integration account is the *sole* satisfactory account of the metaphysics of primitive introspection: equally satisfactory options may be available to those who make different assumptions.

The integration account is a version of the acquaintance view. As such, it entails that there is a relationship of *constitutive dependence* between the target and the state of primitive introspection. However, whereas extant acquaintance accounts spell out such a relationship in terms of *unilateral* dependence (the primitive introspective state depends on its target, but the target is independent of the introspective state), the integration account spells it out in terms of *interdependence*: not only the primitive introspective state depends on the target; once introspected, (some aspects of) the target depend on the introspective state.

At a first approximation, the idea is the following. Once introspected, the target phenomenal state becomes a *proper part* of the primitive introspective state. In virtue of this, the state of primitive introspection inherits some of the target phenomenal state's phenomenal properties. This explains, first, the dependence of the primitive introspective state on the presence of a phenomenal state (which is entailed by the no-introspective-hallucination claim) and, second, our intuition that the phenomenology of the primitive introspective state somehow derives from the phenomenology of its target. So far, the explanatory power of the integration account is roughly equal to that of other acquaintance accounts. What makes it different from (and, on my view, better than) other acquaintance accounts is that, on the integration account, some of the phenomenal properties that the target phenomenal state has once it is introspected depend on the presence of the primitive introspective state. This explains why, once introspected, the target phenomenal state depends on the state of primitive introspection (which is entailed by the phenomenal modification claim).

To make the account more precise and clarify its explanatory power some set-up is in order. In §4.1 I lay out the assumptions on which the integration account is based. In §4.2 I articulate a more precise and more explanatory formulation of the view.

4.1. Preliminary assumptions

4.1.1. *The ontological status of phenomenal states.* In thinking about the metaphysics of primitive introspection, it may be helpful to specify what kind of entity we have in mind when we talk about phenomenal states. As noted in Chapter 3, there are at least five options. First, a phenomenal state may be seen as a *particular*, a bearer of phenomenal properties (I called this the *Particular View*). Accordingly, a mild pain sensation is characterized as a particular which has the phenomenal property associated with pain (i.e. phenomenal painfulness) and, perhaps, the phenomenal property associated

with mildness (i.e. phenomenal mildness). Second, a phenomenal state may be characterized in terms of *phenomenal property instance*. On this view, your mild-pain sensation is a phenomenal-property instance, or phenomenal trope (I called this the *Phenomenal Property Instance View*). Third, a phenomenal state may be described as a Kimean event (Kim 1993), i.e. the instantiation of a phenomenal property by a subject at a time (I called call this the *Kimean View*). On this view, the pain sensation you are currently undergoing is characterized as the instantiation of phenomenal painfulness by you (the subject) now. Fourth, the phenomenal state may be seen as a Davidsonian event, or as a Davidsonian state.²⁷ On this view, your pain sensation is an unrepeatable individual, irreducible to more fundamental constituents. Fifth, a phenomenal state may be characterized mereologically, as a phenomenal part of one's overall experience (call this the *Mereological View*). Consider your current overall experience, which encompasses, among other things, your seeing the screen before you, your hearing children playing in the courtyard, and your feeling pain in your right knee. The pain sensation is here seen as a phenomenal *part* of your current overall experience (as are the screen seeing and the children hearing).²⁸

All I say in what follows can fit all five views about the metaphysical status of phenomenal states. I will often talk as if phenomenal states are particulars (bearers of phenomenal properties), mostly for expository reasons. This seems to me the easiest way to talk about the metaphysical structure of primitive introspection. When needed, I will specify how the account is to be adjusted to fit different views of the metaphysical status of phenomenal states.

4.1.2. Kinds of phenomenal properties. Phenomenal states have properties, some of which are phenomenal (e.g. reddishness, painfulness) other non-phenomenal (e.g. occurring on Tuesday). It might be useful to divide phenomenal properties into three kinds: qualitative, quantitative, and relational.

Qualitative phenomenal properties are phenomenal properties such as reddishness, bluishness, painfulness, burning painfulness, stabbing painfulness, and so on. They define the qualitative aspect which characterizes each kind of phenomenal state: they are that which constitutes the difference between, say, an experience of redness and an experience of blueness (or

²⁷ As noted in Chapter 3, the two views may or may not collapse into each other. For ease of exposition, I will run them together here.

²⁸ There are surely other options to be explored, but the ones mentioned here strike me as the theoretically most interesting.

between an experience of saxophone sound and an experience of stubbing pain). Arguably, two qualitatively different phenomenal states (phenomenal states with different qualitative phenomenal properties) are phenomenal states of different kinds.

Quantitative phenomenal properties are properties in virtue of which qualitative phenomenal properties are modulated. Each dimension of phenomenal variation is associated with a quantitative phenomenal property. Quantitative phenomenal properties are phenomenal properties that can be quantified: for any quantitative phenomenal property Q , a phenomenal state ϕ , that can have Q , can be more or less Q . Quantitative phenomenal properties are *phenomenal* because they make a difference to the phenomenology of one's experience. *Intensity* is the paradigmatic example of quantitative phenomenal property. A pain sensation, for instance, may be more or less painful. When you stub your toe on the couch, you have a sudden excruciating pain sensation in your toe, which gradually diminishes. In this case, the pain sensation, initially *extremely intense*, becomes *less intense* with the passage of time. Similarly, other kinds of phenomenal state may vary in intensity. A visual experience may be more or less bright, an auditory experience more or less loud, an itchy experience more or less itchy, an anger experience more or less intense, and so on and so forth. The saturation and brightness dimensions of variation in color experiences are also quantitative phenomenal properties—they could perhaps be classified as species of intensity.²⁹ *Clarity* (in the sense of definiteness as spelled out in §1.2) may also be numbered among the quantitative phenomenal properties. As long as a certain phenomenal state can be phenomenally *more or less determinate*, its clarity or definiteness is a phenomenal property that can be quantified (it is a quantitative phenomenal property). Given that a phenomenal state can be more or less salient, *salience* is also a candidate for being a quantitative phenomenal property. Whether salience can legitimately be numbered among quantitative phenomenal properties partly depends on whether a phenomenal state's saliency can be measured only relatively to other phenomenal states (ϕ_1 is more salient than ϕ_2) or there is also an absolute measure of salience (ϕ 's degree of salience is n) instead. As for intensity, for example, not only a phenomenal state can be more or less intense than another phenomenal state: one can also establish a scale, say from 1 to 10, such that a phenomenal

²⁹ Indeed, whether it is in virtue of one and the same quantitative phenomenal property, i.e. intensity, that all qualitative properties are modulated along the dimension of variation we intuitively call intensity is controversial. It may be argued that saturation, brightness, loudness, degrees of pain, itch, anger, and so on, are different quantitative phenomenal properties rather than distinct exemplifications of one and the same property (intensity).

state can be said to have a certain degree of intensity within that range. (If I ask you how intense the pain sensation in your toe is, in a scale from 1 to 10, right after you have stubbed it on the couch, you may reply, say, 6; the pain intensity may drop to 2 after a while.) If a similar scale can be established for salience, then salience may be considered as a quantitative phenomenal property as well.

Qualitative and quantitative phenomenal properties are both kinds of non-relational phenomenal properties. *Relational* phenomenal properties are phenomenal properties phenomenal states have in virtue of bearing certain relations to other phenomenal states. Among relational phenomenal properties there are: phenomenal unity (φ_1 is unified with $\varphi_2, \varphi_3, \varphi_4, \dots,$ and φ_n), mereology (φ_1 is part of the same whole as $\varphi_2, \varphi_3, \varphi_4, \dots,$ and φ_n), temporality (φ_1 is [experienced] before/after φ_2), and salience (φ_1 is more/less salient than φ_2). Relational phenomenal properties are *phenomenal* because they make a difference to the phenomenology. Phenomenal states φ_1 and φ_2 which are part of the same overall experience are not only unified, but also *experienced as* unified: not only there is something it is like to have φ_1 and something it is like to have φ_2 —there is also something it is like to have φ_1 and φ_2 *together* (Bayne 2010: 11). Moreover, being experienced together with φ_2 may have an effect on φ_1 's phenomenology: the latter might have been (slightly) different had φ_1 been experienced together with φ_3 rather than with φ_2 . For similar reasons, φ_1 's being part of the same whole as $\varphi_2, \varphi_3,$ and φ_4 has an effect on its phenomenology: had it been part of the same whole as $\varphi_5, \varphi_6,$ and φ_7 —rather than $\varphi_2, \varphi_3,$ and φ_4 — φ_1 's phenomenology would have been different. Your experience of the sound of the saxophone, for example, has probably a slightly different phenomenology when the sound of the saxophone is experienced together with the sound of a xylophone and that of an electric guitar, from when it is experienced together with the sound of a piano and that of a clarinet. φ_1 's being experienced after φ_2 rather than after φ_3 also makes a difference to its phenomenology. Eating a piece of cheesecake after feeling hungry has a different phenomenology from eating a piece of cheesecake after having already eaten a first piece of cheesecake. As for salience, we have already seen in §1.2 that difference in salience makes a difference to the phenomenology: φ_1 's phenomenology is different when φ_1 's is more salient than $\varphi_2, \varphi_3,$ and φ_4 , from when it is as salient as φ_2, φ_3 and less salient than φ_4 . The phenomenology of your experience of the saxophone is different when the sound of the saxophone is the most salient aspect of your phenomenal field, from when the sound of the piano is the most salient aspect.

One could spell out the relationship between qualitative, quantitative, and relational phenomenal properties in terms of different orders.³⁰ Qualitative phenomenal properties are first-order phenomenal properties. Quantitative and relational phenomenal properties are second-order phenomenal properties—they are properties of phenomenal properties.

If phenomenal states are particulars, as per the *Particular View*, then they are bearers of quantitative, qualitative, and relational phenomenal properties. Things are similar if they are phenomenal parts of the overall experience (as per the *Mereological View*): each phenomenal part has qualitative, quantitative, and relational phenomenal properties.

If phenomenal states are phenomenal-property instances, as per the *Phenomenal Property Instance View*, then, plausibly, they are instances of *qualitative* phenomenal properties. Quantitative and relational phenomenal properties are properties of qualitative-phenomenal-property instances (which squares well with the idea that they are second-order phenomenal properties).

If phenomenal states are Kimean events, then at least two options suggest themselves. On the first option, call it the *True-Kimean View*, a phenomenal event consists in the subject's instantiating a certain *qualitative* phenomenal property. Quantitative and relational phenomenal properties are properties of phenomenal events: they are properties of instantiations of qualitative phenomenal properties.³¹ This option is consistent with Kim's own conception that events have *one* constitutive property (which, together with the object which instantiates it and the time at which the property is instantiated, individuates the relevant event) but can themselves exemplify several properties (Kim 1993: Ch. 3).³² It also goes well with the idea that qualitative phenomenal properties are first-order, whereas quantitative and relational phenomenal properties are second-order. On the second option, call it the *Pseudo-Kimean View*, phenomenal events are *co*-instantiations of several phenomenal properties: qualitative, quantitative and relational, by a subject at a time. If so, quantitative and relational phenomenal properties are

³⁰ I am indebted to Takuya Niikawa for suggesting this to me.

³¹ Of course, there are further options in the logical space: that phenomenal events are instantiations of quantitative phenomenal properties (qualitative and relational phenomenal properties are properties of phenomenal events), and that phenomenal events are instantiations of relational phenomenal properties (qualitative and quantitative phenomenal properties are properties of phenomenal events). These options, though, do not seem worth considering.

³² The view is consistent with Kim's conception if, for any phenomenal state ϕ , there is just one qualitative phenomenal property which individuates ϕ .

also constitutive: they contribute to the individuation of phenomenal events. This option seems fully legitimate, although it drifts considerably apart from the Kimean spirit.

If, instead, phenomenal states are Davidsonian events, they do not involve instantiation of phenomenal properties. However, arguably, the gist of our distinction between qualitative, quantitative, and relational properties may be adapted to fit a Davidsonian framework nonetheless. For instance, a version of the Davidsonian account may be developed such that there are three kinds of Davidsonian event, qualitative, quantitative, and relational. On this view, what is primitively introspected is a triad constituted by three co-occurring Davidsonian phenomenal events—a qualitative, a quantitative and a relational event. Another version could be that there is only one kind of Davidsonian event (namely the one which corresponds to the instantiation of *qualitative* phenomenal properties in the Kimean framework), but Davidsonian events can have quantitative and relational phenomenal properties.³³ In what follows I will set the *Davidsonian View* aside for expository reasons. Besides, the *Davidsonian View* of phenomenal states strikes me as the most idiosyncratic. At any rate, although that would require some extra theoretical work, my theory of primitive introspection can be adapted to the *Davidsonian View* as well.

4.1.3. Essential vs. accidental properties. Some of a phenomenal state's properties are essential to it, other accidental. Notwithstanding its well-known problems (Fine 1994), the modal characterization of essential and accidental properties will suffice for the present purpose. Roughly, essential properties are properties something must have—did it not have its essential properties, x would not be x , i.e., it would not exist (at x 's place there would be something else, $\neq x$, or nothing). Accidental properties are properties something has, but could have lacked—although x actually has F , it could have existed even if it lacked F . However, you may plug your preferred characterization of essential and accidental properties into my account, provided that it entails the modal characterization as roughly presented above. Accordingly, the essential properties of a phenomenal state ϕ are properties it must have: if it did not have them, ϕ would not exist. The accidental properties of a phenomenal state ϕ are properties that ϕ has but could have lacked. Importantly, a phenomenal state is individuated by its essential properties: ϕ and ψ are the same phenomenal state iff they have exactly the same essential properties.

³³ Again, further options may well be explored.

Therefore, even if a phenomenal state's accidental properties change, it remains the same state.

Which of a phenomenal state's properties are essential to it, and which are accidental? Many options are available. For one thing, as noted, phenomenal states have both phenomenal and non-phenomenal properties. Accordingly, here is a preliminary menu of options:

- (a) *All* of a phenomenal state's properties (that is, all the phenomenal and all the non-phenomenal ones) are essential.
- (b) All the *phenomenal* properties, but only *some* non-phenomenal properties, are essential.
- (c) *Some* phenomenal properties and *some* non-phenomenal properties are essential.
- (d) *All and only* the *phenomenal* properties are essential (all non-phenomenal properties are accidental).
- (e) *Some but not all* phenomenal properties are essential and all non-phenomenal properties are accidental.

The list is not exhaustive—the left-out options, however, strike me as so implausible not to be worth considering.³⁴ Option (a) entails a version of essentialism about phenomenal states which is perhaps too strong. It entails, for instance, that occurring at the same time as the Pope's *Urbi et Orbi* blessing is essential to the excruciating pain in your toe. Options (b) and (c)—that is, the options that entail that *some but not all* non-phenomenal properties are essential—may be adopted, for example, by those whose preferred ontology of phenomenal states implies that token phenomenal states do not extend through time but are instantaneous.³⁵ On this view, occurring at time *t*, which is non-phenomenal, is an essential property of ϕ — ϕ is individuated partly by (some of) its phenomenal properties, and partly by the property of occurring at time *t*, which is non-phenomenal. Options (d) and (e) may be adopted by those who have the intuition that only *phenomenal* properties should play a role in the individuation of *phenomenal* states (all non-phenomenal properties are not essential).

However, as it will become clear in the next section, our present interest is somewhat independent of whether some (and, if, so, which and how many)

³⁴ Consider, for instance, the option that all and only the non-phenomenal properties of a phenomenal state are essential: once you accept that phenomenal states have phenomenal properties, a view that entails that phenomenal properties play no role in the individuation of phenomenal states sounds pointless.

³⁵ The view under consideration is a sort of four-dimensionalism about phenomenal states. Thanks to Yaojun Lu for suggesting this option to me.

non-phenomenal properties are essential to a phenomenal state. For what we are concerned with, in developing an account of the metaphysics of primitive introspection, are the properties of a phenomenal state which can be primitively introspected, namely its *phenomenal* properties. Therefore, what is crucial here is to establish which *phenomenal* properties are essential to a phenomenal state: different assumptions in this respect will lead to different accounts of the metaphysics of primitive introspection. Accordingly, here is a more relevant menu of options:

- (f) All of a phenomenal state's phenomenal properties are essential.
- (g) Among the phenomenal properties, all and only the *non-relational* (i.e. qualitative or quantitative) phenomenal properties are essential.
- (h) Among the phenomenal properties, all and only the *qualitative* phenomenal properties (i.e. all its non-relational and non-quantitative phenomenal properties) are essential.
- (j) *Some but not all* qualitative phenomenal properties are essential and all non-qualitative phenomenal properties are accidental.

Again, the list is not exhaustive—I selected the options which strike me as the most relevant.³⁶ Here I assume that (h) is true: all and only the qualitative phenomenal properties of a phenomenal state are essential to it—the quantitative and the relational phenomenal properties are accidental. On this view, the painful aspect of the pain sensation in your toe—phenomenal painfulness—is essential to it. Its intensity, by contrast, is *not* essential (it is accidental), as well as the property of, say, being unified with children hearing and chocolate tasting and the property of occurring after a visual experience of a written sheet of paper.³⁷

³⁶ For instance, I have not considered the option that all non-qualitative phenomenal properties are essential because it seems obvious enough to me that quantitative and relational phenomenal properties are not sufficient to individuate phenomenal states.

³⁷ A few considerations aimed at seeing how this interacts with some potential views of the ontological status of primitive introspection's target. Option (h) is compatible with all views of the metaphysical status of phenomenal states except the *Pseudo-Kimean View*. On the *Pseudo-Kimean View*, recall, phenomenal states are phenomenal events characterized as co-instantiations of qualitative, quantitative, and relational phenomenal properties. Accordingly, phenomenal states are individuated by all their qualitative, quantitative, and relational phenomenal properties. Therefore, on the *Pseudo-Kimean View*, not only qualitative phenomenal properties, but also quantitative and relational phenomenal properties are essential, which is inconsistent with (h). On the other hand, (h) fits particularly nicely with the *True-Kimean View*, according to which phenomenal states are phenomenal events characterized as instantiations of a qualitative phenomenal property (by a subject at a time). The *True-Kimean View* entails (h), for it entails that the property which individuates a phenomenal event is a qualitative phenomenal property.

On option (g), only relational phenomenal properties are accidental. This entails that not only qualitative phenomenal properties, but also quantitative phenomenal properties like intensity are essential. Accordingly, the pain sensation you have right after stubbing your toe on the couch (which is very intense) and the one you have one minute later, when the pain in your toe has cooled down a bit (which is slightly less intense), are different phenomenal states.³⁸ On option (f), all the phenomenal properties of a phenomenal state—qualitative, quantitative, and relational—are essential to it.³⁹ Finally, option (j) is that only some (but not all) qualitative phenomenal properties are essential. I am not sure about the theoretical usefulness of this last option. For one thing, it seems to me that any account which takes only a subset of a phenomenal state’s qualitative properties as essential would do so on arbitrary grounds—there does not seem to be any matter of fact concerning which qualitative phenomenal properties of a phenomenal state should be considered as essential.

Those who resist option (h) to favor options (f) or (g) may do so on the basis of intuition-driven or theoretically-driven considerations. Some may have the intuition that the pain sensation you have when you have just stubbed your toe is radically different from the sensation you have later, when the pain has cooled down sensibly and is about to disappear. On the basis of intuitions of this kind, it may be argued that intensity is essential to a phenomenal state: phenomenal states with different intensities are different phenomenal states. Similarly, it may be argued that being more or less determinate makes an essential difference to phenomenal states and that, therefore, clarity is an essential phenomenal property. If so, at least some quantitative phenomenal properties are essential. On the other hand, some philosophers may maintain that at least some relational properties are essential to a phenomenal state, for theoretical reasons: strong holistic theories of the unity of consciousness, for instance, entail that being unified

³⁸ And this *not* just in virtue of the fact that the relevant sensations occur at different times, as the four-dimensionalist about phenomenal states would argue. On the view under consideration here, the relevant sensations are distinct phenomenal states because they have different intensities, and intensity contributes to the individuation of phenomenal states.

³⁹ The latter matches particularly nicely with the *Pseudo-Kimean View* of the metaphysical status of phenomenal states. On the *Pseudo-Kimean View*, phenomenal states are individuated by all their qualitative, quantitative, and relational phenomenal properties, which is exactly what (f) entails. It is also perfectly consistent with the other views. Its consistency with the *True-Kimean View*, though, is not straightforward. However, one could have a (slightly complex) account on which not only constitutive properties are essential to phenomenal events: some of the phenomenal properties exemplified by the phenomenal event itself (rather than by the *object* which is constitutive of the event) are essential too.

with the phenomenal states it is actually unified with is essential to a given phenomenal state.⁴⁰

I do not think there are definitive reasons to reject (f) or (g)—indeed I believe that these options are worth exploring. However, I think that there are some reasons to prefer (h). For one thing, although the pain sensation you have right after stubbing your toe and the one you have when the pain has diminished are obviously different, the difference between the two sensations seems to be a difference in degree, rather than a difference in kind. Whereas the difference between a burning pain and a stubbing pain, or the difference between a pain sensation and a tickling sensation strike me as obviously essential, I am inclined to see the difference between qualitatively identical sensations of different intensities as accidental. Moreover, there are theoretical reasons to resist the idea that quantitative phenomenal properties are essential. If quantitative phenomenal properties such as intensity and clarity were essential, there would be a multiplication of phenomenal states, which friends of parsimony might want to avoid.⁴¹ Degrees of intensity and clarity are innumerable (perhaps infinite, if they belong to a continuum). Accordingly, if phenomenal states are partly individuated by their degree of intensity and clarity, the number of phenomenal states is extremely high (perhaps infinite)—much higher than if quantitative phenomenal properties were accidental. Finally, against the idea that relational phenomenal properties are essential, is the intuition that a certain phenomenal state may well have existed independently of having the relational properties it actually has: it may have existed independently of being unified with the phenomenal states with which it is actually unified, independently of being part of the overall experience is actually part of, independently of occurring before and after the phenomenal states which actually occur after and before it, and independently of being more or less salient than the phenomenal states it is actually more or less salient than. You may well have the same visual experience of the screen before you even if you were not concomitantly hearing children playing in the courtyard and feeling pain in your right knee. Similarly, you may have felt the very same pain sensation in your toe even if the visual experience of the written sheet of paper had not preceded it.

As I said, these are far from being knock-down reasons to prefer (h) over (f) and (g). The integration account, however, is grounded in the assumption that a phenomenal state's qualitative properties are essential and that its

⁴⁰ I think that theories such as Tye's (2003) and Bayne and Chalmers' (2003) imply that unity is an essential phenomenal property.

⁴¹ I owe this argument to Jorge Morales.

quantitative and relational phenomenal properties are accidental. Alternative accounts may be developed which are grounded on different assumptions.⁴²

4.2. *Integration*

I have argued that a satisfactory account of the metaphysics of primitive introspection must not only entail the metaphysical impossibility of introspective hallucination. It must also explain the phenomenal change a phenomenal state undergoes upon being primitively introspected. Particularly, it must explain (1) why such phenomenal change occurs at every instance of primitive introspection and (2) how primitive introspection is possible notwithstanding such phenomenal modification. Extant acquaintance accounts do entail the impossibility of introspective hallucination, but they fail to explain phenomenal modification. The account developed in this section, what I call the *integration account* of the metaphysics of primitive introspection, satisfies both requirements.

Roughly, the idea is the following. By being partly constituted by the target phenomenal state, the primitive introspective state inherits the target's essential phenomenal properties (i.e. its qualitative phenomenal properties). Therefore, the phenomenology of the primitive introspective state partly depends on that of the target phenomenal state. This explains the metaphysical impossibility of introspective hallucination (i.e. the fact that one cannot be in a state of primitive introspection without any phenomenal state being present). On the other hand, once it is primitively introspected, some of the *accidental* phenomenal properties of the target phenomenal state (i.e. some of its relational or quantitative phenomenal properties) depend on the state of primitive introspection directed at it. Importantly, salience (i.e. the phenomenal property which is modified at every instance of primitive introspection), which is a relational phenomenal property, is among the target phenomenal state's accidental properties. Therefore, once introspected, the target's phenomenal property of salience depends on the state of primitive introspection. This explains the nomological necessity of introspective phenomenal modification. The remainder of this subsection is devoted to developing the integration account more precisely and thoroughly.

Integration is a relationship between phenomenal states (i.e. conscious states with phenomenology). Here is my characterization of *integration*:

A phenomenal state ϕ is integrated in a phenomenal state ι iff:

⁴² I briefly explore a potential alternative account in *Appendix 2*.

- (i) φ is a proper part of ι ;
- (ii) ι inherits all of φ 's essential properties;
- (iii) (some of) φ 's accidental phenomenal properties depend on ι ;
- (iv) ι refers to φ .

(i) entails that, when φ is integrated in ι , ι is constituted by φ , although not fully constituted by it: there are parts of ι which are not parts of φ . Therefore, ι is not identical to φ , but is partly constituted by φ .

That ι *inherits* all of φ 's essential properties—i.e., (ii)—means that, when φ is integrated in ι , all of φ 's essential properties are also ι 's essential properties, and this in virtue of φ 's being part of ι .

Although all of φ 's essential properties are also ι 's essential properties, the reverse does not apply: there are some essential properties of ι which are *not* also essential properties of φ . For if, as we are assuming here, phenomenal states are individuated by their essential properties, if ι and φ shared all of their essential properties (that is, all and only φ 's essential properties were also essential properties of ι), then ι and φ would be identical—they would be the same state. But, on the integration account, ι and φ are *not* the same state. For one thing, as noted, although φ partly constitutes ι , there are some parts of ι that are not parts of φ . For another thing, given that ι refers to φ —as per condition (iv), the two states must be at least conceptually distinct. Therefore, since the two states are not identical, there must be some essential properties of ι which are *not* also essential properties of the target phenomenal state. Spelling out *which* are the essential properties of ι that are *not* also φ 's essential properties is not straightforward and exceeds our present purpose. The explanatory power of the integration account remains intact even if this specific question is left open. (I nonetheless try to make some progress and explore potential answers to this question in *Appendix 1*).

As a consequence of (ii), ι cannot be present without φ being also present: without φ , some of ι 's essential phenomenal properties would not be present, therefore (since, recall, phenomenal states are individuated by their essential phenomenal properties) ι could not exist.

As per (iii), when φ is integrated in ι , some of φ 's accidental phenomenal properties depend on ι . Thus, were φ not integrated into ι , some of its accidental phenomenal properties would have been different.

Finally, as per (iv), when φ is integrated into ι , ι refers to φ : in other words, ι is directed to, or is about, φ .

My substantial claim about the metaphysics of primitive introspection is that the relationship between the target and the state of primitive

introspection is integration: when one primitively introspects, the target phenomenal state is integrated in the state of primitive introspection. This means that, when the subject primitively introspects a certain phenomenal state,

- (i) the target phenomenal state becomes a proper part of the state of primitive introspection;
- (ii) all of the target phenomenal state's essential properties are inherited by the state of primitive introspection;
- (iii) (some of) the target phenomenal state's accidental phenomenal properties depend on the state of primitive introspection;
- (iv) the state of primitive introspection refers to the target phenomenal state.

Condition (i) entails that the state of primitive introspection is partly constituted by the target phenomenal state, which makes the integration account a version of the acquaintance view of the nature of primitive introspection.

As noted, condition (ii) entails that the integrating state cannot exist independently of the integrated state. By applying this to primitive introspection we have that the state of primitive introspection cannot exist independently of the target phenomenal state. This implies *no introspective hallucination*: one cannot be in a state of primitive introspection if no phenomenal state is present. In virtue of (ii), the no-introspective-hallucination desideratum is also explained. It is because it inherits the essential phenomenal properties of the target phenomenal state that the state of primitive introspection cannot exist if no phenomenal state is present: if no phenomenal state is present, some of the essential phenomenal properties of the state of primitive introspection cannot be present; since phenomenal states are individuated by their essential properties, and the state of primitive introspection is a phenomenal state, if some of its essential phenomenal properties are absent, the state of primitive introspection does not exist.

Condition (iii) is that some of the target phenomenal state's accidental phenomenal properties depend on the state of primitive introspection. As noted in §4.1.3, we are assuming here that a phenomenal state's relational and quantitative phenomenal properties are accidental. We saw in §1.2 that, when it is primitively introspected, the target phenomenal state undergoes a change in phenomenology: every time one primitively introspects a preexisting phenomenal state, the latter undergoes a change in salience. Moreover, at least sometimes, it undergoes a change in intensity and clarity. Such phenomenological change depends on the state of primitive introspection. As

noted in §4.1.2, salience is a relational phenomenal property; intensity and clarity are quantitative phenomenal properties. Therefore, the phenomenal properties that change when a certain phenomenal state is primitively introspected are accidental phenomenal properties. Condition (iii) accounts for the fact that the phenomenological change undergone by the target phenomenal state depends on the state of primitive introspection. It depends on the state of primitive introspection because the phenomenal properties that change are accidental and, when a phenomenal state is integrated in a state of primitive introspection, such accidental properties depend on the state of primitive introspection.

Finally, condition (iv) is aimed to rule out cases in which (i), (ii), and (iii) are satisfied in the absence of primitive introspection. It may be argued that the relationship between a phenomenal state and the overall experience to which it belongs satisfies all of (i), (ii), and (iii).⁴³ On some views of the unity of consciousness, a phenomenal state ϕ is a phenomenal part of the overall experience E (i). In virtue of having ϕ as a part, E inherits ϕ 's essential properties (ii) and thereby depends on ϕ for its existence. The other phenomenal states which compose E may have an effect on ϕ 's accidental properties (even on those associated with salience). If, for instance, during the jazz concert the volume of all instruments is lowered, except that of the sax, the sound of the sax is likely to become more salient. If it is always the case that some of ϕ 's accidental phenomenal properties depend on the other phenomenal states which compose E , then condition (iii) is satisfied too. But, of course, E is *not* a state of primitive introspection. What E lacks and the state of primitive introspection has is *reference* to the target phenomenal state: the state of primitive introspection is *about* the target phenomenal state, but E , the overall experience, is *not* about any of its parts.

As noted in §1.2, to satisfy the phenomenal modification desideratum, an account of the nature of primitive introspection must (1) account for the fact that phenomenal modification associated with salience occurs at every instance of primitive introspection and (2) explain how one can primitively introspect notwithstanding phenomenal modification. The integration account satisfies both requirements. Condition (iii) explains why such a change in accidental phenomenal properties occurs at *every time* one primitively introspects—as per (1). Every time it is integrated, some of the target phenomenal state's accidental phenomenal properties associated with salience depend on the presence of the state of primitive introspection: the target phenomenal state's being phenomenally more salient depends on the

⁴³ Thanks to Angela Mendelovici and Geoff Lee for this remark.

fact that it is primitively introspected. When not primitively introspected, the target phenomenal state is less salient. Therefore, every time it is primitively introspected, the target phenomenal state undergoes a change in the phenomenal properties associated with salience—it becomes more salient.

Condition (iii) also explains why the target phenomenal state, *post-change* in phenomenal properties associated with salience, depends on the state of primitive introspection. The accidental phenomenal properties associated with salience, those that change when the target phenomenal state is primitively introspected, are those which depend on the state of primitive introspection. Therefore, the target phenomenal state, *as modified* by the act of primitive introspection, i.e., with the accidental properties it has upon being primitively introspected, depends on the state of primitive introspection.

However, the integration account is compatible with the target phenomenal state's having independent existence from the state of primitive introspection. Since the target phenomenal state ϕ has the same essential properties before and while being introspected, and phenomenal states are individuated by their essential properties, post-change ϕ is the same state as pre-change ϕ , though it is different in that (some of) its accidental properties have changed. The integration account is thus consistent with the intuition, accommodated by the other acquaintance accounts, that phenomenal states can exist independently of being introspected.

That said, we can see how the integration account satisfies condition (2) of the phenomenal modification desideratum. Even though the target phenomenal state undergoes some phenomenal changes upon being primitively introspected, such phenomenal changes do not concern its essence: they only concern its accidental phenomenal properties. Therefore, what is integrated into the state of primitive introspection is exactly the same phenomenal state which preexists the act of primitive introspection and is targeted by it. Primitive introspection is possible notwithstanding phenomenal modification because only the accidental phenomenal properties are modified by the act of primitive introspection—the essential properties which individuate the target phenomenal state remain unchanged.

The integration account, then, satisfies both the no-introspective-hallucination desideratum and the phenomenal modification desideratum. It satisfies the former in virtue of condition (ii). It satisfies the latter in virtue of condition (iii). For these reasons, I take it to be a good account of the metaphysics of primitive introspection.

Conclusion

In this chapter, I have proposed an account the nature of primitive introspection. A satisfactory account of primitive introspection must fit two desiderata: the *no introspective hallucination* desideratum and the *phenomenal modification* desideratum. I have argued that neither the inner sense account nor extant versions of the acquaintance account satisfy both desiderata. The inner sense account may satisfy *phenomenal modification*, but it does not satisfy *no introspective hallucination*. Extant versions of the acquaintance account do satisfy *no introspective hallucination*, but they do not satisfy *phenomenal modification*.

I have then developed what I have called the *integration account* of primitive introspection. The integration account is a version of the acquaintance view. I showed that the integration account satisfies both the *no introspective hallucination* desideratum and the *phenomenal modification* desideratum. This is a good reason to think that the integration account is a satisfactory account of the metaphysics of primitive introspection.

The integration account relies on some assumptions about the metaphysics of the mind which I have tried to spell out thoroughly. As noted, based on different assumptions, alternative accounts of the metaphysics of primitive introspection could be explored. Thus, it may be that the integration account is not the *sole* candidate account of the metaphysics of primitive introspection. Provided that they satisfy the abovementioned desiderata, alternative accounts, to be yet developed, may constitute potential competitors to the integration account which has been developed here.

Appendix 1. Unshared essential properties

Condition (ii) of the integration account is that, when φ is integrated in ι , ι inherits all of φ 's essential properties. Applied to primitive introspection, condition (ii) says that, when the target phenomenal state is primitively introspected (and thereby integrated in the state of primitive introspection), all of its essential properties are inherited by the state of primitive introspection. This implies that, in primitive introspection, all the essential phenomenal properties of the target phenomenal state are also essential properties of the state of primitive introspection.

Now, the question arises whether the essential properties of the target phenomenal state *exhaust* the essential properties of the state of primitive introspection. That is, whether all the essential properties of the state of primitive introspection are also essential properties of the target phenomenal

state or there are some essential properties of the state of primitive introspection that are *not* also essential properties of the target phenomenal state.

Since, as we are assuming here, phenomenal states are individuated by their essential properties, if the primitive introspective state and the target phenomenal state shared all of their essential properties (that is, all and only the target state's essential properties were also essential properties of the state of primitive introspection), then they would be identical—they would be the same state. As I pointed out in §4.2, however, on the integration account the state and the target of primitive introspection are *not* the same state. Therefore, there must be some essential properties of the state of primitive introspection which are *not* also essential properties of the target phenomenal state.

The distinction between the state and the target of primitive introspection (that is, the idea that they are not the same state but distinct states) is not merely dogmatic but is theoretically motivated. Even if the relationship between them is constitution rather than causality (as an inner sense account would entail), we can (and should) still *distinguish* the target phenomenal state and the state of primitive introspection. For every state of primitive introspection aims at a phenomenal state: to be introspective, a mental state must be *about* a phenomenal state. Therefore, at any instance of primitive introspection, we must conceptually distinguish between the introspective state and the target state it aims at. If all and only the target state's essential properties were also the introspective state's essential properties, we couldn't distinguish them (since phenomenal states are individuated by their essential properties). Since, on the integration account, all of the target state's essential properties are also essential properties of the primitive introspective state, some of the introspective state's properties must not be also essential properties of the target state. But what are those unshared essential properties?

First of all, the question arises whether the unshared essential properties are *phenomenal*.⁴⁴ As noted in §4.1.3, it is an assumption of the integration account that, among the phenomenal properties of a phenomenal state, all and only the *qualitative* ones are essential. On this assumption, if the unshared essential properties are phenomenal, they must be qualitative. Yet, it is far from straightforward *what* these unshared *qualitative* phenomenal properties are supposed to be, since, at least *prima facie*, there do not seem to be any *qualitative* phenomenal properties to the primitive

⁴⁴ My reflection on this issue has benefitted a lot from insightful comments by Luke Roelofs.

introspective state on top of the qualitative phenomenal properties of the target phenomenal state. When you primitively introspect a pain sensation in your right knee, there do not seem to be any qualitative phenomenal properties to your primitive introspective state on top of the painful character the introspective state inherits from the introspected pain sensation.

What about their being non-phenomenal? The integration account is based on the assumption that, among a phenomenal state's *phenomenal* properties, all and only the qualitative ones are essential. As it has been developed up to now, however, the integration account is neutral about whether a phenomenal state's qualitative phenomenal properties exhaust its essential properties—that is, whether all the essential properties of a phenomenal state are (qualitative) phenomenal. One option would be to endorse the idea that not only the qualitative phenomenal properties, but also some *non-phenomenal* properties are essential. On this view, all the qualitative phenomenal properties *and some non-phenomenal properties* are essential to a phenomenal state. This seems to me a promising way to go. Spelling out what these unshared non-phenomenal essential properties are, though, is not obvious. Here is a potential track to explore. As noted, what helps us distinguish a primitive introspective state from the phenomenal state it is directed at is that the primitive introspective state is indeed *directed at* the relevant phenomenal state—it is *about* or *refers to* that phenomenal state. This suggests that *being directed at a phenomenal state* may be the property which fundamentally distinguishes a primitive introspective state from a merely conscious phenomenal state. This is especially plausible if we assume first-order representationalism about phenomenal consciousness (Dretske 1995; Tye 1995). First-order representationalism about consciousness is, roughly, the theory that a mental state is phenomenally conscious in virtue of its representing things and their properties. On this view, conscious states are intentionally directed toward the objects in their subject's environment, but they are not (at least not necessarily) themselves the intentional object of a conscious state. So, in a first-order representationalist framework, the essential property that the state of primitive introspection has, but the target phenomenal state lacks, is that of *being directed at a phenomenal state*.

If, however, we adopt a different theory of phenomenal consciousness, this explanation may no longer be available. On a self-representationalist view about phenomenal consciousness (Kriegel 2009), for a mental state to be conscious it is not sufficient that it be intentionally directed toward the objects in one's environment: it must also be intentionally directed toward *itself*. On this view, then, the phenomenal state which is the target of primitive introspection does have the property of being directed at a phenomenal state:

it is directed at itself. Since this is (part of) what makes such a phenomenal state conscious, the target phenomenal state has that property *essentially*. Therefore, in a self-representationalist framework, *being directed at a phenomenal state* cannot be the unshared essential property we are looking for. However, there is still a difference to be drawn, in the self-representationalist framework, between the kind of reference that characterizes the (merely conscious) phenomenal state and the kind of reference that characterizes the state of primitive introspection. Whereas the state of primitive introspection refers to a phenomenal state *attentively*, the merely conscious phenomenal state refers to a phenomenal state (that is, it refers to itself) *inattentively*. Importantly, attentiveness is a necessary feature of a primitive introspective state (as noted in Chapter 3, attention is a necessary condition of primitive introspection). Therefore, in a self-representationalist framework, the unshared essential property of the primitive introspective state is *being attentively directed at a phenomenal state*.

The idea hinted at in the previous paragraph is thus that attentive reference to a phenomenal state is a non-phenomenal essential property of the state of primitive introspection. Proponents of *phenomenal intentionality*, though, would probably object that attentive reference to a phenomenal state is *phenomenal* reference—it is a phenomenal kind of intentionality. On their view, if attentive reference to a phenomenal state is the extra essential property of the primitive introspective state, then such essential property is phenomenal. Indeed, a view along these lines strikes me as antecedently plausible. A full development of it, though, would require careful further theoretical work, which exceeds the scope of this dissertation but would be worth carrying out in the future. Particularly, it would require specifying whether the relevant essential phenomenal property is *qualitative* or rather of a yet different type (if so, there would be a fourth kind of phenomenal property besides qualitative, quantitative, and relational).

The considerations sketched in this appendix are tentative and should be considered as explorative hypotheses for further development of the integration account rather than assertive solutions to the problem of unshared essential properties.

Appendix 2. Sketch for an alternative account of primitive introspection's metaphysics

I have insisted that the integration account may not be the *sole* satisfactory account of the metaphysics of primitive introspection. The integration account is grounded in some assumptions about the metaphysics of the mind—that phenomenal properties come into three kinds, that phenomenal states have essential and accidental properties, and that, among a phenomenal state's phenomenal properties, all and only the qualitative ones are essential. Based on different assumptions, however, alternative accounts may be developed. Here I draw the sketch of a potential alternative account of the metaphysics of primitive introspection. Let us call it the *identity account*.

On the identity account, the relationship between the state and the target of primitive introspection is *identity*: they are the *same state*.⁴⁵ As the other acquaintance accounts, the identity account entails that there is a dependence relationship between the target and the state of primitive introspection. As on the integration account, such dependence is not unilateral but bilateral—it is *interdependence*. However, differently from all other versions of the acquaintance account, the identity account implies that the phenomenal state that pre-exists the primitive introspective act, that is, the phenomenal state which constitutes the initial target of that act, is *not the same state* as the phenomenal state which is actually primitively introspected. When introspective attention is directed toward it, the originally targeted phenomenal state is somehow 'destroyed' and replaced by the state of primitive introspection (which coincides with the new, actual target of the primitive introspective act).

An account along these lines may find the sympathies of those who, like Franz Brentano, believe that phenomenal states cannot be the object of observation (i.e. the object of one's introspective attention). He writes:

In observation, we direct our full attention to a phenomenon in order to apprehend it accurately. But with objects of inner perception this is absolutely impossible. This is especially clear with regard to certain mental phenomena such as anger. If someone is in a state in which he wants to observe his own anger raging within him, the anger must already be somewhat diminished, and so his original object of observation would have disappeared. The same impossibility is also present in all other cases. It is a universally valid psychological law that we can never focus our *attention* upon the object of inner perception. (Brentano 1874: 30)

⁴⁵ Although I had already been considering a view along these lines for a while, the identity account was explicitly suggested to me by François Récanati.

Arguably, underlying this reasoning is the rejection of the idea that quantitative and relational properties are merely accidental. The reasoning seems to be the following:

- (P1) Introspective attention modifies a phenomenal state's phenomenology.
- (P2) If a phenomenal state undergoes a change in phenomenology, it ceases to exist.
- (C) Once introspective attention is directed upon it, a phenomenal state ceases to exist.

The integration account, as other extant acquaintance accounts, rejects (P2): the fact that some aspects of its phenomenology change does not entail that a phenomenal state cease to exist. On the integration account, if what changes is a quantitative or relational phenomenal property, the relevant phenomenal state continues to exist.

On the identity account, instead, *any* change in phenomenology entails the destruction of a phenomenal state. Arguably, on this account, *all* the phenomenal properties of a phenomenal state (not only the qualitative ones, but also the relational and quantitative ones) are essential to it.

On the identity account, then, a subject is in a state of primitive introspection ι directed toward a phenomenal state φ only if $\iota = \varphi$. Arguably, further conditions should be articulated for the account to be viable, for the identity relationship may not be sufficient to explain why ι is an introspective state at all. At the very least, the account needs a story about what makes ι intentionally directed toward φ if ι and φ are the same state. The core idea, at any rate, is that, once one primitively introspects, the state of primitive introspection and the target phenomenal state are identical.

As the other acquaintance accounts, the identity account satisfies the no-introspective-hallucination desideratum. It also partly satisfies the phenomenal modification desideratum, for it accounts for (1) the fact that, at any time one primitively introspects, the phenomenology of the target phenomenal state changes. However, the identity account does not seem to have a straightforward explanation of (2) how primitive introspection is possible at all, given phenomenal modification. For if the initial target of primitive introspection is destroyed by the very act of primitive introspection, then primitively introspecting any phenomenal state seems to be impossible. (This indeed seems to be, roughly, Brentano's conclusion.)

The main challenge for the identity account is thus to find a solution to the problem of how, if introspective attention destroys its target, primitive

introspection is possible at all. Perhaps something along the following lines may be a potential way to explore. Although the initial target of the primitive introspective act and the actual target of the primitive introspective state are different phenomenal states, they share many of their essential properties. After all, only *a few* phenomenal properties are modified by the introspective act (i.e. those associated with salience and, sometimes, those associated with intensity or clarity). All the qualitative phenomenal properties remain unchanged. Accordingly, at the very least, the original target and the actual target have the same qualitative phenomenal properties. Although they are not the *same* state, they are very *similar*. Therefore, introspectively acquiring information about the actual target, although it does not provide us with *all* the information about the original target, provides us with *some* information about it—enough information for us to claim that by primitively introspecting the actual target we can nonetheless introspectively know the original target.

To be sure, more work needs to be done if the identity account is to satisfyingly fulfill the phenomenal modification desideratum—work I am not going to do here. This appendix’s goal was mainly to point at an alternative way to explain the metaphysics of primitive introspection, a way perhaps worth exploring for those who want to resist the assumption that only qualitative phenomenal properties are essential phenomenal properties of a phenomenal state.

*

To repeat, what I argued for in the bulk of this chapter is that the integration account is *a* good account of the metaphysics of primitive introspection, although I have *not* argued that it is the *only* satisfactory account. Indeed, by the present appendix, I have tried to convey a real aspect of my reflection about the metaphysical structure of primitive introspection, namely that perhaps the integration account is not the only way to go and that *other* accounts may be developed, which may be no less satisfactory than the integration account. Still, I think that the tools deployed here to elaborate the integration account (imposing two desiderata—the no-introspective-hallucination desideratum and the phenomenal modification desideratum—to any account of the metaphysics of primitive introspection; identifying in the phenomenal feature(s) associated with salience what necessarily changes at any instance of primitive introspection; distinguishing between qualitative, quantitative, and relational phenomenal properties; distinguishing between essential and accidental properties of phenomenal states) can constitute a

useful common ground upon which different accounts of the metaphysics of primitive introspection may be built. At the very least, I hope that this chapter's reflections could be helpful in reflecting about the metaphysical structure of primitive introspection.

PART 3: EPISTEMOLOGY

CHAPTER 5

PRIMITIVE INTROSPECTION AS *SUI GENERIS* KNOWLEDGE BY ACQUAINTANCE

In Chapters 3 and 4 I explored the *nature* of primitive introspection. Particularly, I investigated the metaphysical structure of the *state* of primitive introspection, that is, the mental state a subject is in in virtue of primitively introspecting. In this chapter and the next, I will focus on the *epistemology* of primitive introspection. I will argue that primitive introspective states are *epistemically significant*. By being in a state of primitive introspection directed toward a certain phenomenal state, a subject acquires knowledge of that phenomenal state. More specifically, the primitive introspective state *constitutes* a specific kind of knowledge, that is, *knowledge by acquaintance*.

Knowledge by acquaintance is (roughly) knowledge we have of that which we are *directly aware* of. In Chapter 6, I will argue that introspective knowledge by acquaintance (that is, knowledge constituted by the state of primitive introspection) is epistemically significant. In this chapter, I argue that knowledge by acquaintance is a *sui generis* kind of knowledge: it is irreducible to propositional knowledge, as well as to other (putative) kinds of knowledge such as knowing-how and knowledge of a subject matter.

First, I spell out the notion of *knowledge by acquaintance*, partly by appeal to Bertrand Russell's (1912) seminal characterization of it (§1). I also illustrate what it means for knowledge by acquaintance to be *sui generis*, and what it takes for it to be irreducible to other kinds of knowledge (§2). Secondly, I present some cases in which one intuitively seems to have some kind of knowledge which exceeds possession of propositional knowledge (or of any other putative kind of knowledge); on this basis I argue that there is *prima facie* reason to believe that knowledge by acquaintance is a *sui generis* kind of knowledge (§3). I also consider some ways those intuitions may be challenged and explain why, at a closer examination, the challenges are not really threatening (§4). Thirdly, I consider two objections to the claim that knowledge by acquaintance is *sui generis*: the objection from *disunity* (knowledge is a unified rather than scattered notion) and the objection from *mysteriousness* (the notion of knowledge by acquaintance is unintelligible unless reduced to some other kind of knowledge). I show that these potentially threatening objections can be answered and that knowledge by acquaintance being a *sui generis* kind of knowledge remains a live option on the table (§5).

Finally, I briefly explain why sui generis knowledge by acquaintance is important (§6).

*

1. Knowledge by acquaintance

The notion of *knowledge by acquaintance* is introduced by Bertrand Russell in chapter 5 of *The Problems of Philosophy*.^{1:2} Knowledge by acquaintance is a kind of what Russell calls ‘knowledge of things’, which he contrasts with ‘knowledge of *truths*’. Knowledge of truths is the sort of knowledge one has when one “know[s] *that* something is the case” (1912: 69). Knowledge of truths, then, is *propositional knowledge*: knowledge partly consisting in the subject forming a *judgment* about what is known. Knowledge of things may come in two kinds: knowledge of things *by acquaintance* and knowledge of things *by description*. Knowledge by acquaintance is “essentially simpler than any knowledge of truths, and logically independent of knowledge of truths”. Knowledge by description, by contrast, “always involves [...] some knowledge of truths as its source and ground” (Russell 1912: 72-73). Knowledge by acquaintance is then spelled out in terms of *direct awareness*: “we have *acquaintance* with anything of which we are directly aware, without the intermediary of any process of inference or any knowledge of truths” (Russell 1912: 73). Knowledge by acquaintance is therefore a kind of knowledge whose object is known *immediately*, and that does not depend on the subject’s forming *any judgment* about what is known. What can be known by acquaintance, on Russell’s view, is restricted to *sense data* (i.e. mental items, like *color* and *shape*, by being aware of which, according to Russell, we are aware of physical objects and their properties) and *universals* (i.e. “general ideas, such as *whiteness*, *diversity*, *brotherhood*, and so on” [1912: 81]).³

¹ It is actually pre-introduced at the end of chapter 4, but chapter 5 is where he spells the it out more thoroughly. Also, the notion had already been presented in his earlier (1910) article ‘Knowledge by Acquaintance and Knowledge by Description’.

² In fact, the idea that direct apprehension constitutes a kind of knowledge which is not only distinct from, but also more basic than, propositional knowledge was introduced much earlier. The idea is present in Plato and Aristotle as well as in later authors such as Aquinas and Spinoza (Hayner 1969; Peterson 2008: 92-93). Thanks to Ben Koons and Kara Richardson for pointing out to me that knowledge by acquaintance was already discussed by Aristotle and Aquinas respectively.

³ Russell also explores the hypothesis that we are acquainted with our *selves* (Russell 1912). Although he attributes a fairly high credence to that hypothesis, he is very careful not to

Russell illustrates the difference in *kind* between knowledge by acquaintance and knowledge by description by way of an example along the following lines. By perceiving a table, Russell argues, you are acquainted with a color sense datum. In virtue of being acquainted with it, you know that color sense datum: you know it by acquaintance. Now, you may learn many truths about the sense datum: that it is yellow, that it is bright, that it is light, etc. But your knowing all these facts does not improve, or affect in any way, your knowledge by acquaintance *of the color sense datum* itself, knowledge you have already got by being acquainted with it. So, no matter how many truths you know about a given sense datum, that will not make you know *it* (by acquaintance) any better.

More generally, knowledge by acquaintance is ‘objectual’ knowledge, in the sense that what is known by acquaintance is an *item*, rather than a proposition.⁴ It consists in a *direct awareness* of its ‘object’, where the relevant awareness is *direct* both epistemically and metaphysically. It is *epistemically* direct in that, by being acquainted with *x*, a subject *S* gets a *non-inferential* access to *x*. Therefore, being acquainted with *x* does not depend on having epistemic access to anything else (importantly, it does not depend on the subject’s forming any *judgment* about the relevant ‘object’). It is *metaphysically* direct in that, when *S* is acquainted with *x*, no state or process mediates between *x* and *S*’s awareness of *x* (Gertler 2011). To better understand the notion of metaphysical directness, compare acquaintance with perception. It is often maintained that perception consists in representing objects in one’s environment (and their features). On this view, when one perceives a certain object *o*, one has a mental representation of *o* which is distinct from and caused by *o*. The relation between one’s awareness of *o* (the mental representation) and *o*, in this case, is *not* direct: a causal process mediates between them. By contrast, if one is *acquainted* with *o*, there is nothing mediating between *o* and one’s awareness of *o*: the relationship is metaphysically direct.⁵

The contrast with perceptual representation suggests a further way in which acquaintance is direct. Differently from perceptual knowledge (as

commit to it. I will leave this issue aside here, since it does not directly bear on to the present chapter’s focus.

⁴ I leave the question open here whether the relevant item is an object, a property, a trope, or something else. ‘Objectual’ should therefore be read as a broader notion than ‘object-directed’, possibly including ‘property-directed’, ‘trope-directed’, and so on. *Mutatis mutandis*, the same applies to ‘object’ in the next two sentences.

⁵ As we saw in Chapter 4, this metaphysically direct relationship is often spelled out by acquaintance theorists in terms of *constitution*. I will come back to this in §5.

characterized by the representationalist model sketched in the previous paragraph), knowledge by acquaintance lacks any representational medium.⁶ Whereas, in the perceptual case described above, the subject's awareness of *o* is mediated by a perceptual representation of *o*, when the subject is acquainted with *o*, no representation mediates between *o* and the subject's awareness of *o*. So, besides being epistemically and metaphysically direct, acquaintance is also *representationally* direct.

Knowledge by acquaintance, then, is by its nature non-propositional: the 'object' is immediately apprehended, and no description is attached to it (where attaching a description to an object *o* implies ascribing a certain property *F* to it, and thereby forming the proposition 'o is *F*'). Indeed, no amount of propositional knowledge has any effect on knowledge by acquaintance: when you know something by acquaintance, no matter how many truths you learn about the thing, that will leave your knowledge by acquaintance of it unchanged. Knowledge by description, by contrast, is constitutively propositional. The distinction between knowledge by acquaintance and knowledge by description is mirrored, in some languages other than English, by the existence of two different words for 'knowledge': *connaître* and *savoir* in French, *conoscere* and *sapere* in Italian, *kennen* and *wissen* in German.⁷ Whereas *savoir* (and analogs) takes a that-clause ('Je sais *que tu es à Paris*'), *connaître* only takes a direct object ('Je connais *Jeanne*').

Having been neglected for a while after Russell, the notion of *acquaintance* has regained momentum in more recent literature (Conee 1994; Fumerton 1995, 2009; BonJour 2000, Gertler 2001, 2012; Chalmers 2003; Pitt 2004; Horgan and Kriegel 2007; Balog 2012), and a deeper analysis of the acquaintance relation and of its role in metaphysics and epistemology seems to be emerging. However, virtually none of those interested in acquaintance seems to address the notion of *knowledge by acquaintance* as Russell himself conceived it. And indeed, the idea that knowledge by acquaintance is a *sui generis* kind of knowledge, independent from and irreducible to propositional knowledge, is now almost unanimously discarded. What most contemporary philosophers call 'knowledge by acquaintance' is *propositional* knowledge which is *directly based on* the relation of acquaintance. It is interesting to remark that both the *Internet Encyclopedia of Philosophy* and the *Stanford Encyclopedia of Philosophy* entries on the distinction between knowledge by acquaintance and knowledge by description attribute the idea of knowledge

⁶ Thanks to Andrew Lee for calling my attention to this notion.

⁷ The same remark is made by Russell (1912: 70). Spanish and Portuguese have two distinct words for 'knowledge' too. Unfortunately, I do not know about other languages.

by acquaintance being a *sui generis* kind of knowledge to confusion or equivocation:

The traditional account of knowledge by acquaintance is susceptible to being misunderstood or conflated with merely being directly acquainted with something [...] For a subject to be directly acquainted with something only requires for the subject to have unmediated access to the object of awareness. Knowledge by acquaintance that something is the case, however, [...] is a kind of *knowledge*, which requires the subject to hold a belief under the right conditions. (DePoe 2018: §1, emphasis in original)

It is tempting to suppose that Russell equivocates between the relation of acquaintance and the special kind of knowledge of truth (foundational knowledge) whose sole source is acquaintance. [...] In order to guard against confusing *acquaintance* on the one hand with *foundational knowledge of truths acquired by acquaintance* on the other [...] we should be more careful than Russell and restrict *knowledge by acquaintance* to foundational knowledge of *truths*. By contrast, one can have *acquaintance* with items that are not truths, items that cannot be said to be true or false. (Hasan and Fumerton 2017: §1, emphasis in original)

And contemporary acquaintance theorists often explicitly reject Russell's "idiosyncratic" idea that acquaintance suffices for knowledge. This quote from Gertler (2011) is an example:

The epistemic features of Russell's view are questionable as well. [...] [H]e appears to take acquaintance with sense data to suffice for knowledge of those sense data. [...] It is hard to understand how one could *know* one's sense data simply by being aware of them in the minimal way required for ordinary perception. Knowledge of an object seems to require thinking *about* the object, which in turn involves some *way* of thinking of it. [...]

Contemporary acquaintance theorists [...] deny that acquaintance suffices for knowledge [...] and construe introspective knowledge as knowledge of truths. (Gertler 2011: 92-94)

The idea that Russell's conception of knowledge by acquaintance is due to equivocation is perhaps motivated by the fact that the expression 'knowledge by acquaintance' is ambiguous, due to ambiguity between a causal and a constitutive reading of 'by' (compare: 'I made him happy by waving my hand' vs. 'I said hello by waving my hand').⁸ On the first reading, 'knowledge by acquaintance' means 'knowledge which is caused, or produced, by acquaintance'. On the second reading, 'knowledge by acquaintance' means 'knowledge which is constituted by acquaintance'. I agree that the expression

⁸ Uriah Kriegel, personal communication.

'knowledge by acquaintance' is ambiguous in the way just described. However, different from contemporary epistemologists and acquaintance theorists, I believe that the constitutive reading of the expression is *not* the result of equivocation. On the contrary, it is a well-aware choice which is part of a substantial thesis: regardless of whether there is knowledge which is *caused* by acquaintance, there is knowledge which is *constituted* by acquaintance.

To my knowledge, the only exception to this skeptical trend about the constitutive approach to knowledge by acquaintance is Earl Conee.⁹ Conee (1994) explicitly endorses the Russellian spirit when he introduces the notion of knowledge by acquaintance in his physicalist reply to Jackson's (1982) knowledge argument.¹⁰ Similarly to proponents of the so-called *ability hypothesis* (Lewis 1990), Conee argues that, although Mary does acquire *new knowledge* when she exits the black-and-white room, such knowledge is *not* knowledge of a *fact*. However, Conee rejects the central tenet of the ability hypothesis, namely that *knowing what it's like* consists in the possession of a certain sort of abilities (i.e. remembering, imagining, and recognizing the experience).¹¹ Instead, he maintains, *knowing what it's like* consists in *acquaintance* with the experience, where "[a]cquaintance constitutes a *third category of knowledge, irreducible* to factual knowledge or knowing how. Knowledge by acquaintance of an experience requires only a maximally direct cognitive relation to the experience." (Conee 1994: 136, my emphasis). Knowledge by acquaintance, on Conee's view, is neither knowledge of a fact nor knowledge of an ability; rather, it is knowledge of a *property* (a phenomenal property). From inside her black-and-white room, Mary *does know* all the facts about phenomenal redness, including the fact that it is a property of experiences. However, what she *does not* know from inside the black-and-white room is *phenomenal redness*, i.e. *the property itself*, and this

⁹ An earlier defense of the Russellian conception of knowledge by acquaintance was attempted by Hayner (1969), though Hayner's arguments strike me as too weak. Hayner cites Ducasse (1953) as a proponent of a distinction between two kinds of knowledge, 'experient knowledge' and 'scient knowledge', which is similar to Russell's.

¹⁰ As most readers will know, Jackson's argument rotates around Mary the neuroscientist, who knows all the physical facts about color and color vision, although, having grown up in a black-and-white room, she has never had any visual experience of color. When Mary leaves her black-and-white room and sees a red object, Jackson argues, Mary learns a *new fact*: she learns *what it is like* to see red. Since Mary already knew all the relevant physical facts before leaving her room, Jackson concludes, what Mary learns cannot be a physical fact.

¹¹ He argues that ability possession is neither necessary nor sufficient for *knowing what it's like*.

simply because phenomenal redness is not a property of any of her visual experiences.

In the bulk of this chapter, I will defend the Russellian thesis that knowledge by acquaintance is a *sui generis* kind of knowledge. The thesis I defend is ‘Russellian’ rather than ‘Russell’s’ because, while being true to the spirit of Russell’s view, it involves a notion of *knowledge by acquaintance* which differs in part from Russell’s. For one thing, it does not commit to the existence of *sense data* or to the claim that we can be acquainted with universals.

Indeed, what I am interested in here is knowledge by acquaintance of *phenomenal states*, i.e. conscious mental states with phenomenology (i.e. states there is something *it is like* to be in). The focus of this chapter is therefore knowledge by *introspective* acquaintance. This, of course, mainly because introspective acquaintance is the kind of acquaintance which is yielded by primitive introspection. There are, however, independent reasons for restricting the present study to introspective acquaintance. First, it is *this* kind of acquaintance which is most discussed in the current debate about knowledge by acquaintance. Second, if there is knowledge by acquaintance at all, introspective acquaintance is, arguably, the most plausible candidate to constitute it. To be sure, my focusing on introspective acquaintance does not presuppose that there may not be other kinds of knowledge by acquaintance, such as, for instance, perceptual acquaintance and intellectual (or intuitive) acquaintance (which, though, I am not going to explore here).^{12;13}

I have characterized my main thesis as the thesis that knowledge by acquaintance is a *sui generis* kind of *knowledge*. However, some readers may be unsympathetic to the idea that anything other than propositional knowledge can be claimed to be full-fledged knowledge. For this reason, my ultimate thesis, as I will articulate it in Chapter 6, will be that knowledge by acquaintance is a *sui generis* kind of *epistemic standing*, where the notion of epistemic standing is understood in terms of *epistemic significance*: knowledge by acquaintance is a *sui generis* epistemic phenomenon and it is epistemically

¹² The former could be spelled out in terms of a version of naïve realism, the theory according to which the relationship between the perceiver and a perceived object is direct in the sense that the perceiver’s perceptual experience is partly constituted by the perceived object (see e.g. Campbell 2002). The latter could be explained, for example, in terms of direct apprehension of abstract objects (*cf.* Chudnoff 2013 and Bengson 2015).

¹³ The application of the present discussion of knowledge by acquaintance may be expanded even further. For instance, one could try to explore the idea that we know *moral values*, or *aesthetic values*, by acquaintance.

significant.¹⁴ In this framework, knowledge is a kind of epistemic standing, along with, for instance, *understanding*. I believe that knowledge by acquaintance can legitimately be classified not only as a *sui generis* kind of epistemic standing, but also as a *sui generis* kind of *knowledge*. However, for those who are not sympathetic to this idea, I am happy to retreat to the less committal position. To simplify the exposition, and to be consistent with the terminology most commonly used in the debates I mention, throughout this chapter I stick to the idea that knowledge by acquaintance is a *sui generis* kind of *knowledge*. I will come back to the conceptual distinction between knowledge and epistemic standing in Chapter 6, where I will develop the constructive side of my main thesis. All I say in the present chapter's discussion can be reformulated in terms of epistemic standing rather than knowledge.

In what follows I put forward my defense of the thesis that knowledge by acquaintance is a *sui generis* kind of knowledge (or epistemic standing). In §2 I present what have been sometimes argued to be distinct *sui generis* kinds of knowledge: propositional knowledge, knowing how, and knowledge of a subject matter. I illustrate what it means for a kind of knowledge to be *sui generis* and explain what it takes for knowledge by acquaintance to be irreducible to other kinds of knowledge. In §3 I present some cases in which one intuitively seems to have some kind of knowledge which exceeds possession of propositional knowledge (or of any other putative kind of knowledge); on this basis I argue that there is *prima facie* reason to believe that knowledge by acquaintance is a *sui generis* full-fledged kind of knowledge. In §4 I answer some objections to the effect that the examples in §3 do not even provide *prima facie* motivation in favor of my thesis. In §5 I rebut two objections to the claim that knowledge by acquaintance is *sui generis*: the objection from *disunity* (knowledge is a unified rather than scattered notion) and the objection from *mysteriousness* (the notion of knowledge by acquaintance is unintelligible unless reduced to some other kind of knowledge). I conclude that these potentially threatening objections can be answered and that knowledge by acquaintance being a *sui generis* kind of knowledge remains a live option on the table.

2. Kinds of knowledge

The central concern of epistemology has been, for the most part, *propositional knowledge* or *knowledge-that*. Quite straightforwardly, propositional

¹⁴ More on this in Chapter 6.

knowledge involves the subject bearing a relation to a proposition (knowing that *p*). Knowing that French Revolution occurred in 1789, knowing that there is a red rose before me, knowing that $2+3=5$, and knowing that I am hungry right now are all examples of propositional knowledge. Propositional knowledge was traditionally analyzed in terms of more basic elements: *justified true belief*. Ever since Edmund Gettier (1963) showed that a belief being true and justified does not suffice for it to be knowledge, epistemologists have been looking for the further basic element which makes justified true belief *knowledge*. A huge debate about the analysis of knowledge and the nature of justification followed—though this falls out of my present concern. What is relevant here is that, although knowledge is typically analyzed in terms of more basic notions, and although some of those notions (belief, justification) are *epistemic*, propositional knowledge is usually *not* analyzed in terms of any other *kind of knowledge*.¹⁵

In this sense, propositional knowledge is a *sui generis* kind of knowledge.^{16;17} A kind of knowledge is *sui generis* if it is *irreducible* to any other kind of knowledge and it is *not* a *species* of any other kind of knowledge. Kind *A* is reducible to kind *B* iff *A* can be fully specified in terms of *B*: there is nothing to *A* over and above its being *B*.¹⁸ For instance, heat is reducible to particle kinetic energy: there is nothing over and above to a certain quantity of heat than the aggregate motion of a set of particles. Something could be irreducible to something else but be a species of it. Baking does not reduce to cooking (arguably because baking involves *a certain way* of cooking), although nobody would deny that baking is (a kind of) cooking.¹⁹ A *sui generis* kind of knowledge, then, is neither reducible to other kinds of knowledge nor a species of any other kind of knowledge.²⁰ If there are two (or more) *sui generis*

¹⁵ Some (Williamson 2000) reject the idea that propositional knowledge should be analyzed at all: knowledge is a fundamental epistemic notion and other epistemic notions such as belief and justification are to be analyzed in its terms. The debate around whether knowledge should be treated as the basic epistemic notion is orthogonal to the question I am outlining here. As I said, propositional knowledge being analyzable in terms of more basic epistemic notions is different from its being reducible to, or a species of, another kind of knowledge.

¹⁶ Indeed, some take it to be the *only* kind of knowledge.

¹⁷ However, there is no universal consensus on this. Some philosophers (e.g. Hartland-Swann 1956), for example, attempt to show that propositional knowledge is a species of, or reducible to, knowledge-how. (On the notion of *knowledge-how* see next paragraph.)

¹⁸ Some would probably say that in this case *A* and *B* are not really two different kinds: there is only kind *B* and *A* is at most a *putative* kind. I have no problem with this: my choice to speak in terms of *kind* vs. *sui generis kind* rather than *putative kind* vs. *kind* is purely terminological.

¹⁹ Stanley and Williamson (2001: 434) make the same point with *murdering* and *killing*.

²⁰ Though it may be described as a species of the kind *knowledge*.

kinds of knowledge, although they have in common whatever makes them both kinds of *knowledge*, they are, in an important sense, fundamentally different.

Some philosophers have argued that there is at least another *sui generis* kind of knowledge, besides propositional knowledge, i.e. *knowledge-how*: knowing how to ride a bike, knowing how to play Leonard Cohen's *Suzanne* on the guitar, knowing how to make a proper pizza, knowing how to be a good parent, and so on and so forth. The notion of knowing-how was introduced in contemporary epistemological debate by Gilbert Ryle (1949).²¹ Ryle himself defended the view that knowledge-how and knowledge-that are fundamentally different, and that the former is independent from the latter. Outside epistemology, the notion of *knowing-how* gained special momentum in philosophy of mind after Jackson put forward the abovementioned knowledge argument. As mentioned in the previous section, an attempted strategy to refute Jackson's argument against physicalism consists in appealing to the *ability hypothesis* (Lewis 1990; Nemirow 1990): Mary does gain new knowledge upon exiting her black-and-white room; however, the knowledge she gains is not *factual* knowledge. Rather, what she acquires is *knowledge-how*: she learns how (i.e. acquires the ability) to remember, imagine and recognize a reddish experience.

The *ability hypothesis* draws on Ryle's conception of knowledge-how as a *sui generis* kind of knowledge. Proponents of the *ability hypothesis* thus adopt a *non-reductivist* approach:²² knowledge-how is distinct from and irreducible to knowledge of facts (and thereby distinct from and irreducible to propositional knowledge). Non-reductivists' positive view usually analyzes knowledge-how in terms of *abilities*: knowing how to α is possessing the ability to α . For instance, knowing how to play Leonard Cohen's *Suzanne* on the guitar is possessing the ability to play Leonard Cohen's *Suzanne* on the guitar. The ability account has been challenged though, on the basis that having an ability has been claimed to be neither necessary nor sufficient for knowing-how. A guitarist who lost both hands still knows how to play Leonard Cohen's *Suzanne*, although she lost (along with the hands) the ability to do so (Stanley and Williamson 2001: 416). By mere chance, a novice trampolinist may successfully perform a difficult somersault at the very first attempt; arguably,

²¹ Although notions in the vicinity can be found throughout the history of thought, e.g. the distinction between *technê* and *episteme*, the distinction between practical and theoretical knowledge, and the distinction between procedural and declarative knowledge (Fantl 2017).

²² This approach is often also called *anti-intellectualist*, to underline the idea that knowledge-how does not require prior judgment formation or running through a series of previously known propositions.

nobody would deny that he was able to perform the somersault, although most would agree that he does *not* know how to perform it (Carr 1981: 53).

These challenges have led some philosophers to adopt a *reductivist* approach to knowledge-how.²³ Stanley and Williamson's (2001) is perhaps the most influential argument for reductivism about knowledge-how in the recent literature. Their strategy consists in treating attributions of knowledge-how along with other knowledge attributions containing embedded questions—i.e. attributions of what is sometimes called *knowledge-wh*: knowing *what* to prepare for dinner, knowing *whom* to ask for suggestions, knowing *where* to park one's car, knowing *when* to get ready, knowing *whether* to take one's coat, knowing *why* to take part in the march. Attributions of knowledge-wh occur not only in the *untensed* form displayed by the latter examples (where the embedded clause is untensed), but also in tensed forms like the following: knowing what Vincenzo is preparing for dinner, knowing whom Jill will ask for suggestions, knowing where Camille parked her car, and so on. These are, arguably, attributions of propositional knowledge. And indeed, it has often been argued that knowledge-wh is reducible to knowledge-that: knowing *what* Vincenzo is preparing for dinner is knowing *that* Vincenzo is preparing *pizza* for dinner; knowing *whom* Jill will ask for suggestions is knowing *that* Jill will ask *Jane* for suggestions; knowing *where* Camille parked her car is knowing *that* Camille parked her car *in front of the pastry shop*; and so on. Knowledge-wh attributions are thus attributions of propositional knowledge, where the relevant proposition is the true answer to the embedded question.²⁴ Like knowledge-wh attributions, attributions of knowledge-how may occur in the tensed form too: knowing how Jacque rides a bike, knowing how Leonard Cohen played *Suzanne*, knowing how Vincenzo will make a proper pizza, and so on. Untensed knowledge-how attributions can indeed be rephrased as tensed ones: I know how to play *Suzanne* iff I know, of some way *w*, that *w* is for me a way to play *Suzanne*. Stanley and Williamson take these observations to show that, notwithstanding the superficial linguistic dissimilarities, attributions of knowledge-how are in fact syntactically and semantically similar to attributions of propositional knowledge. Therefore, according to Stanley and Williamson, the syntactic and semantic structure of knowledge-how attributions, far from supporting the claim that knowledge-how is non-propositional, is rather evidence that knowledge-how attributions are

²³ These challenges have also led other philosophers to develop alternative or more sophisticated non-reductivist accounts. Reviewing those, though, falls out of the scope of this chapter.

²⁴ This, arguably, is the received view about the analysis of knowledge-wh. However, it has been challenged (Schaffer 2007).

attributions of propositional knowledge. They conclude that knowledge-how is a species of knowledge-that: it consists in the relationship between a subject and a proposition involving ways of doing something.

So, although the received view is that knowledge-how is a *sui generis* kind of knowledge, there are arguments in the extant literature to the effect that it is not, and that it is rather a species of propositional knowledge. It is not my present purpose to evaluate such arguments here, though, as we will see, reflecting on some of the above considerations may help us get a better grasp of the issues concerning knowledge by acquaintance as a *sui generis* kind of knowledge.

A third kind of knowledge which has received attention in recent literature is what we may call *knowledge of a subject matter*, or *knowledge of a body of information*. This kind of knowledge has sometimes been called ‘objectual knowledge’ because it takes an *object* (the relevant subject matter or body of information), rather than a proposition, as a complement: knowing *mathematics*, knowing *politics*, knowing *quantum mechanics*, and so on. Since there are other kinds of knowledge, which are distinct from knowledge of a subject matter, that may be legitimately labeled ‘objectual’ (knowledge by acquaintance being an example), I stick to the label ‘knowledge of a subject matter’ here, to avoid confusion. The notion of *knowledge of a subject matter* is characterized by analogy with the notion of *objectual understanding*. Literature on epistemic value has recently displayed growing interest in the notion of *understanding* (Kvanvig 2003). Three kinds of understanding have been distinguished: *propositional* understanding (understanding that the position and the momentum of a particle cannot be simultaneously measured with maximal precision), *understanding-why* (understanding why the position and the momentum of a particle cannot be simultaneously measured with maximal precision), and *objectual* understanding (understanding quantum mechanics). The latter involves understanding a subject matter, a structured body of information. Kvanvig (2003) argues that understanding, besides having a number of beliefs about a certain subject matter, involves grasping the dependence, explanatory, and coherence relationships between them. Kvanvig further argues that, since knowledge does not necessarily involve grasping those relationships, understanding is not only distinct from, but also more valuable than knowledge. Brogaard (2005) argues, *contra* Kvanvig, that knowledge and understanding are not distinct: knowledge *is* understanding. What Kvanvig takes to be the difference between knowledge and understanding is, in fact, what makes the difference between *propositional* knowledge or understanding and *objectual* knowledge or understanding (i.e. knowledge or understanding of a subject matter): whereas propositional

knowledge (or understanding) can be piecemeal (one may know a proposition independently of knowing any other proposition and of grasping any relationship between different propositions), knowledge (or understanding) of a subject matter requires having a system of beliefs and grasping the dependence, explanatory, and coherence relationships among them.

Independently of how the dispute between Kvanvig and Brogaard should be assessed (thus independently of whether knowledge and understanding are distinct or differ in value), what is relevant for the present purpose is that there is a third kind of knowledge, i.e. knowledge of a subject matter, besides propositional knowledge and knowledge-how. As with the other kinds, the question arises whether knowledge of a subject matter is *sui generis* or if it reduces to one or more other kinds of knowledge. If grasping explanatory and coherence relationships does not reduce to propositional knowledge, then this may be taken as evidence for knowledge of a subject matter being *sui generis*. However, those who take grasping such relationships to reduce to a set of abilities—e.g. the ability to manipulate information and use it in counterfactual thought (Grimm 2011)—may be inclined to favor a view according to which knowledge of a subject matter reduces to a combination of propositional knowledge and knowledge-how. Finally, those who aim at a unified account of the notion of knowledge, and are convinced by arguments, such as Stanley and Williamson’s, to the effect that knowledge-how is a kind of propositional knowledge, will seek to reduce knowledge of a subject matter to propositional knowledge.

To sum-up: besides knowledge by acquaintance, at least three other kinds of knowledge should be distinguished: propositional knowledge (which has dominated epistemological literature), knowledge-how, and knowledge of a subject matter. For each kind, at least two opposite stances are available: on one view, the relevant kind of knowledge can be reduced to one or more other kinds of knowledge, or it is a species of another kind of knowledge; on the alternative view, the relevant kind of knowledge cannot be so reduced, nor is it a species of any other kind of knowledge—it is *sui generis*.

3. Propositional knowledge does not exhaust knowledge by acquaintance

Suppose I describe my old friend Claudia to you in the finest detail. I tell you about how she looks physically (perhaps I also show you pictures of her) and the particular gestures she makes while she talks; I tell you how I met her, how we grew up together at school, how we fell out of touch for years and then met again like time had not passed; I tell you all the little and less-little troubles she went through and how she helped me to overcome mine; I tell

you about all the thoughts and emotions she confessed to me (to be sure, I would not do that in reality, but this is a thought experiment) and the thoughts and emotions I had in response. In fact, I tell you everything I know about Claudia. Do you know Claudia? Well, perhaps, you have heard so much about her that you almost feel like you know her. But would you be prone to claim that you do know Claudia? My guess is that you would not. In order for you to know *Claudia*, rather than just a lot of things *about Claudia*, you need to *meet* her. And indeed, even when we have a large amount of propositional information about someone, we get to know something more (arguably, something quite important) upon being perceptually acquainted with them: we know *that person*, rather than a mere bunch of *propositions* (or facts) about that person—or so it seems to me.

To be sure, the acquaintance relationship at work in this example is different from the one which constitutes the focus of the present study. As I said in §1, I am here focusing on a relationship one has with one's *phenomenal states*, i.e. *introspective* acquaintance. The example just provided, by contrast, features a *perceptual* relationship, i.e. a relationship with external, non-phenomenal objects. Whether perceptual relationships are genuine acquaintance relationships is a matter of debate. Recall, acquaintance (in the technical sense adopted here) is a relation which is both epistemically and metaphysically direct. Now, although not uncontroversial, that perceptual relations are direct in the epistemic sense is often accepted.²⁵ But that perceptual relations are direct in the metaphysical sense is much more contentious.²⁶ What is, then, the point of the above example? For one thing, if it turns out that perceptual encounter *does* involve perceptual acquaintance, then examples such as the one sketched above will constitute direct *prima facie* evidence for the claim that knowledge by acquaintance (at least in its perceptual form) is a *sui generis* kind of knowledge. If, however, perceptual encounter does *not* involve acquaintance, the example above can nonetheless help better understand in which sense knowledge by acquaintance *seems* to exceed propositional knowledge. At the very least, I hope, it will help the reader understand the reasoning underlying my preliminary motivation for thinking that propositional knowledge cannot

²⁵ One way to spell out the epistemic directness of perception is in terms of immediate justification: perceptual experience provides immediate justification for some of our beliefs about the perceived objects (Pryor 2005; Feldman 2003; Huemer 2001). As noted, this position is not uncontroversial but, arguably, it is less controversial than the thesis that perceptual experience is partly *constituted* by the perceived objects (i.e. that there is a metaphysically direct relation between the perceptual experience and the perceived object).

²⁶ Naïve realists seem to endorse a claim along these lines; however, naïve realism is far from being mainstream.

exhaust knowledge by acquaintance. Indeed, the kind of reasoning I will invite you to rehearse, in a few paragraphs, for the case of introspective acquaintance has a similar structure to the reasoning underlying the above example.

A second case I wish to draw your attention to is the following. Imagine you are going through a particularly tough period of your life, say, because of an ill-ended relationship. I am your friend, so you come to me to find comfort. As you tell me about all the vicissitudes and despair you are going through, I recognize much of the vicissitudes and despair *I* went through when *my* relationship ill ended. Back in time, I found myself in a very similar situation to yours and experienced very similar experiences to those I hear from you now. Trying to console you, I might begin by sincerely saying “I know how you feel...” and, since you know about my past, you take my words at face value and believe that I really know what you are experiencing.²⁷ But consider, instead, the case in which I have never gone through such a crushing situation. Another close friend of mine did, though, and I stood by him when that happened. He kept telling me about his vicissitudes and despair, and he did that in such detail, that I can now sincerely claim that I know (almost) everything about what he experienced throughout that tough period of his life. When you come to me to find comfort, and you tell me about your vicissitudes and despair, I may try to console you starting by saying “I know how you feel...”. I am sincere—I have heard so much about this kind of situation that I believe I know almost everything one can know about what it is like to go through it. You appreciate my support. But, deep down, you believe I cannot *really* know how you feel. Although I know a lot of facts about how people may feel in this sort of situation, there is something—arguably, something quite important—I miss until I myself undergo a similar experience—until I am myself *acquainted* with it.

As in the previous example, whether the relationship at work in this example is genuine acquaintance may be a matter of debate. For one thing, it seems that the notion of *experience* is used here in a looser sense than in the introspective acquaintance debate. Acquaintance theorists sometimes think of experiences in terms of *events*: they are instantiations of phenomenal properties (Gertler 2012). For instance, my visual experience of the blue sky is an event consisting in the instantiation of phenomenal blue.²⁸ Now,

²⁷ To be sure, on a rigorous account of the story, it is much more likely that I know something *similar*, rather than *identical*, to what you are experiencing, and that this is what you believe.

²⁸ Of course, other options are available for an account of the metaphysics of experience (see discussion in Chapter 3). Rehearsing other potential options again here is superfluous though. My point is simply that the individuation criteria for experience typically used by

although the end of your relationship is an event, it is not itself an experience because it does not consist in the instantiation of a *phenomenal* property. What in our example is referred to by your end-of-relationship ‘*experience*’ is, arguably, a much more complex phenomenon than the instantiation of a certain phenomenal property. It very likely involves the instantiation of *several* phenomenal properties at the same time and at *different* times, along with thoughts and desires with different contents (which, arguably, even if we assume that they have phenomenology, may not *reduce* to instantiation of cognitive-phenomenal properties), following and interlacing with one another in a complex mental *process* that takes place throughout a fairly *long span* of time.

However, again, this example may nonetheless help us understand the kind of reasoning underlying the intuition of acquaintance’s epistemic significance. Propositional knowledge is such that the information it conveys can be *transmitted* to others. What the example highlights is that the information my friend can transmit to me by expressing a bunch of propositions about how he feels is not sufficient for me to have all knowledge I need to truly claim that I know *what you are experiencing* (although it may be sufficient for me to truly claim that I know a lot of things *about* what you are experiencing). For me to truly claim *that*, I need to *have* the experience: I need to go through all the vicissitudes it implies firsthand.²⁹

My last example is a case of genuine introspective acquaintance. Imagine you are conversing with someone who has never had a pain sensation.³⁰ You decide to try and explain to them what it is like to feel pain. Such a hard task! How to describe the phenomenology of a pain sensation if not by saying that it is, well, *painful*? You gather up ideas to try and formulate a non-trivial description of a sample pain sensation. You may come up with something like “unpleasant sensation associated with the feel of an urge to

philosophers of mind imply that what is referred to by ‘experience’ in my example does not coincide with what most philosophers would call ‘experience’ in the technical sense.

²⁹ This probably involves being introspectively acquainted with a number of experiences. Perhaps this is partly what legitimates the broader use of ‘being acquainted with an experience’ in the example at issue. This is just a speculation, though, and analyzing the example further really exceeds my present purpose.

³⁰ The example may sound a bit unrealistic to you—it is unlikely that you can find an adult person who has never had a pain sensation—but consider that there exists a rare condition, called *congenital analgesia*, or *congenital insensitivity to pain*, in which the patient does not experience physical pain. You can, at any rate, design an analogous example which revolves around a different phenomenal state—a phenomenal state such that it is more likely to find an adult person who has never had it (e.g. orgasm).

complain, cry, scream, or otherwise express discomfort”.³¹ This, perhaps, may help your interlocutor to know *something* (some facts) about pain sensations (that they are unpleasant, that they make you feel like expressing discomfort). However, first, it is not clear that someone who has never felt pain can fully grasp the notions of unpleasantness and discomfort, which (especially the latter) are closely related to pain. Moreover, even admitting your interlocutor can grasp the notions of unpleasantness (they might have had unpleasant but not painful sensations) and feeling the urge to express discomfort (although this really sounds quite implausible, since discomfort is typically defined as a species of pain), your description is very far from exhausting what pain feels like. You may attempt alternative descriptions. My guess, however, is that you will never be able to find a non-trivial description which can convey the information your interlocutor would acquire if only they got acquainted with a pain sensation.

So, for any phenomenal property ϕ , which you now have, but I have never had in my life, no amount of propositional knowledge you may convey to me can fill the gap between my epistemic position and your epistemic position with respect to ϕ . You may describe ϕ to me in the finest detail, transmitting to me all the propositional knowledge you have about ϕ . You may tell me that ϕ is *P*, that ϕ is *Q*, that ϕ is *R*, and so on and so forth, and I may thereby come to know that ϕ is *P*, that ϕ is *Q*, that ϕ is *R*, and so on and so forth. Yet, although now we both know the same propositions about ϕ , there still is an important epistemic asymmetry between us, with respect to ϕ : you *know* ϕ , whereas I do not, and this partly in virtue of the fact that you have *experienced* ϕ , and thereby become acquainted with ϕ , whereas I have not.³² When I myself experience ϕ , I come to know something I did not know before, something on top of the propositional knowledge about ϕ I acquired from you: I come to know ϕ itself, on top of all the propositions about ϕ I already knew before experiencing ϕ .

The same model of reasoning which underlay the previous examples applies to the present case. The idea is that direct acquaintance with an experience provides one with knowledge of that experience, knowledge that cannot be acquired by merely learning an (even very large) number of

³¹ Nota bene: I am not referring here to the *disposition* to express discomfort (that would be off the point, since it is knowledge of pain *phenomenology* we are discussing, not knowledge of its functional role), but rather the *feeling* (i.e. the phenomenology) associated with one’s urge to express discomfort.

³² I say that the epistemic difference depends *partly* on your having experienced ϕ because experiencing ϕ , although necessary, may not be sufficient for knowing ϕ —something more (e.g. attending to ϕ) may be required.

propositions about that kind of experience. No amount of propositional knowledge about ϕ can change one's epistemic position with respect to knowledge by acquaintance of ϕ . This suggests that knowledge we get by being acquainted with an experience is *irreducible to* propositional knowledge: knowledge by acquaintance cannot be fully specified in terms of propositional knowledge—there is something to knowledge by acquaintance over and above propositional knowledge. The example also suggests that knowledge by acquaintance is *not* a *species* of the propositional knowledge kind either. As noted, there is no proposition you could transmit to me which could convey your knowledge by acquaintance of ϕ . This strongly suggests that knowledge by acquaintance is not a species of propositional knowledge. If those suggestions turn out correct, then there are good reasons for thinking that knowledge by acquaintance is a *sui generis* kind of knowledge.

4. Challenging intuitions

Let me consider some objections that may come to your mind at this point. First, you may think that I *could* indeed get a grasp of ϕ if you only provided me with the *right* pieces of propositional knowledge. More specifically, if you describe ϕ to me by highlighting the differences and similarities between ϕ and other experiences I have had, I may be able to narrow down on ϕ and thereby know it. For instance, you could tell me that ϕ is unpleasant, but different from pain; it has something in common with itch, but it is not quite itch—for one thing, it does not dispose you to scratch; it is somewhat similar to tickle, but it is not a tickle sensation either, and so on. If you tell me enough about how ϕ resembles and differs from other phenomenal properties, you may think, I will ultimately come to know ϕ itself even though I have not experienced it.³³ However, although—of course—there is a sense in which I know ϕ better than before upon being told what it resembles and what it differs from, such additional knowledge does not exhaust what I would know if I myself *had* the experience—if I knew it *by acquaintance*. Although knowledge about differences and similarities between ϕ and other phenomenal properties I have had may help me get a *better* grasp of ϕ , it will not help me get a *full* grasp of ϕ , or, at least, it will not help me get the right kind of grasp which is needed for me to know ϕ by acquaintance. Arguably, unless you tell me that ϕ is exactly like, say, itch (and I have already experienced itch), what I know is simply a fine-grained *description* of ϕ , rather than ϕ itself. Your telling me about the relevant differences and similarities

³³ The objection could be seen as a more general instance of Hume's *missing shade of blue* case.

can, at best, enable me to visualize or imagine some phenomenal properties which are similar to or different from φ and thereby extrapolate some information about φ 's phenomenology. But *merely* visualizing or imagining the phenomenal properties which are similar to or different from φ , plus extrapolation, is not sufficient for me to *know* φ *by acquaintance*. You may object that this *is* indeed sufficient for me to know φ by acquaintance, if by extrapolation I come to *have* the relevant experience—which, arguably, would involve having an imaginative phenomenal experience of φ . However, for one thing, it seems more plausible that coming to know about the differences and similarities between φ and other experiences one has been acquainted with enables one to *think* about φ , by way of a very fine-grained *description*, rather than *experience* φ itself. Moreover, even if knowing the differences and similarities enabled one to experience φ , it is clear that part of the new knowledge that one would acquire would be knowledge by acquaintance, i.e. knowledge one gets by being directly aware of φ .

Although you may be persuaded that there is something about φ that I cannot know by merely listening to your description of φ , you may still think that this does not entail that the relevant bit of knowledge I lack must be non-propositional. It may well be that there are some propositions about φ which I cannot know by testimony simply because you are not able to *express* them to me. Still, what I lack is *propositional* knowledge, rather than a different, irreducible kind of knowledge. If you *were* able to express the relevant proposition(s)—the objection goes, I would have full knowledge of φ . However, the point here is really not about proposition *expressibility*. For it may well be that the phenomenal property I have never had is *itch*, for which you do have a word, by which you can express what you know about it. Still, the word 'itch' will not help me know anything more about itch until I experience it firsthand.

Here is an alternative and perhaps clearer way to make the point.³⁴ Consider again your conversation with the person who has never felt physical pain. Part of the reason why this person cannot come to know what pain feels like (what it is like to *feel* pain, to have a pain sensation) simply upon being *told* by you is, arguably, that they lack the right *concepts*, possession of which would enable them to understand what you are talking about. More specifically, the pain-deprived person lacks the *phenomenal concepts* associated with pain experiences (phenomenal concepts such as PAIN, UNPLEASANTNESS, DISCOMFORT, and so on), possession of which would enable them to understand or grasp the propositions you are trying to convey.

³⁴ My reflection on this point benefitted from discussion with Luca Gasparri.

This consideration may help outline a subtler analysis of the pain-deprived person's epistemic situation. There *is* in fact, arguably, something about pain sensations that that person *can* come to know by simple testimony. They can come to know, for example, that pain sensations are typically provoked by bodily damage (burning, cutting, bumping, scratching, and so on). They can also come to know that pain sensations typically cause people to behave in a certain way (crying, screaming, groaning, moaning, complaining, making ugly face expressions, and so on). Arguably, the pain-deprived person can come to know all these facts independently of their *having* any pain experience because coming to know these facts only requires deployment of *non-phenomenal* concepts (particularly, it does not require the deployment of any phenomenal concept associated with pain phenomenology).

Accordingly, in the case of pain, we may distinguish the *phenomenal* concept PAIN_e, that is associated with the *phenomenology* of pain sensations, and the *non-phenomenal* (functional) concept PAIN_f, that is associated with pain sensations' functional role. The formation or acquisition of both kinds of concept requires *acquiring information* about what they are associated with: to form or acquire a certain concept one needs to acquire the information that constitutes its content. In the case of non-phenomenal concepts, one can acquire such information independently of one's being introspectively acquainted with any experience. The pain-deprived person, for instance, can form the non-phenomenal concept PAIN_f even if they have never had a pain sensation, simply by acquiring information about pain sensations' functional role. Such information can be acquired by testimony: once you have conveyed to them the appropriate set of propositions concerning the causes and effects of pain sensations, the pain-deprived person will be able to understand the non-phenomenal concept PAIN_f. Things are different when it comes to *phenomenal* concepts. For to acquire the information that constitutes the content of a *phenomenal* concept (and have such information available for one to come to *possess* the concept in the sense articulated in Chapter 2, that is, to have that concept available to be deployed in personal-level cognition and action-guidance), one needs to *have* the relevant experience and be introspectively *acquainted* with it. (This is closely related to what I argued for in Chapter 2: phenomenal-concept acquisition is grounded in introspection. More specifically, it is grounded in primitive introspection, that is, introspective acquaintance.) Accordingly, the pain-deprived person *cannot* possess (and thereby deploy) the phenomenal concept PAIN_e, unless they come to be introspectively acquainted with a pain experience, because being introspectively acquainted with a pain experience is necessary for them to

acquire the information which constitutes the content of such a phenomenal concept (that is, information about pain sensations' phenomenology).

At this point, you might perhaps have reasons to think that *introspective acquaintance* with pain experiences is somewhat responsible of the epistemic asymmetry between you and the pain-deprived person. If so, you may grant that, for any phenomenal state φ that you have had, and I have not, to fill the gap between my knowledge of φ and your knowledge of φ I need to be acquainted with φ . Still, you may object that the further piece of knowledge I acquire by being acquainted with φ (rather than by merely hearing a bunch of propositions about φ) does not constitute a distinct, *sui generis* kind of knowledge. Rather, it is just a new piece of good old propositional knowledge—that φ is *so-and-so*, where *so-and-so* is a property one can only attribute to φ if one is acquainted with φ . True, this is a somewhat *special* piece of propositional knowledge (one can only know that φ is *so-and-so* if one is acquainted with φ), but it is still propositional—and thereby not special in kind.

Those who want to defend this line of thought may want to spell it out in terms of what Chalmers (2003) calls 'direct phenomenal concepts'. Recall, direct phenomenal concepts are phenomenal concepts that are formed *upon* attending to the experience they are associated with, experience which (according to Chalmers) partly constitutes their content. Therefore, direct phenomenal concepts are not possessed *prior* the act of introspective acquaintance directed at the relevant experience but are deployed *during* the act of introspective acquaintance. The objector may argue that being introspectively acquainted with φ implies entertaining a propositional content involving a pure demonstrative (THIS) and a direct phenomenal concept (SO-AND-SO). The idea would thus be that knowledge by introspective acquaintance is special propositional knowledge whose content is the proposition 'this is so-and-so', where, to repeat, the concepts composing the relevant proposition are the pure phenomenal demonstrative THIS and the direct phenomenal concept SO-AND-SO, which is formed upon being acquainted with φ and is partly constituted by φ .³⁵ Accordingly, knowing φ by acquaintance requires possessing and deploying the direct phenomenal concept SO-AND-SO.

I am strongly inclined to resist this reasoning. In Chapter 2 I mentioned that, whereas I am open to the idea that primitive introspection requires the deployment of pure phenomenal demonstratives (where the deployment of a pure phenomenal demonstrative coincides with performing an act of attention directed toward the target phenomenal state), I tend to refuse the idea that

³⁵ Assuming that propositions are structured and made up of concepts.

primitive introspection depends on the deployment of direct phenomenal concepts. Here I explain why I think so. Direct phenomenal concepts, by definition, are phenomenal concepts that one can only have if one is acquainted with the experience such concepts are associated with. Accordingly, the direct phenomenal concept SO-AND-SO can only be possessed upon being introspectively acquainted with ϕ . Arguably, this is so because being introspectively acquainted with ϕ provides one with information about ϕ 's phenomenology, information that constitutes the content of the direct phenomenal concept SO-AND-SO. But if this is so, then it cannot be the case that *being introspectively acquainted* with ϕ itself requires possessing and deploying the direct phenomenal concept SO-AND-SO. To avoid vicious circularity, we must accept that being acquainted with an experience theoretically *precedes* the formation of *any* concept associated with that experience (including direct phenomenal concepts).³⁶ Such information acquisition about ϕ 's phenomenology is enabled by one's being introspectively acquainted with ϕ , that is, by one's being in a state of primitive introspection directed toward ϕ . Therefore, by being in a state of primitive introspection directed toward ϕ , one acquires information about ϕ 's phenomenology.³⁷ Such information is thereby available to form the content of the direct phenomenal concept associated with ϕ , that is, the phenomenal concept SO-AND-SO.

On the model I propose, such information acquisition is constituted by one's knowing ϕ by acquaintance. Knowledge by acquaintance is therefore *prior* to direct-phenomenal-concept-acquisition: it is by coming to know ϕ by acquaintance that one can acquire the information about ϕ 's phenomenology that is relevant to the formation of the direct phenomenal concept SO-AND-SO associated with ϕ . Quite obviously, then, knowledge by acquaintance cannot itself require the possession and deployment of direct phenomenal concepts. Therefore, knowledge by acquaintance cannot be propositional in form, not even in the special kind of way implied by the direct-phenomenal-concept view.

A final objection perhaps worth mentioning here is that, even if knowledge by acquaintance is not reducible to, or a species of, *propositional* knowledge, it may be reducible to, or a species of, *another* kind of knowledge:

³⁶ I say "theoretically precedes" because it might be that, in some cases, the introspective acquaintance and the formation of the direct phenomenal concept occur simultaneously or quasi-simultaneously (that is, it might be that there is no significant *temporal* priority of the introspective acquaintance over the concept formation).

³⁷ I will say more about the relationship between primitive introspection and information acquisition in Chapter 6, where I will articulate the idea that knowledge by acquaintance provides one with information which is non-propositional in form.

knowledge-how.³⁸ The idea would be that knowledge of ϕ by acquaintance reduces to one's ability to remember, recognize and imagine ϕ . However, first, a view along these lines, besides being minoritarian, would need to deal with arguments, such as Conee's (1994), to the effect that knowledge-how is neither sufficient nor necessary for knowledge by acquaintance. Here I briefly summarize Conee's argument. Consider Martha, who has extraordinary interpolation abilities: she has the power to visualize any shade of color she has not yet seen by interpolation between two shades she did see. Suppose that Martha has never experienced the shade *cherry red*, but she is told that cherry red is the shade between *burgundy red* and *fire engine red*, which she has experienced before. Thank to her extraordinary interpolation abilities, Martha can now visualize cherry red, and thereby come to know what it is like to see cherry red (she becomes acquainted with the relevant phenomenal property). However, by hypothesis, Martha knew already *how* to visualize cherry red before actually visualizing it—since she already possessed the *ability* to visualize by interpolation *any* shade of color she has not experienced yet. Therefore, Conee concludes, knowing how to visualize ϕ is *not sufficient* to know ϕ by acquaintance. As Conee acknowledges, the Martha example does not yet constitute a case against reducing knowledge by acquaintance to knowledge-how, because, as noted, the latter implies also the ability to *remember* and *recognize* the relevant experience (which Martha does *not* have until she visualizes cherry red). However, Conee argues, knowledge-how is *not necessary* for knowledge by acquaintance. Suppose that Mary, the black-and-white neuroscientist of Jackson's knowledge argument, has no visual imagination. As before, once Mary exits the black-and-white room, and sees a red object, she knows what it is like to see red—she knows phenomenal redness by acquaintance. However, since she does not have visual imagination, she is unable to remember, recognize and visualize phenomenal redness. Therefore, the ability to remember, recognize and visualize phenomenal red is not necessary to know phenomenal red by acquaintance. More generally, knowing how to remember, recognize and visualize a phenomenal property is only possible if one has *already* become acquainted

³⁸ I take the view that knowledge by acquaintance is reducible to knowledge of a subject matter (as characterized in §2) so implausible not to be worth of consideration. Incidentally, it is perhaps worth pointing out that my argument against the reducibility of knowledge by acquaintance to propositional knowledge will also be, *a fortiori*, an argument against the reducibility of knowledge by acquaintance to knowledge of a subject matter. For even if knowledge of a subject matter is irreducible to propositional knowledge (say, because on top of knowing a number of propositions it involves grasping the coherence and explanatory relationships they bear to each other), it surely does at least partly involve possession of some propositional knowledge.

with the relevant phenomenal property—if one already *knows* the relevant phenomenal property by acquaintance.

In addition to the abovementioned arguments, there is a further and perhaps more compelling reason to resist the idea that knowledge by acquaintance reduces to knowledge-how. Arguably, if knowledge-how is *sui generis* and analyzed in terms of abilities, it is *dispositional*: if one knows how to α , one is disposed to α when placed under the right conditions. Knowledge by acquaintance, by contrast, is *manifest*: being disposed to remember, recognize and imagine ϕ is not sufficient for one to know ϕ by acquaintance. For one may be disposed to remember, recognize and imagine ϕ even when one is not acquainted with ϕ . But to know ϕ by acquaintance, one *must* be acquainted with ϕ . Knowledge by acquaintance is thus manifest, rather than dispositional. Therefore, reducing knowledge by acquaintance to knowledge-how does not seem to be a promising move.³⁹

If my replies are sound, the examples analyzed in §3 show that there is *prima facie* motivation to think that knowledge by acquaintance is *sui generis*. This *prima facie* motivation may be defeated though. In the next section, I consider two objections against the claim that knowledge by acquaintance is *sui generis*, what I call the *objection from disunity* and the *objection from mysteriousness*, and suggest a way to rebut them.

5. Knowledge by acquaintance is *sui generis*: rebutting *disunity* and *mysteriousness*

5.1. The objection from disunity

You may think that there are independent theoretical reasons for believing that knowledge by acquaintance must be reducible to, or a kind of, propositional knowledge. For one thing, you may think that, other things being equal, we should favor a *unified* account of knowledge over an account according to which knowledge is a disjunctive notion that may refer to different kinds of thing. The objection, therefore, is that my view that knowledge by acquaintance is a *sui generis* kind of knowledge implies a *disunified* account of knowledge, one on which knowledge is unnecessarily treated as a disjunctive notion. The objection could also be seen as an appeal to *parsimony*: we should not multiply the number of kinds of thing (in this case, the kinds of knowledge) beyond necessity.

³⁹ Thanks to John Morrison for suggesting this reply to me.

First, I tried to show in the previous section that positing a distinct kind of knowledge—knowledge by acquaintance—besides propositional knowledge is *not* unnecessary or superfluous, for it constitutes the *best explanation* of some epistemic facts about our relationship with our own experiences. If knowledge by acquaintance were just a kind of propositional knowledge, it would be hard to explain why knowledge *of* a phenomenal property can only be acquired by being acquainted with that property.

Second, an epistemological view which is widespread enough to be considered the ‘received’ view entails that the notion of knowledge is *already* disunified (in the sense specified by the objection), that is, it is disunified independently of whether knowledge by acquaintance reduces to propositional knowledge or not. The relevant epistemological view is that *knowledge-how* is a *sui generis* kind of knowledge.⁴⁰ To be sure, as noted in §2, the received view about knowledge-how has been challenged, most notably by Stanley and Williamson (2001). However, Stanley and Williamson’s stance is heterodox and minoritarian. Moreover, it has been compellingly criticized. Arguably, the main problem with Stanley and Williamson’s view is that it is threatened by circularity (Koethe 2002). Recall, their view is that *S* knows how to α just in case *S* knows, of some way *w*, that *w* is for *S* a way to α . Yet, suppose you are watching me playing *Suzanne* on the guitar. You can demonstratively refer to the way I play the song and thereby come to know that *that* is a way for you to play *Suzanne*. However, you certainly do not thereby come to know how to play the song. To overcome this problem, Stanley and Williamson introduce what they call ‘practical modes of presentation’: *S* knows how to α just in case *S* knows, of some way *w*, that *w* is a way for *S* to α *and* *S* thinks of *w* under a practical mode of presentation, where “[t]hinking of a way under a practical mode of presentation [...] entails the possession of certain complex dispositions.” (Stanley and Williamson 2001: 429). Koethe (2002: 326-27), however, develops a compelling argument to the effect that we cannot explain the notion of *thinking of w under a practical mode of presentation* without either (i) engaging an infinite regress (thinking of *w* under a practical mode of presentation entails *knowing how* to instantiate *w*) or (ii) admitting that at least some instances of knowledge-how do not reduce to propositional knowledge. Stanley and Williamson’s intellectualism (i.e. the reduction of knowledge-how to propositional knowledge) has also been criticized, among others, by Rumfitt (2003) Cath (2011), Ren (2012) Brown (2013), Carter and Pritchard (2015a), Carter and

⁴⁰ Analogous considerations might be made about knowledge of a subject matter. I focus on knowledge-how here because the relevant philosophical debate has reached a much more advanced stage.

Pritchard (2015b), Carter and Pritchard (2015c), Brownstein and Michaelson (2016), and Carter and Navarro (2017).

So, there seem to be reasons to be skeptical about Stanley and Williamson's reduction of knowledge-how to propositional knowledge. Moreover, even if Stanley and Williamson could provide an explanation of practical modes of presentation which does not already presuppose the notion of knowledge-how, it is not clear that the resulting reductive account would be necessarily better off, in terms of unity, than the non-reductive account. For on the reductive account, knowledge-how and knowledge by acquaintance are *special* kinds of propositional knowledge: the former involves practical modes of presentation; the latter (as noted in §3 and §4) involves incommunicable propositions (propositions which can only be entertained by those who are acquainted with the relevant experience). On this view, knowledge-how and knowledge by acquaintance are two distinct kinds of propositional knowledge and, arguably, they are irreducible to one another. Knowing an experience by acquaintance has nothing to do with knowing ways under practical modes of presentation: I may well know ϕ by acquaintance without knowing that w is a way for me to experience ϕ —indeed, it sounds dubious to me that there even be *at all* a way for me to experience ϕ (at least in the sense of *way* which is relevant to Stanley and Williamson's account of knowledge-how). And knowing that w is a way for one to α has nothing to do with being acquainted with any experience: arguably, a zombie (who by definition does not have any experience and therefore cannot be acquainted with any experience either) may well know how to play *Suzanne* on the guitar. Furthermore, neither knowledge-how nor knowledge by acquaintance reduces to propositional knowledge *tout-court*, for each of them requires something more than justified true belief (plus anti-Gettier condition): the former requires that some of the belief's constituents be presented under a practical mode of presentation, the latter requires that some of the belief's constituents be experiences with which the subject is acquainted. The reductive view therefore implies a somewhat *disunified* account of propositional knowledge, on which propositional knowledge is a disjunctive notion which may refer either to propositional knowledge *tout-court*, or to knowledge-how, or to knowledge by acquaintance. Now, it is not clear why this view would be better off, in terms of unity, than the view according to which *knowledge* rather than *propositional* knowledge is disjunctive. The reductive view may save the unity

of knowledge, but only on pain of committing to the disunity of propositional knowledge.⁴¹

5.2. *The objection from mysteriousness*

The second objection I would like to address is the objection from *mysteriousness*. The idea is the following. As far as other putative *sui generis* kinds of knowledge are concerned, one or more possible analyses (with necessary and sufficient conditions), and thereby explanations in terms of more basic elements, are available. Propositional knowledge, for instance, may be analyzed in terms of justified true belief (plus an anti-Gettier condition).⁴² Knowledge-how, as we saw, may be analyzed in terms of abilities.⁴³ Knowledge of a subject matter may be perhaps analyzed in terms of propositional knowledge plus grasp of dependence, explanatory, and coherence relationships. Knowledge by acquaintance, by contrast—the objection goes, if it is not reducible, nor a species of, any other kind of knowledge, does not seem to be analyzable. Therefore, it cannot be explained in terms of more basic elements. Its primitivity makes it mysterious and this renders suspicious its existence as a *sui generis* kind of knowledge, since all other *sui generis* kinds of knowledge *can* be analyzed in terms of something else.

However, knowledge by acquaintance *can* be analyzed. One knows a phenomenal state φ by acquaintance iff one bears a relation of acquaintance, i.e. *direct awareness*, to φ (where the awareness is both epistemically and metaphysically direct in the sense described in §3). Being acquainted with φ is thus both necessary and sufficient for one to know φ by acquaintance. As mentioned in Chapter 4, considerable efforts have been made to explain what direct awareness is and to specify the metaphysical structure of the mental state the subject is in by virtue of being acquainted with a phenomenal state. Typically, the idea is that, in virtue of being acquainted with a phenomenal state φ , the subject is in an introspective state ι which is partly constituted by φ . As noted, a variety of ways have been explored to explain what the

⁴¹ Moreover, the non-reductive view (according to which knowledge-how and knowledge by acquaintance are *sui generis* kinds of knowledge) allows for there being a notion of *knowledge* such that, although irreducible to one another, propositional knowledge, knowledge-how and knowledge by acquaintance are distinct species of one and the same kind—i.e. *knowledge*.

⁴² I am clearly oversimplifying: the literature on the analysis of propositional knowledge is so huge that, obviously, it makes no sense to summarize it here. The important point is simply that multiple analyses of propositional knowledge are available.

⁴³ Here too I omit potential specifications and alternative analyses.

relevant constitutive relationship involves. Brie Gertler (2001), for instance, proposes the following account of introspective acquaintance: one is introspectively acquainted with one's current phenomenal state token φ iff one has an occurrent mental token ι which is such that: (i) φ is embedded in ι ; (ii) ι refers to φ ; and (iii) (ii) is true partly in virtue of (i).⁴⁴ On Gertler's view, then, the relationship of acquaintance is analyzed in terms of the more basic notions of *embedding* and *reference*.

Katalin Balog's (2012) *quotational* account suggests an alternative analysis of the acquaintance relationship. Balog argues that the constitutive relationship between the introspective state and the introspected state is reflected at the neural level: when ι is constituted by φ , the neural states that realize ι are the same neural states that realize φ . By applying Balog's account to the notion of introspective acquaintance at issue here (she focuses on the constitutive relationship of acquaintance holds between an experience and a direct phenomenal concept), we have that one is introspectively acquainted with one's current phenomenal state token φ iff one is in a phenomenal state ι such that (i) the neural states that realize ι are the same neural states that realize φ ; (ii) ι refers to φ ; and (iii) (ii) is true partly in virtue of (i).⁴⁵ The acquaintance relationship is here analyzed in terms of *identity of neural realizers* and *reference*.

In Chapter 4 I put forward my own analysis of the constitutive relationship of acquaintance. Recall, on the account I propose, one is introspectively acquainted with a phenomenal state token φ iff one is in a phenomenal state ι such that φ is integrated in ι , where a phenomenal state φ is integrated in a phenomenal state ι iff:

- (i) φ is a proper part of ι ;
- (ii) ι inherits all of φ 's essential properties;
- (iii) (some of) φ 's accidental phenomenal depend on ι ;
- (iv) ι refers to φ .

Here too the notion of constitutive relationship of acquaintance is analyzed in terms of more basic notions: *proper part*, *essential vs. accidental phenomenal properties*, *metaphysical dependence*, and *reference*.

⁴⁴ Where, recall, a mental state φ is embedded in a mental state ι iff (a) ι cannot be present if φ is not also present and (b) φ can be present even if ι is not present.

⁴⁵ I am admittedly stretching a bit Balog's view. However, what I outline in the main text is a legitimate potential analysis of acquaintance—an analysis in terms of identity of neural realizers and reference. This is all that matters for my present purpose.

Thus, analyses of the acquaintance relationship (i.e. analyses with necessary and sufficient conditions and in terms of more basic notions) *are* available—and new ones may be developed. Since knowledge by acquaintance is analyzed in terms of the relationship of acquaintance (S knows ϕ iff S is acquainted with ϕ), we thereby have several options for an analysis of knowledge by acquaintance. Therefore, knowledge by acquaintance should not be considered as more mysterious than other kinds of knowledge in this respect.

6. Why *sui generis* knowledge by acquaintance matters

A last note about what *motivates* my attempt to defend the thesis that knowledge by acquaintance is a *sui generis* kind of knowledge. First of all, as I pointed out, the idea that knowledge by acquaintance is a *sui generis* kind of knowledge is discarded as patently false in most current philosophical debates. Swimming against this alleged orthodoxy's tide, I believe that the Russellian thesis deserves deeper philosophical attention and should be reintroduced in the epistemological scene.

Contemplating the idea that knowledge by acquaintance could be considered as a full-fledged kind of knowledge may help shed light on the concept of knowledge itself. Most of the analytical epistemology debate about the nature of knowledge has focused on reductive analyses of *propositional* knowledge, that is, on finding out propositional knowledge's necessary and sufficient conditions.⁴⁶ Although exploring alternative analyses of propositional knowledge is undoubtedly valuable and useful to get a better grasp of what knowledge may involve, it is not obvious that a reductive analysis of propositional knowledge will ultimately tell us what *knowledge is*. For it may be that not all knowledge is *propositional* knowledge. If so, a reductive analysis of propositional knowledge will not exhaust the more general concept of knowledge. So, if knowledge by acquaintance turns out to

⁴⁶ For instance, knowledge has traditionally been defined in terms of *justified true belief* (S knows that p iff S believes that p , p is true, and S is justified in believing that p). After Edmund Gettier's (1963) famous challenge to this traditional analysis (he compellingly argued that justified true belief is not sufficient for knowledge), alternatives have been explored. Here are some examples:

- S knows that p iff S believes that p , p is true, S is justified in believing that p , and all S 's grounds for believing p are true (Clark 1963);
- S knows that p iff S believes that p , p is true, S is justified in believing that p , and S 's justification for believing that p is undefeated (Lehrer and Paxson 1969);
- S knows that p iff S believes that p , p is true, and the fact that p is causally connected with S 's belief that p in an appropriate way (Goldman 1967).

be full-fledged knowledge, then the concept of knowledge cannot be defined in terms of a specific analysis of propositional knowledge: a more general and comprehensive definition must be provided.

Besides its 'instrumental' (if you will) interest for the broader epistemological debate about the concept of knowledge, knowledge by acquaintance is also of epistemological interest in itself. For example, the following issues would deserve to be explored: is knowledge by acquaintance analyzable? (Compare: propositional knowledge can be analyzed, say, in terms of justified true belief.) Does the relationship of acquaintance always constitute knowledge, or must further conditions apply? (Compare: having a justified true belief may not be sufficient for having knowledge.) Does knowledge by acquaintance have special epistemic properties? (Compare: some propositional introspective knowledge has sometimes been claimed to be infallible or incorrigible.)

Finally, reintroducing the Russellian notion of knowledge by acquaintance in the current epistemological debate might be useful to better understand a variety of views which appeal to the acquaintance relationship (e.g. acquaintance with external objects in perception, acquaintance with abstract objects in intuition, acquaintance with moral values, etc.) and analyze their epistemological consequences.

Conclusion

In this chapter I have defended the thesis that knowledge by introspective acquaintance is a *sui generis* kind of knowledge.

First, by appealing to the reader's intuitions, I have tried to show that there is *prima facie* reason to believe that knowledge by acquaintance is a *sui generis* kind of knowledge: no amount of propositional knowledge can fill the gap between the epistemic position of a person who has been acquainted with a certain phenomenal state ϕ and that of a person who has never been acquainted with ϕ .

Second, I have considered two objections against the thesis that knowledge by acquaintance is a *sui generis* kind of knowledge, what I have called the *objection from disunity* and the *objection from mysteriousness*. I hope to have shown that both objections can be resisted. If so, our *prima facie* reason to think that knowledge by acquaintance is *sui generis* remains undefeated.

If the argument developed in this chapter is sound, there are good reasons to believe that, if knowledge by acquaintance is genuine knowledge, it is *sui generis*. In the next chapter I am going to motivate the antecedent of this conditional. I will argue that knowledge by acquaintance, as conceived here (i.e. as a kind of knowledge irreducible to propositional knowledge) is full-fledged knowledge, or, at least, that it is *epistemically significant*.

CHAPTER 6

KNOWLEDGE BY ACQUAINTANCE IS EPISTEMICALLY SIGNIFICANT

In the previous chapter I tried to show that knowledge by acquaintance is *sui generis* (particularly, that it is not reducible to propositional knowledge). In this chapter, I argue that, even though it is not reducible to propositional knowledge, knowledge by acquaintance is nonetheless epistemically significant.

First, I clarify what it takes for a phenomenon to be epistemically significant (§1). Rather than providing a set of necessary and sufficient conditions, I consider a number of phenomena which are typically considered as epistemically significant and try to individuate the *marks* of epistemic significance, that is, features possession of which provides us with good reasons for considering something as epistemically significant. I propose that there are at least two main marks of epistemic significance: information acquisition and epistemic evaluability. The exploration of these two notions occupies the first and the second half of the rest of the chapter respectively. In §2 I provide a characterization of information, in accordance with current literature in epistemology and philosophy of information. In §3 I highlight the relationship between propositional knowledge and the acquisition of information and I suggest that, in general, epistemic significance tends to correlate with information acquisition. In §4 I argue that knowledge by acquaintance involves information acquisition. In this section, I also make a digression about what I take to be a special epistemic property of knowledge by introspective acquaintance, what I call *complete and perfect grasp*—an epistemic property that, in a sense, is analogous to an epistemic property that has traditionally been attributed to some introspective propositional knowledge, namely, *infallibility*. The interim conclusion of §4 is that, regardless of whether it does possess such a special epistemic property, knowledge by acquaintance displays the first mark of epistemic significance (i.e. information acquisition). In §5 I introduce the notion of epistemic evaluability. In §6 I explore the relationship between epistemic evaluability and epistemic significance. I suggest that either epistemic evaluability or epistemic value is a mark of epistemic significance. In §7 I argue that primitive introspection is both epistemically evaluable and epistemically valuable. Accordingly, knowledge by acquaintance, which is constituted by the state of

primitive introspection, displays the second mark of epistemic significance, i.e. epistemic evaluability. I conclude that, knowledge by acquaintance displaying both marks of epistemic significance, there are good reasons for thinking that knowledge by acquaintance is epistemic significant.

*

1. Two marks of epistemic significance

In chapter 5 I have defended the thesis that knowledge by acquaintance is a *sui generis* kind of *knowledge*. Some readers, however, may reject this claim *a priori*: they might think that knowledge by acquaintance cannot be full-fledged *knowledge* because knowledge is by definition propositional. This issue is likely terminological. Regardless of whether knowledge by acquaintance deserves to be called *knowledge*, there still is, arguably, a substantial disagreement between those who, like me, endorse the Russellian thesis and argue that knowledge by acquaintance is a *sui generis* kind of knowledge and those who reject it. Those who resist the Russellian thesis may do so *not* for merely terminological reasons, but because they think that a phenomenon like what Russell called ‘knowledge by acquaintance’ cannot have any epistemic significance: it does not constitute any epistemic achievement. Since knowledge is epistemically significant (if not knowledge, what else could be?), what Russell called ‘knowledge by acquaintance’ is not knowledge.

The *essence* of the Russellian idea, I suggest, revolves around knowledge by acquaintance’s *epistemic significance*. Accordingly, the main aim of those who want to defend the Russellian idea should be to show that knowledge by acquaintance is epistemically significant: it has a relevance for the subject which is eminently *epistemic*. This is the purpose of the present chapter. In this chapter, I defend the thesis that knowledge by acquaintance is epistemically significant. This framing of the thesis conveys the gist of the claim I want to make but does not meddle in the terminological controversy by committing to the idea that something non-propositional can be legitimately called ‘knowledge’.

Roughly, my argument proceeds by highlighting fundamental similarities between knowledge by acquaintance and other epistemically significant phenomena. I will try to show that knowledge by acquaintance shares with other epistemically significant phenomena at least two features that make them epistemically significant. To be sure, it is not my intention to

provide a set of necessary and sufficient conditions for something to be epistemically significant.¹ More modestly, I will, first, draw a list of epistemically significant phenomena (§1.1) and, second, isolate two features that seem to make those phenomena epistemically significant (§1.2). Such features are what I call ‘marks’ of epistemic significance: they are features such that, if a phenomenon displays them, and especially if it displays the conjunction of them, we would be naturally inclined to classify that phenomenon as epistemically significant. To be sure, there may be other marks of epistemic significance besides those I point at. However, display of the two marks highlighted in this section is already sufficient to give us good reasons for attributing epistemic significance to a phenomenon, or so I argue.

1.1. Some epistemically significant phenomena

If there is any epistemically significant phenomenon at all, that is indubitably (propositional) *knowledge*. Knowing that there is a cup before me, knowing that two plus two equals four, knowing that French Revolution occurred in 1789, and knowing that I am hungry are all epistemically significant states. Indeed, as noted in Chapter 5, most literature in epistemology has focused on (propositional) knowledge (its definition, possibility, analysis, value, and so on). This, arguably, suggests that (propositional) knowledge should be considered as the epistemically significant phenomenon par excellence.

Another phenomenon which has been discussed quite extensively in recent epistemological literature and is typically considered as epistemically significant is *understanding*. Like many philosophical notions, there is no universally accepted definition of the notion of understanding. One option may be to say that understanding involves having a number of beliefs, or a number of pieces of propositional knowledge, and grasping the dependence, explanatory, and coherence relationships between them (*cf.* Kvanvig 2003). As mentioned in Chapter 5, at least three species of understanding have been distinguished: *propositional understanding* or understanding-that (I understand that nothing can travel faster than light), *understanding-why* (I understand why continuous economic growth is unsustainable), and *objectual understanding* (I understand psychology). Independently of how understanding in general, and different species of understanding in

¹ Incidentally, I am not even sure that such a set can be provided at all. It might be that the borders of the notion of epistemic significance are a bit too blurred to allow for a necessary-and-sufficient-condition definition immune from counterexamples.

particular, are to be characterized, that all of them are epistemically significant could be fairly considered as uncontroversial.

Another phenomenon that is by and large considered as epistemically significant is *knowledge of a subject matter* (knowing physics, knowing economics, knowing psychology, and so on). As noted in Chapter 5, some might want to reduce the notion of knowledge of a subject matter to that of propositional knowledge; others might want to reduce it to the notion of objectual understanding. Regardless of how this issue is settled, what seems to be largely agreed upon is that knowledge of a subject matter is an epistemically significant phenomenon.

So, we may quite safely assume that, at the very least, propositional knowledge, understanding, and knowledge of a subject matter are epistemically significant phenomena. Besides these, there are other phenomena that some would consider as epistemically significant but whose epistemic significance may perhaps be more controversial than that of the three phenomena listed above.

Consider knowing-how. Those who think that knowledge-how is full-fledged knowledge will probably also think that knowledge-how is epistemically significant. This will be particularly straightforward for those who reduce knowledge-how to propositional knowledge. Even among those who think that knowledge-how is *not* full-fledged knowledge (say, because they believe that only propositional knowledge is full-fledged knowledge and knowledge-how does not reduce to propositional knowledge), some may nonetheless take it to be epistemically significant (perhaps on the basis of the fact that coming to know how to do something entails some kind of *learning*).

On the assumption that knowledge can be analyzed in terms of more basic components (say, as Gettier-proof justified true belief), some may think that not only knowledge, but also some proper subset of the components of knowledge is epistemically significant too. *Justified belief*, for instance, may be considered as epistemically significant, on the basis that, like knowledge, justified belief conforms to some epistemic norm(s) (e.g. rationality). One might have innumerable justified *false* beliefs, though. Some might take the possibility of justified false beliefs to undermine the plausibility of justified belief's being epistemically significant. Arguably, the background assumption would be that factivity is required for epistemic significance. What about having a *true belief*? That is certainly factive. Some may take true belief to be epistemically significant, based on the fact that true belief, like knowledge, involves the subject's having an accurate representation of reality. However, true beliefs can be formed by pure luck. It may be argued, for example, that

for a true belief to be legitimately considered as epistemically significant, it must be the result of the exercise of an epistemic virtue such as rationality, rather than the outcome of a mere stroke of luck. Finally, what about mere *belief*? Many would probably reject the idea that mere belief is epistemically significant in itself (independently of being true or justified), on the basis that mere belief does not display any epistemically relevant feature (although there might be a way to resist this claim, which I am going to consider—but not endorse—in §1.2): what is the point, from an epistemological perspective, of having a however huge amount of false and unjustified beliefs?

Finally, there is a sense in which we could think about *perceptual states* and *introspective states* as epistemically significant. On the assumption that perceptual experience is representational, by having a perceptual experience the subject represents worldly objects and their properties. Accordingly, perceptual experience has accuracy conditions. If the perceptual representation is accurate, it tells us something about our environment. Indeed, perceptual experience has traditionally been considered as an important epistemic ground upon which the complex of our knowledge is built. In this framework, it is natural to think of perceptual experience as epistemically significant. For similar reasons, introspection may be listed among the bearers of epistemic significance. Classical (hard-core Cartesian) foundationalists may want to develop an epistemic system such that the totality of our knowledge relies upon our introspective judgments. Accordingly, they would certainly consider introspection as epistemically significant. To be sure, less extremist epistemologists may think of introspection as epistemically significant too—they might consider it as *one* of the grounds in the epistemic system. Some might have the intuition that only what I call *reflective introspection* can be legitimately considered as epistemically significant, some might be open to the idea that also *primitive introspection* is epistemically significant. (The purpose of the present chapter is precisely to motivate the latter idea.)

So, here is an interim list of potential epistemically significant phenomena (the list is not supposed to be exhaustive): propositional knowledge, understanding-that, understanding-why, objectual understanding, knowledge of a subject matter, knowledge-how, justified belief, true belief, perception, introspection. Arguably, while considering the phenomena just listed, we have the intuition that at least *some* of them (if not *most* of them, or *all* of them) are similar enough, in a relevant respect, to be considered as members of the same ‘family’: they seem to share a number of features and this seems to indicate that they are species of one and the same kind of phenomenon.

Knowledge by acquaintance, I maintain, should appear in the above list too. When you primitively introspect a phenomenal state ϕ (be it a pain sensation, a reddish experience, a wine-taste experience, or what have you) you are in a state which is epistemically relevant (or so my intuition goes). Those who share this intuition will agree that knowledge by acquaintance (and, *a fortiori*, the state of primitive introspection which constitutes such knowledge) should be numbered among the epistemically significant phenomena.²

In the next subsection, I will try and investigate what the features that are shared by the phenomena listed above could be.

1.2. Marks of epistemic significance

A mark of F , as I understand it here, is a feature such that a phenomenon's possessing it gives us reasons for thinking that that phenomenon is F . Arguably, a mark of F is a feature that is shared by at least some F s, and that could be reasonably considered as part of *what makes* an F an F . Accordingly, a mark of epistemic significance is a feature that is shared by at least some epistemically significant phenomena and that could be reasonably considered as part of what makes epistemically significant phenomena epistemically significant. So, a first pass at the individuation of the marks of epistemic significance consists in trying and isolate the features that epistemically significant phenomena share.

What are the features that at least some of the phenomena listed in §1.1 (at the very least, propositional knowledge, understanding in its various forms, and knowledge of a subject matter) have in common? It might be suggested that *being the topic of a sustained debate in the epistemological literature* could be considered as a symptom of something's being epistemically significant. After all, epistemology is the study of epistemic phenomena, so it is plausible to think that, if something is the topic of a debate in epistemology, there are good reasons to think that it has some epistemic significance. Indeed, many of the notions listed at the end of §1.1 have been the object of sustained epistemological debate. Propositional knowledge has always been the lead in the debate, but knowing-how,

² Throughout the rest of the chapter, I will switch freely between talking about knowledge by acquaintance being epistemically significant and primitive introspection (or the primitive introspective state) being epistemically significant. This because, on my view, the two entail each other—knowledge by acquaintance is constituted by the state of primitive introspection.

knowledge by acquaintance, and, more recently, understanding, have all received a substantial amount of attention.

Being the object of epistemological debate may be a *prima facie* symptom of epistemic significance. However, *ultima facie*, it is probably *not* to be considered as a *mark* of epistemic significance. For one thing, some phenomena that most of us would intuitively classify as epistemically significant have been virtually ignored by epistemological literature until recently (I am thinking especially about *understanding*). Moreover, it may be argued that it is something's being epistemically significant that makes it the object of epistemic debate—rather than the other way around. So, while being the object of epistemological debate may help us identify epistemically significant phenomena, it is not adequate to constitute a *mark* of epistemic significance.

A different proposal may be the following. A feature that at least some of the phenomena listed in §1.1 seem to share is that they have an effect on the subject's epistemic position, or attitude, with respect to something (a proposition, an event, an object, a fact, a subject matter, and so on), where the notion of epistemic attitude is interpreted, roughly, in terms of credence or subjective probability. One changes epistemic attitude with respect to x just in case one's credence in a proposition about x changes. Accordingly, one may hypothesize that one mark of epistemic significance is *change in one's epistemic attitude*. This feature, however, is possessed not only by phenomena listed in §1.1, but also by other phenomena. Coming to *doubt* that p , for instance, implies changing one's epistemic attitude with respect to p : it implies decreasing one's credence in p . Coming to *suspect* that p is another example: when one comes to suspect that p , one increases a bit one's credence in the (previously zero-credence assigned) proposition that p . Mere *belief* also entails a change in one's epistemic attitude: when one comes to believe that p , one's credence in p increases and one's epistemic attitude towards p thereby changes.

If one thinks that doubt, suspicion, and belief are *not* epistemically significant (an issue, to be sure, which I am not going to take a stance on here) change in epistemic attitude, although possessed by some epistemically significant phenomena, may not be that which partly *makes* epistemically significant phenomena epistemically significant. If so, mere change in one's epistemic attitude is not a mark of epistemic significance. Since I do not want my account to *entail* that controversial cases such as doubt, suspicion, and belief are cases of epistemically significant phenomena, I will leave change in epistemic attitude aside here and look for other marks of epistemic significance.

It may be plausible to think that the mark of epistemic significance is *factivity*. From this angle, epistemic attitudes such as doubt or belief are *not* to be considered as epistemically significant. Most of the candidates listed in §1.1, however, are. Knowledge, of course, is factive, as well as true belief. Understanding is usually taken to be factive (although there are exceptions: Gordon (2018) cites Baker (2003) and Zagzebski (2001) as representatives of the idea that factivity is not necessary for understanding). Knowledge of a subject matter, as far as it partly involves propositional knowledge, has a factive dimension. Veridical perception and introspection are factive too. Mere justified belief, though, is *not* epistemically significant if the latter *requires* factivity (that is, if factivity is the *sole* mark of epistemic significance). Those who have the strong intuition that mere justified belief *is* epistemically significant will have to reject the idea that factivity is essential to epistemic significance. They may accept that factivity is one potential mark of epistemic significance, but argue that there are *other* marks, display of which may be sufficient for something to be legitimately considered as epistemically significant.

Things get a bit trickier when it comes to knowledge-how, if we assume that knowledge-how does not reduce to propositional knowledge. For factivity, as typically understood, entails truth, but truth is a feature that only propositions (or propositional states, or attitudes) can display. Now, one may have a more general notion of factivity, one that does not imply truth and thereby allows for knowledge-how to be factive (perhaps one revolving around the idea that, for *S* to know how to α , *S*'s attempts to α must be *successful*). Alternatively, one may give up the idea that knowledge-how is epistemically significant (given that it does not display factivity as typically understood). Finally, one may want to do justice to the intuition that knowledge-how *is* epistemically significant, while avoiding an *ad hoc* characterization of factivity. The latter option pushes one to explore further potential marks of epistemic significance.

Before moving on to another potential mark, let me elaborate a little bit further on the idea of factivity. As noted, true belief is a factive epistemic attitude. However, some may have the intuition that mere true belief should not be considered as epistemically significant, given that a belief could be true by a pure stroke of luck. For instance, I may have the true belief that it is raining outside, but that belief be such that its formation has nothing to do with its raining outside: it is not based on looking out of the window or on listening to the weather report but is, say, caused by my taking a rain-belief-inducing drug. In such a case, my belief is true, but by mere luck: I might

well have had the same belief in a beautiful sunny day. Therefore, the intuition goes, my belief should not be considered as epistemically significant.

One way to deal with this intuition is to supplement the idea of factivity with a stronger connection between the epistemic attitude and that which the epistemic attitude is directed to (say, between the true belief and what is truly believed)—a connection such that the epistemic attitude is factive, but not by mere luck. On my view, such a connection is best spelled out in terms of *information transmission*. Accordingly, a phenomenon, process, or state is epistemically significant if it puts one in the position to acquire information about that which the relevant phenomenon (process, or state) is directed to, where *acquiring* information about x implies *receiving* information which is *generated* by x . I will say more about information and information acquisition in §2. For now, it is important to point out that information acquisition, while entailing factivity (one can only acquire information that p if p is true: there is no false or mistaken information), it implies a stronger connection between the epistemic phenomenon (process, or state) and its target: it implies that the information received is generated by the target of the epistemic phenomenon (process, or state). From this angle, the true rain-belief above is not epistemically significant because it does not involve any information acquisition: my belief that it is raining outside is not formed upon receiving the information generated by the event of raining. If, by contrast, I form the true belief that it is raining outside by *seeing* the rain, then there is transmission of information from the raining event to my belief. My true belief is, in this case, epistemically significant.

To be sure, I do not claim to have thereby established that information acquisition is required for epistemic significance. To repeat what I have already expressed at the beginning of this section, it is not my intention to establish a set of necessary or sufficient conditions for epistemic significance. My aim here is to explore potential marks of epistemic significance—features display of which could be reasonably seen as a symptom of the presence of an epistemically significant phenomenon. True, I am ultimately going to propose that information acquisition, rather than mere factivity, is a mark of epistemic significance. I do not, however, thereby claim to have refuted the idea that mere factivity can be a mark of epistemic significance in itself: some people—who do not feel the pull of the intuition that lucky true beliefs are not epistemically significant—may have good reasons to think so.

Besides information acquisition, another feature that the phenomena listed in §1.1 seem to share is *epistemic evaluability*. The idea is, roughly, that we can individuate some epistemic value(s) setting a standard of evaluation for those phenomena. Intuitively, phenomena such as propositional

knowledge, true belief, knowledge of a subject matter, and understanding can be evaluated from an epistemic point of view. For instance, philosophers have wondered what (if any) is the extra value of knowledge with respect to mere true belief, given that, arguably, truly believing that p and knowing that p have the same consequences on the subject's actions.³ More recently, some philosophers have argued that understanding is more valuable than knowledge (Kvanvig 2003). It thus seems that epistemically significant phenomena can be evaluated. As we will see in §5, the value(s) attributed to these phenomena is (are) eminently epistemic.

In what follows, I will assume that information acquisition and epistemic evaluability are marks of epistemic significance: if a phenomenon displays them, then there are good reasons for thinking that that phenomenon is epistemically significant. It may be that displaying only one of those marks (either information acquisition or epistemic evaluability) is enough for something to be considered as epistemically significant. Here, however, I am going to make the stronger assumption that, for us to have good reasons to think that a phenomenon is epistemically significant, such phenomenon should display *both* marks.

In the remaining of this chapter, I will argue that there are good reasons for thinking that knowledge by acquaintance is epistemically significant. My argument runs as follows:

- P1. If a phenomenon displays (i) information acquisition and (ii) epistemic evaluability, then there are good reasons for thinking that such a phenomenon is epistemically significant.
- P2. Knowledge by acquaintance displays both (i) and (ii).
- C. There are good reasons for thinking that knowledge by acquaintance is epistemically significant.

§§2-4 are devoted to showing that information acquisition is a mark of epistemic significance and that knowledge by acquaintance involves information acquisition. §§5-7 are devoted to showing that epistemic evaluability is a mark of epistemic significance and that knowledge by acquaintance is epistemically evaluable.

Throughout my discussion of knowledge by acquaintance and information acquisition in §4, I will also point at what I take to be a special epistemic property of knowledge by introspective acquaintance, what I call

³ This is usually called the *Meno problem*, since the question why knowledge is more valuable than true belief is asked by Socrates in Plato's *Meno*.

complete and perfect grasp—an epistemic property that, in a sense, is analogous to an epistemic property that has traditionally been attributed to some introspective propositional knowledge, namely, *infallibility*.

2. Information acquisition

Although consensus as to a unified notion of information is far from having been reached, the characterization articulated in Claude Shannon's (1948) communication theory has turned out particularly useful for philosophical thinking, at least since Fred Dretske (1981) exploited it in elaborating his theory of knowledge. Accordingly, information is characterized in terms of *reduction of uncertainty*, or *exclusion of possibilities*. To be sure, this does *not* necessarily provide a *definition* of information: it does not purport to tell us what information is. However, it enables us to quantify information and calculate the *amount* of information which is associated with a given event or state of affairs. It also enables us to evaluate whether information about a state of affairs is transmitted to a receiver and, if so, how much information gets transmitted.⁴

The amount of information associated with a given event or state of affairs is a function of the number of alternatives that the occurring of such event or state of affairs contributes to rule out: the more possibilities are excluded, the more information. Take two boxes: box A containing eight numbered balls, box B containing sixteen numbered balls. The number of possibilities excluded by your drawing a three from box A is seven. The number of possibilities excluded by your drawing a three from box B is fifteen. Therefore, the amount of information associated with your drawing a three from box B is greater than the amount of information associated with your drawing a three from box A.⁵

When we consider information *transmission*, we must distinguish the *source* (*s*), i.e. what *generates* the information, and the *destination* or *receiver* (*r*), i.e. that which *receives* the information. Both the source and the destination have a certain amount of information associated with them. The drawing of a three from box A is a source of information—it generates a certain amount of information: it excludes the possibilities that the drawn number is

⁴ The present section's discussion heavily draws on Dretske (1981).

⁵ Here I am focusing on possibilities that are equally likely. Things are more complex, though, because the amount of information associated with a certain event is also a function of the *likelihood* of that event: the more likely the event, the less information associated with it. Although this is very important in communication theory and its application in epistemology, it is not crucial for our present purpose.

a one, a two, a four, a five, a six, a seven, and an eight. When you come to know that a three was drawn, you are a receiver of information: you receive the information that a three was drawn. You can thereby rule out all the other possibilities (i.e. that the drawn number is a one, a two, a four, a five, a six, a seven, and an eight). More generally:

You acquire information about a state of affairs s iff you can rule out the possible worlds in which $\neg s$.⁶

When you come to know that an odd number was drawn from box A, you still acquire information. However, the amount of information you receive in this case is smaller, for the number of possibilities you can exclude is lower: you can only rule out four possibilities, i.e. that the drawn number is a two, a four, a six, and an eight.

Thus, there may be information which is generated at the source but does not get to the receiver. In the case just discussed, the information that the drawn number is odd does get to the receiver, but the information that the drawn number is a three does not. Only part of the information generated by the event of drawing a three from box A is acquired by the receiver in this case. It may also be the case that there is more information associated with the receiver than is transmitted by the source. Imagine that the way you come to know that an odd number was drawn is by receiving a sheet of paper with “odd” written in red on it. To the event of your reading the red “odd” is associated not only the information that an odd number was drawn, but also that “odd” is written in red (it is not written in black, or in blue, or in green, or in any other color).

Call the information associated with the source $I(s)$ and the information associated with the receiver $I(r)$. The information *transmitted* $I_s(r)$ is the information which is generated at s and received by r . It is the overlap between $I(s)$ and $I(r)$. $I_s(r)$ is transmitted from s to r by a *signal*. The information that is generated at s but not received by r is called *equivocation*. The information that is associated with r but is not generated by s is called *noise*. Consider again the case where a three is drawn and you receive the information that an odd number was drawn by reading “odd” written in red on a sheet of paper. Here $I_s(r)$ is the information that an odd number was drawn. That a three was drawn is the equivocation and that “odd” is written in red is the noise. The following diagrams represent all the possible kinds of relationship between

⁶ Thank to Alfredo Tomasetta for suggesting this simple characterization of information acquisition to me.

$I_s(r)$, $I(s)$, and $I(r)$ (E stands for equivocation, N stands for noise; gray areas represent absence of information).




Fig. 1: All $I(s)$ is transmitted, but there is some $I(r)$ that is not $I(s)$. Accordingly, there is some noise, but no equivocation.

$$I_s(r) = I(s); I_s(r) < I(r)$$




Fig. 2: Not all $I(s)$ is transmitted, but all $I(r)$ is also $I(s)$. There is some equivocation, but no noise.

$$I_s(r) = I(r); I_s(r) < I(s)$$




Fig. 3: All and only $I(s)$ is transmitted. There is neither equivocation nor noise—the information transmission is optimal.

$$I_s(r) = I(s) = I(r)$$




Fig. 4: Not all $I(s)$ is transmitted and some $I(r)$ is not $I(s)$. There is both equivocation and noise.

$$I_s(r) > I(s) < I(r)$$




Fig. 5: No information is transmitted. All information is either equivocation or noise.

We saw that the amount of information associated with a given state of affairs is a function of the number of possibilities excluded by that state of affairs. However, exactly *how many* possibilities are excluded by a state of affairs is typically extremely hard, if possible at all, to calculate (Dretske 1981: 53). In the box example above, I presented an oversimplified description of the situation by saying that drawing a three from box A eliminates seven

possibilities (i.e. the possibilities that the drawn number is a one, a two, a four, a five, a six, a seven, and an eight). In fact, the occurrence of that event rules out many more possibilities. For there are many more alternative ways the three-drawing may occur. For example, the draw may be performed by a blindfolded child, or by a drawing machine, or by a dog, and so on. The ball with a “3” written on it may be of a variety of different colors. There may not even have been a draw in the first place, but, say, a horse race, or a football match. There is thus a huge amount of alternative ways things may have gone and therefore a huge amount of possibilities the occurrence of the three-drawing rules out. Calculating the exact amount of information associated with such an event is, therefore, extremely hard a task—to say the least.

Does this imply that the notion of information, so characterized in terms of reduction of uncertainty, is bound for theoretical uselessness after all? No, it does not. This notion of information can still, at the very least, play an important role in the analysis of information transmission. It enables us to *compare* the amount of information generated at a source and the amount of information associated with the receiver and thereby *calculate* whether the amount information transmitted is equal to or smaller than the amount of information generated (Dretske 1981: 54). Even though we have a hard time calculating the exact amount of information which is generated by the occurrence of the three-drawing from box A (because listing all the possible alternatives is just too hard a task), we can still draw important conclusions about how much of the information generated by that event is transmitted to you (the receiver). We can say that, if you come to know that a three was drawn, you acquire less information than is generated by the three-drawing event (because that event also generates information about the color of the ball, how the drawing took place, and so on). We can also say that, if you come to know that a three was drawn, and I come to know that an odd number was drawn, you have received more information than I have about the relevant event.

It thus seems that the amount of information transmitted by a signal partly depends on how finely-grained alternatives such a signal individuates. The signal transmitting the information that a three was drawn individuates alternatives in a finer-grained way than the signal transmitting the information that an odd number was drawn. Therefore, it transmits a greater amount of information.

To sum up, acquiring information about a state of affairs s entails ruling out the possible worlds in which s is not the case. The amount of information transmitted by a signal from a source s to a receiver r is a function of the number of possibilities about s which can be ruled out at r in virtue of the

receiving of that signal. The information at s may exceed the information transmitted to r ; the exceeding information, in this case, is called *equivocation*. Some information at r may not be generated by s ; here the exceeding information is called *noise*. When there is neither equivocation nor noise, there is *optimal information transmission*.

3. Information acquisition and epistemic significance

Information acquisition is a feature that most epistemically significant phenomena (i.e. those listed in §1.1) share.

That propositional knowledge typically involves information acquisition seems uncontroversial.⁷ Notwithstanding widespread disagreement as to how exactly propositional knowledge should be analyzed, our intuitions seem to converge on the idea that lack of information acquisition usually implies lack of knowledge. This is consistent with the typical interpretation of Gettier cases.⁸ Smith is justified in believing that Jones will take the job and that Jones has ten coins in her pocket; therefore, she justifiably believes that the woman with ten coins in her pocket will get the job. However, not Jones but Smith has actually got the job and, as it turns out, Smith has ten coins in her pocket. Although Smith has the justified true belief that the woman with ten coins in her pocket will get the job, she does not know that. One way to (at least partially) explain our intuition that Smith lacks knowledge is by noticing that Smith's belief about who has got the job is not formed upon acquisition of information, for no information is transmitted from the source event (the woman with ten coins in her pocket getting the job) to the receiver (Smith).

Although knowledge seems to bear a constitutive connection with information acquisition, the latter is only seldom mentioned in specific analyses of knowledge. Some philosophers, however, do have integrated the notion of information in their analyses of knowledge. Most notably, Fred Dretske (1981) gives information an essential role in his definition of (perceptual) knowledge: " K knows that s is $F = K$'s belief that s is F is caused (or causally sustained) by the information that s is F ." (1981: 86). Although Dretske's is the most famous and perhaps the most thorough information-based theory of knowledge, other authors have given information an important role in their account of knowledge as well. Richard Foley (2012), for

⁷ I say 'typically involves' rather than 'entails' because it is not straightforward that analytic knowledge (that is, knowledge of analytic propositions) implies information acquisition. One might wonder, however, whether analytic knowledge is epistemically significant.

⁸ Cf. Dretske (1981: 96-99).

instance, argues that having (the right kind of) information is what turns true belief into knowledge: “Whether a true belief counts as knowledge thus hinges on the importance of the information one has and lacks.” (2012: 5). Information-involving accounts of knowledge have also been provided, among others, by Floridi (2006), Adams (2010), and Hannon (2014).

Plausibly, not only propositional knowledge, but also other epistemically significant phenomena involve information acquisition. Take *knowledge of a subject matter*. For one thing, knowledge of a subject matter partly involves propositional knowledge; given that propositional knowledge implies information acquisition, knowledge of a subject matter implies information acquisition too. This is also suggested by our intuitions. Would you attribute to me knowledge of physics if I did not possess any information about physics at all?

Understanding also seems to involve information acquisition. For you to *understand that* nothing can travel faster than light, you must, at the very least, acquire the information that nothing can travel faster than light. Possibly, you also need acquire further information about physical facts that are relevant to light travelling and light speed. Similarly, for you to *understand why* continuous economic growth is unsustainable, you need possess the information that continuous economic growth is unsustainable, as well as some information concerning how the economy works—at the very least, information about the finiteness of Earth’s resources, Earth’s waste-absorption capacity, and the relationship between resources regeneration pace and human consumption pace. Finally, for you to understand psychology (*objectual understanding*), you must possess information about how the human mind works.

Perceiving clearly involves information acquisition. By seeing the cup in front of you, you acquire visual information about its shape and color. By grabbing it, you acquire tactile information about its shape and texture. By taking a sip from it, you acquire gustatory information about the taste of the coffee the cup contains. Indeed, arguably, providing us with information about our surroundings is the primary function of perception.

Introspecting, arguably, involves acquiring information about one’s experiences. When you introspectively attend to your stomachache experience to report its phenomenal character to the doctor, you acquire information about that experience. When you introspectively try and figure out what your current emotion is, or what your current mood is, you collect information about your emotional experience, or mood experience.

I am not sure what to say about *knowing how*. On the one hand, I have the intuition that for me to learn (i.e. come to know) how to play *Suzanne* on the guitar, I must acquire some information—arguably, information about the way to move my fingers and hands. On the other hand, it does not seem to me that such information is propositional in form—it does not seem to consist in the ordered series of propositions: <I should place my left first finger on the fourth string, first fret, my left second and third finger on the second fret, second and third string respectively>, <with my right hand, I should pick the first and the fifth string with my thumb and fourth finger respectively>, <I should pick the third string with my second finger>, and so on. For one thing, this information is not sufficient for me to know how to play *Suzanne* (if, for instance, I do not know how to play the guitar in the first place). For another thing, this information does not seem to be necessary for me to know how to play the song: I, who know how to play *Suzanne* on the guitar, needed to actually play it, and look at what my fingers were doing, in order to be able to write the ordered series of propositions above. So, although the idea that there is information acquisition in knowledge-how strikes me as attractive, an adequate characterization of the kind and form of the information acquired in knowledge-how, is not straightforward and would exceed the scope of the present work.

Mere justified belief and mere true belief, in contrast, do *not* entail information acquisition. As noted in §1, *justified beliefs* may be false. You may justifiably believe that *a* is *F*, while *a* not being *F*. Yet, one can only acquire information that *a* is *F* if *a* is in fact *F*: information acquisition is factive. Since mere justified belief is not factive, having a justified belief does not imply information acquisition. True beliefs, by contrast, *are* factive. However, as I have shown in §1.2, true beliefs may be formed by mere luck, independently of the presence of any information transmission from what the belief is about to the person holding the belief. Therefore, mere true beliefs do not imply information acquisition.

Although *some* of them (mere justified belief, mere true belief, and, perhaps, knowledge-how) do not necessarily involve information acquisition, several of the phenomena that I listed in §1.1 *do* display information acquisition. Importantly, all the *paradigmatic* epistemically significant phenomena (propositional knowledge, understanding, and knowledge of a subject matter) involve information acquisition. Moreover, information acquisition seems, intuitively, to be (part of) what *makes* the latter epistemically significant at all—as long as epistemically significant phenomena are supposed to ‘tell’ us something about the world. This seems to give us at least *prima facie* motivation to think that involving information

acquisition is *sufficient* for something to be considered as epistemically significant.⁹

You might think that this conclusion is too strong: information acquisition alone is insufficient for epistemic significance—it must be supplemented with a further condition. Although I am inclined to believe that information acquisition *is* sufficient for epistemic significance, I am happy to retreat to a weaker thesis, namely the thesis that information acquisition is a *mark* of epistemic significance, where a *mark* of *F*, as explained in §1, is intended here as a feature such that *x*'s displaying it gives us good reasons for thinking that *x* is *F*. So, if a phenomenon displays information acquisition, this gives us good reasons for thinking that such a phenomenon is epistemically significant.

4. Knowledge by acquaintance involves information acquisition

In the previous section I have argued that information acquisition is a mark of epistemic significance. In this section, I argue that knowledge by acquaintance displays such a mark—it involves information acquisition. If so, then there are good reasons for thinking that knowledge by acquaintance is epistemically significant. Along the way, I will also defend some increasingly stronger claims about the relationship between knowledge by acquaintance and information acquisition. I will argue that, first, knowing a phenomenal state ϕ by acquaintance provides one with the *maximal amount* of information about ϕ 's phenomenology and, second, that knowing ϕ by acquaintance provides one with *all and only* information about ϕ 's phenomenology. The latter is a special epistemic property of knowledge by acquaintance, that I call *complete and perfect grasp*.

4.1. Knowledge by acquaintance and information acquisition

Recall our characterization of information acquisition above:

You acquire information about a state of affairs *s* iff you can rule out the possible worlds in which $\neg s$.

If we apply this to the specific case of information about phenomenal states, we have that

⁹ Incidentally, it may be worth noting that if involving information acquisition is sufficient for being an epistemically significant phenomenon, then information acquisition is itself an epistemically significant phenomenon. Thanks to Uriah Kriegel for pointing this out to me.

You acquire information about a phenomenal state ϕ iff you can rule out the possible worlds in which $\neg\phi$.

Therefore, if being in a certain introspective state enables you to exclude the possible worlds in which $\neg\phi$, this is sufficient for you to acquire information about ϕ and is thereby a good reason for thinking that such introspective state is epistemically significant.

Now, it seems that primitively introspecting a phenomenal state ϕ , and thereby knowing ϕ by acquaintance,¹⁰ does enable you to rule out the possible worlds in which $\neg\phi$. For one thing, when you attend to ϕ in primitive introspection, you can rule out the possible worlds in which there is *no phenomenal state* at all. As noted in Chapter 4, there can be no hallucination in primitive introspection. Therefore, if you have a state of primitive introspection as of a certain phenomenal state, there must be a phenomenal state your introspective state is directed at. At the very least, when you primitively introspect a phenomenal state you acquire the information that *a* phenomenal state is present—that there is *a* phenomenal state rather than *no* phenomenal state. Accordingly, you can rule out the possible worlds in which there is no phenomenal state at all.

For another thing, being introspectively acquainted with ϕ enables you to rule out the possible worlds in which there is a phenomenal state which is *different* from, and incompatible with, ϕ . For instance, your primitively introspecting a homogeneously greenish visual experience enables you to rule out the possible worlds in which you have a visual experience as of any other color.

These two considerations are sufficient to show that knowledge by introspective acquaintance involves information acquisition. Knowledge by acquaintance, then, displays our first mark of epistemic significance.

It may be pointed out that, if that is it, information acquisition in knowledge by acquaintance is quite meager—the reader may be disappointed. Indeed, although the two considerations above show that knowledge by acquaintance satisfies the minimal requirement for information acquisition (i.e. the exclusion of *some* possibilities), this is not the whole story. In what follows, I argue that not only knowing ϕ by acquaintance enables you to rule out some possibilities; it enables you to exclude *all* the relevant alternatives. Accordingly, it provides you with the *maximal* amount of information about ϕ 's phenomenology (§4.3). Moreover, knowing ϕ by acquaintance provides you

¹⁰ Recall, introspective knowledge by acquaintance is the kind of knowledge one has by merely primitively introspecting a certain phenomenal state.

with *all and only* the information about φ 's phenomenology which is generated by φ 's occurrence (§4.4). (Nota bene: you may reject the two latter claims without also dismissing the main idea that knowing φ by acquaintance enables you to acquire some information about φ .)

4.2. *The format of information in knowledge by acquaintance*

As noted in §3, information is transmitted from a source (information generator) to a receiver (information receptor). In the case of phenomenal-state introspection,¹¹ what generates the relevant information is the occurrence of the target phenomenal state φ and what receives the information is the introspective state—call it ι . What is transmitted from φ to ι in phenomenal state introspection is information about φ 's phenomenology.

The receiving introspective state may be a state of primitive introspection or a state of reflective introspection. States of reflective introspection receive information in classificatory (potentially propositional) form and may directly yield introspective *propositional* knowledge. States of primitive introspection, instead, receive information in non-classificatory (thereby non-propositional) form and constitute introspective knowledge *by acquaintance*.

The distinction between classificatory and non-classificatory form of information may be elucidated through Dretske's notion of *digital* vs. *analog* information. Here is how he characterizes the distinction:

I will say that a signal (structure, event, state) carries the information that s is F in *digital* form if and only if the signal carries no additional information about s , no information that is not already nested in s 's being F . If the signal *does* carry additional information about s , information that is *not* nested in s 's being F , then I shall say that the signal carries this information in analog form. When a signal carries the information that s is F in analog form, the signal always carries more specific, more determinate information about s than it is F . (Dretske 1981: 137, emphasis original)

The information that s is G is *nested* in s 's being F iff s 's being F carries the information that s is G (Dretske 1981: 71). For instance, the information that the drawn number is odd is nested in the drawn number's being a three: the drawn number being a three carries the information that the drawn number is odd. If you receive the information that a three was drawn via a sheet of

¹¹ Recall, *phenomenal-state introspection* is the distinctively first-personal method through which one can get knowledge of the phenomenology of one's own current conscious experience.

paper with “3” written on it, you receive that information in digital form. If you receive that same piece of information by looking at a picture of the ball (taken during the drawing), you receive the information that a three was drawn in analog form. In the latter case, you not only receive the information that a three was drawn; you also receive additional information about, say, the color of the drawn ball, the venue where the drawing takes place, the size of the ball with respect to its surroundings, and so on. Your seeing the picture of the drawing therefore provides you with richer, more determinate information about the three-drawing event than your receiving the note.

The process of *digitalization* of analog information is a *classification* process. Fine-grained, specific information is converted into coarser-grained, more generic information. For instance, if I describe to you the content of the three-drawing picture by uttering “A three was drawn”, I am digitalizing (i.e. classifying) the information I receive from the picture in analog form and transmitting it to you in digital form. The digitalization process typically renders the relevant information more manageable, that is, easier to deploy. At the same time, it entails a *loss* of information. The information carried by my utterance “A three was drawn” carries less information than the picture of the drawing: it enables the receiver to exclude a smaller number of alternatives. Dretske deploys the distinction between digital and analog information to explain the difference between the information carried by sensory experiences and cognitive states such as beliefs. Sensory states, on Dretske’s view, carry information in analog form; cognitive states, by contrast, carry information in digital form. As before, digitalization implies classification and thereby loss of information:

Our sensory experience is informationally rich and profuse in a way that our cognitive utilization of it is not. Relative to the information we manage to *extract* from the sensory representation (whatever beliefs may be occasioned by having this kind of sensory experience), the sensory representation itself qualifies as an *analog* representation of the source. It is this fact that makes the sensory representation more like a *picture of*, and the consequent belief a *statement about*, the source. (Dretske 1981: 150, emphasis original, footnote omitted)

Now, back to primitive and reflective introspection. Dretske’s distinction between analog and digital information may turn out useful to understand the difference between the format in which the state of primitive introspection on the one hand and the state of reflective introspection on the other hand, receive information about a target phenomenal state ϕ . The idea is that whereas reflective introspection carries information about ϕ ’s phenomenology in digital form, primitive introspection carries information about ϕ ’s

phenomenology in analog form. This squares well with the idea that reflective introspection is classificatory, whereas primitive introspection is not. When you reflectively introspect that this experience you are having is a papaya-taste experience, you acquire the information that this experience is a papaya-taste experience, information that is digital in form: the information about the rich and complex phenomenology of your experience has been classified as *papaya-taste*. When, instead, you primitively introspect your papaya-taste experience, the information you acquire is analog in form. Accordingly, primitive introspection carries additional information about the phenomenology of your experience—information that exceeds the piece of information that this experience you are having is a papaya-taste experience.

In this framework, propositional introspective knowledge of ϕ , which is yielded by reflective introspection, involves acquisition of information about ϕ 's phenomenology in digital form. Knowledge of ϕ by acquaintance, which is constituted by the state of primitive introspection directed to ϕ , by contrast, involves acquisition of information about ϕ 's phenomenology in analog form.

With the present subsection's discussion, I do not intend to *commit* to the idea that the distinction between the format of the information carried by primitive and reflective introspection respectively maps onto Dretske's analog vs. digital distinction. What I *am* committed to is that the difference between the information carried by a state of primitive introspection and that carried by a state of reflective introspection lies in the fact that the former but not the latter is classificatory in form. I think, however, that Dretske's notion is a helpful tool to elucidate such a difference.

4.3. *Knowledge by acquaintance and amount of information*

Consider a subject S who reflectively introspects that ϕ is F . On the basis of this state of reflective introspection, suppose, S comes to know (i.e. have the piece of propositional knowledge) that ϕ is F . By acquiring the piece of propositional introspective knowledge that ϕ is F , S acquires the information that ϕ is F . S can thereby rule out the possible worlds in which ϕ is not F .

Now, the *amount* of information conveyed by the piece of propositional introspective knowledge that ϕ is F depends of F 's *degree of determinacy*. Roughly, the more determinate, the more information. This was already mentioned in §2 but let us consider it more slowly.

Take the following phenomenal properties: phenomenal crimson, phenomenal red, phenomenal color. Straightforwardly, the first is the most specific (the most determinate), the last is the most generic (the less

determinate). The three properties bear a determinable-determinate relationship to each other: phenomenal color is a determinable of phenomenal red, which is a determinable of phenomenal crimson; phenomenal crimson is a determinate of phenomenal red, which is a determinate of phenomenal color (of course, by transitivity, phenomenal color is also a determinable of phenomenal crimson and phenomenal crimson is also a determinate of phenomenal color).

When you come to know that ϕ is *phenomenal red*, you can thereby rule out the possible worlds in which ϕ is *not* phenomenal red, i.e. the possible worlds in which ϕ is olfactory, pain, anger, phenomenal blue, phenomenal green, phenomenal yellow, and so on and so forth. When you come to know that ϕ is *phenomenal color*, you can still rule out the possible worlds in which ϕ is olfactory, pain, anger, and so on (i.e. the possible worlds in which ϕ is not phenomenal color); however, in this case you *cannot* also rule out the possible worlds in which ϕ is phenomenal blue, phenomenal green, phenomenal yellow, and so on, for your coming to know that ϕ is phenomenal color does not provide you with information about *what particular phenomenal color* ϕ is. Since, recall, the amount of information acquired is a function of the number of possibilities excluded, when you come to know that ϕ is phenomenal color, you acquire *less* information than when you come to know that ϕ is phenomenal red. When, instead, you come to know that ϕ is *phenomenal crimson*, you can not only rule out the possible worlds in which ϕ is olfactory, pain, anger, phenomenal blue, phenomenal green, phenomenal yellow, and so on (as in the case where you come to know that ϕ is phenomenal red); in this case, you can, *in addition*, rule out the possible worlds in which ϕ is phenomenal scarlet, phenomenal vermilion, phenomenal carmine, phenomenal cherry, and so on. Therefore, when you come to know that ϕ is phenomenal crimson, you acquire *more* information than when you come to know that ϕ is phenomenal red (and, *a fortiori*, more information than when you come to know that ϕ is phenomenal color), for you can rule out a greater number of possibilities.

A different angle to put the same idea is the following. Information acquisition can be understood in terms of *reduction of uncertainty*, where uncertainty spans across a certain range of possibilities: the more possibilities in that range are excluded, the more uncertainty is reduced (and, therefore, the more information is acquired). The number of possibilities in the range is partly determined by how coarsely- or finely-grained such possibilities are individuated. Here are some examples of ranges of possibilities (in introspective inquiry) which differ depending on how finely-grained possibilities are individuated: (a) perceptual experience, bodily sensation,

emotional experience, imagination experience, etc.; (b) visual experience, auditory experience, tactile experience, olfactory experience, gustatory experience; (c) phenomenal red, phenomenal green, phenomenal yellow, phenomenal blue, etc.; (d) phenomenal crimson, phenomenal scarlet, phenomenal vermilion, phenomenal carmine, phenomenal cherry, and so on. It is easy to see that, the more finely-grained possibilities are individuated, the larger the range of possibilities across which uncertainty in introspective inquiry spans. When you come to know that your experience is *phenomenal crimson*, the introspective uncertainty which is thereby reduced spans across a range of possibilities which is larger than the range of possibilities across which spans the introspective uncertainty reduced by your coming to know that this experience is a *visual experience*, for in the former case the possibilities are more finely-grained individuated. Therefore, your coming to know that your experience is phenomenal crimson enables you to rule out a greater number of possibilities (and thereby acquire more information) than your coming to know that this experience is a visual experience.

Now, here comes the key idea of this subsection. In primitive introspection, possibilities are *maximally fine-grained* individuated: the relevant range of possibilities includes *all possible different phenomenologies* a subject may instantiate. Therefore, in knowledge by introspective acquaintance, the possible alternatives (ways the target phenomenal state ϕ could be) are individuated by the *maximal degree of determinacy*. By knowing ϕ by acquaintance, you can thereby rule out any possible world in which ϕ is phenomenally different, however slightly, from how it is. For example, by *primitively* introspecting your phenomenal crimson experience, you can not only rule out the possible worlds in which your experience is phenomenal scarlet, phenomenal vermilion, phenomenal carmine, phenomenal cherry, and so on. Put very roughly, in this case you can rule out *all the worlds in which your experience is phenomenally different, however slightly, from how it is*, say, worlds in which it is of a slightly different shade, or slightly phenomenally brighter, or slightly phenomenally more intense, and so on.

Therefore, when you primitively introspect ϕ (and thereby come to know ϕ by acquaintance), you can rule out the greatest number of possibilities about ϕ 's phenomenology: there is no introspective state which could enable you to rule out more possibilities about ϕ 's phenomenology. Accordingly, when you come to know ϕ by acquaintance, you get the *relative* maximal amount of information about ϕ 's phenomenology, where by 'relative' I mean 'relative to other introspective states or processes': no other introspective state or process gives you more information than primitive introspection does. (In §4.4 I will argue that primitive introspection also gives you the *absolute*

maximal amount of information about φ 's phenomenology: it enables you to acquire all the information about φ 's phenomenology that is generated by the occurrence of φ .)

A clarification is in order. What we are concerned with here is information about the *phenomenology* of φ , *not* with information about the *stimulus* that brings φ about. Accordingly, when you primitively introspect your phenomenal crimson experience, your coming to know the experience by acquaintance does *not* allow you to exclude the possible worlds in which the stimulus that causes your visual experience (light reflected by the surface before you) is (slightly) different. For one thing, it may be that some differences in stimulus cannot be detected or are otherwise not reflected in the phenomenology of your visual experience.

In Chapter 4, I have distinguished between qualitative, quantitative, and relational phenomenal properties. Recall, qualitative phenomenal properties are the phenomenal properties that define the qualitative aspect which characterizes each kind of phenomenal state—phenomenal properties such as reddishness, painfulness, burning painfulness, and so on; quantitative phenomenal properties are properties in virtue of which qualitative phenomenal properties are modulated—phenomenal properties such as intensity and clarity; relational phenomenal properties are phenomenal properties phenomenal states have in virtue of bearing certain relations to other phenomenal states—properties such as phenomenal unity, phenomenal mereology, and phenomenal salience. A legitimate question arises about whether knowledge by acquaintance involves information acquisition about all or only some kinds of phenomenal properties. My tentative proposal is the following. Primitive introspecting φ enables you to acquire information about φ 's *qualitative* properties. Accordingly, when you primitively introspect φ , you can thereby rule out the possible worlds in which φ has different (however slightly) qualitative phenomenal properties (e.g. it enables you to rule out the worlds in which your current color experience is of a slightly different shade). Arguably, primitively introspecting φ also gives you information about φ 's *quantitative* phenomenal properties. For instance, when you primitively introspect your current color experience, you can rule out the possible worlds in which the intensity of color phenomenology is different (however slightly). What primitive introspection does *not* provide you with, on my view, is information about φ 's *relational* phenomenal properties. For acquiring information about relational phenomenal properties requires comparing *two or more* phenomenal states, or different temporal stages of the same phenomenal state; but primitive introspection targets *one* phenomenal state (or phenomenal-state temporal stage) at a time. Therefore, information

about relational phenomenal properties is not acquired via primitive introspection.¹² You may object that there might be cases in which a relational property between two phenomenal states *can* be primitively introspected. Suppose you have a visual experience as of a black dot in the left half of your visual field and a visual experience as of a red dot in the right half: could not you introspect the (phenomenal) spatial relation between the two phenomenal states?¹³ I think that your introspecting the phenomenal spatial relation is *not* an instance of *primitive* introspection. Although the linguistic *expression* describing your introspective state directed at the spatial difference may *prima facie* suggest that this is a case of primitive introspection ('introspecting the spatial relation' suggests that the introspective attitude is objectual rather than propositional), a closer examination reveals that it is not. For introspecting the phenomenal spatial relation really amounts to introspecting *that* one phenomenal state bears a certain phenomenal spatial relation to the other. Therefore, your introspective state in this case is in fact propositional.¹⁴

In this framework, then, by knowing ϕ by acquaintance you acquire information about ϕ 's qualitative and quantitative phenomenal properties. Assuming, as specified in Chapter 4, that all and only ϕ 's qualitative phenomenal properties are *essential* to it, we have that knowing ϕ by acquaintance gives you information about all ϕ 's essential phenomenal properties. You thereby know the *phenomenal essence* of ϕ . At this point, you might be worried by a remark of Dretske's about essential properties and information:

And if (as some philosophers suppose) individuals have some of their properties *essentially*, then the possession of these properties by these individuals generates zero information. (Dretske 1981: 12)

The reasoning is that, if essential properties are properties an individual has *necessarily* (if it exists), then there is *no* possible world in which that individual exists and does not have those properties. Accordingly, there are

¹² Though it may be acquired via different means, e.g. introspection *plus* memory. For instance, one may come to know that this phenomenal state is more salient than it was before by introspecting it now plus remembering how it was before.

¹³ Thanks to Uriah Kriegel for this objection.

¹⁴ Cf. Dretske (1993: 266-67): "When perceptual verbs (including the generic 'aware of' and 'conscious of') are followed by abstract nouns (the difference, the number, the answer, the problem, the size, the colour) and interrogative nominals (where the cat is, who he is talking to, when they left), what is being described is normally an awareness of some (unspecified) fact. The abstract, noun phrase or interrogative nominal stands in for some factive clause. Thus, seeing (being conscious of) the difference between *A* and *B* is to see (be conscious) *that* they differ. If the problem is the clogged drain, then to be aware of the problem is to be aware that the drain is clogged."

no possibilities that can be excluded by an individual's possessing its essential properties. Therefore, an individual's possessing its essential properties generates zero information. If no information is generated, no information can be acquired. So, it would seem that, since a phenomenal state φ 's possessing its essential phenomenal properties generates zero information, no information is acquired by primitively introspecting φ .

However, this conclusion does not follow. Even if φ 's possessing its essential properties F , G , and H does not generate any information, φ 's *presence* does. For even if it is necessary that if φ is present, then φ has F , G , and H , *the occurrence* of φ is *not* necessary, but contingent. Therefore, when you primitively introspect φ , you do acquire information about its essential phenomenal properties: you acquire the information that these properties are instantiated (rather than not).

To sum up, by primitively introspecting φ , you acquire the maximal amount of information about φ 's phenomenology, for no other introspective state could enable you to rule out a greater number of possibilities about φ 's phenomenology than your state of primitive introspection directed to φ .

It may be pointed out that, however, the state of primitive introspection may not be the *only* introspective state which provides you with the maximal amount of information about φ 's phenomenology—other introspective states may enable you to rule out the same (number of) possibilities about φ 's phenomenology.¹⁵ Consider again what David Chalmers (2003) calls 'direct phenomenal concepts'. Recall, a direct phenomenal concept is a phenomenal concept which is formed upon attending to the phenomenal state it refers to and whose content is constituted by such a phenomenal state. The direct phenomenal concept Φ , which refers to the phenomenal state φ , is formed upon attending to φ and is partly constituted by φ . Accordingly, its existence depends on φ 's existence. By combining the direct phenomenal concept Φ with the pure phenomenal demonstrative E (i.e. the concept by which you refer to φ demonstratively and non-descriptively, i.e. by a mere act of ostensive attention), you can form the proposition that E is Φ . Therefore, when you introspectively attend to φ , you can acquire the piece of propositional knowledge that E is Φ . Now, since Φ is partly constituted by φ , that piece of propositional knowledge enables you to rule out all the possible worlds in which φ has a (however slightly) different phenomenology. But if it is so, then not only knowledge by acquaintance, but also propositional knowledge

¹⁵ I owe this comment to Marie Guillot.

involving direct phenomenal concepts enables you to acquire the maximal amount of information about φ 's phenomenology.

I am happy to acknowledge that, if there are direct phenomenal concepts (a question I am not going to take a stance on here), then knowledge by acquaintance is not the *only* kind of knowledge that can provide you with the maximal amount of information about φ 's phenomenology. And indeed, all I have argued for in this subsection is the weaker thesis that no piece of propositional knowledge can provide you with a *greater* amount of information than knowledge by acquaintance provides—which is consistent with there being pieces of propositional knowledge which provide you with the same amount of information.

However, there is a thesis in the vicinity to which I am attracted, which is weaker than the thesis that *only* knowledge by acquaintance provides you with maximal information amount, but stronger than the thesis that knowledge by acquaintance provides you with maximal information amount, *but there are special pieces of propositional knowledge* that do exactly the same job. The thesis is the following: *only* knowing φ by acquaintance is a *ground* to acquire the maximal amount of information about φ 's phenomenology—special pieces of propositional knowledge may provide you with the same amount of information, but only because such information is transmitted to them by a state of primitive introspection (which constitutes the relevant knowledge by acquaintance). The idea is that, if the piece of propositional knowledge that E is Φ provides you with maximal information amount, this is only in virtue of the fact that such information is transmitted to the direct phenomenal concept Φ by the primitive introspective state ι directed to φ . The reasoning underlying this idea was already developed in §4 of Chapter 5. The direct phenomenal concept Φ is formed upon attending to φ . So, it is in virtue of your introspectively attending to φ that Φ comes to carry information about φ 's phenomenology. Arguably, the acquisition of such information is prior to Φ 's existence—until the information it is supposed to carry is not acquired, Φ cannot be formed. On my view, this preliminary information acquisition is fulfilled by the state of primitive introspection (which is a pre-conceptual state). Accordingly, the information about φ 's phenomenology is generated by φ and transmitted by φ to the primitive introspective state ι . Once the information reaches ι , it is poised for being transmitted to the direct phenomenal concept Φ .

So, although I am happy with the weaker thesis that knowledge by acquaintance is a way to acquire maximal information amount, but some pieces of propositional knowledge may do the same job, I am inclined to prefer the stronger thesis that even if there are pieces of propositional knowledge

that carry maximal information amount, their doing so is grounded in knowledge by acquaintance.

4.4. Complete and perfect grasp

According to an important tradition in philosophy, tracing back to at least Descartes, introspection has a special epistemic property: it is infallible. This means that the beliefs which are formed through introspection cannot be false: having them entails their truth. As is well known, this tradition has been harshly (and partly rightly) criticized, on the basis of the fact that we are, at least sometimes, *wrong* about our experiences: some of our introspective judgments are false. Indeed, reflective introspection is *not* infallible. I suggest, however, that *primitive* introspection does possess an interesting epistemic property, which is different from infallibility but special nonetheless.

To be sure, primitive introspection *cannot* be infallible. For, strictly speaking, infallibility is a property of truth-apt mental states, such as beliefs: a belief is infallible iff having it entails its truth. Now, of course, primitive introspection cannot be infallible in this sense. Since primitive introspection does not involve any classification, its occurrence is independent of the formation of any belief, or of any truth-apt mental state. Accordingly, truth and falsity do not apply to primitive introspection. What primitive introspection can have, at most, is an epistemic property, call it *infallibility**, which is *analogous* to infallibility, but applies to *non-classificatory* mental states, rather than truth-apt mental states such as beliefs. Here is a tentative characterization of *infallibility**: an introspective state is *infallible** iff, if one has it, an experience must be present. *Infallibility** is therefore equivalent to what I have called, in Chapter 4, *no introspective hallucination*.

That primitive introspection is *infallible** is no news. It was a constraint to our quest for a satisfactory account of the metaphysics of primitive introspection in Chapter 4 and indeed it is entailed by the integration account developed in that same chapter. So, even if it is a real property of primitive introspection, *infallibility** may not strike us as so interesting at this point.

There is, however, a more complex special epistemic property, which implies but does not reduce to *infallibility**, whose articulation may capture the epistemic peculiarity of primitive introspection and help us understand how the traditional idea may interestingly apply to primitive introspection.

What I take to be the most interesting special epistemic property of primitive introspection is that it provides the subject with a *complete and*

perfect grasp of the phenomenology of their current experience. The idea, roughly, is that primitively introspecting φ provides you with *all and only* the information about φ 's phenomenology which is generated by φ 's occurrence.¹⁶ That primitive introspection gives you a 'complete' grasp means that it provides you with *all* the information about φ 's phenomenology; that it gives you a 'perfect' grasp means that it provides you with *only* the information about φ 's phenomenology. Let us consider this more slowly.

By 'providing a grasp' I mean that primitive introspection provides the subject with information about φ 's phenomenology. This is what I have argued for in §4.1.

By 'complete' I mean 'maximally comprehensive' or 'exhaustive'. In §4.3 I have argued that primitively introspecting φ provides the subject with the relative maximal amount of information about φ 's phenomenology, where the amount of information is *maximal* with respect to the amount of information other introspective states (or pieces of introspective knowledge) can carry (this is what I meant by 'relative'). Here I argue for the stronger thesis that primitively introspecting φ also provides the subject with the *absolute* maximal amount of information about φ 's phenomenology. The criterion for *maximality*, here, is not constituted by other potential introspective states (or pieces of introspective knowledge), but by the target phenomenal state itself. Let me explain what this means. The occurrence of a phenomenal state φ generates a certain amount of information about φ 's phenomenology. The thesis defended in §4.3 was that all the information about φ 's phenomenology that *can* be transmitted by φ to an introspective state, be it the *totality* of the information generated, or only a *part* of it, is acquired by primitive introspection: primitive introspection enables one to acquire the maximal amount of information that can be acquired by an introspective state. So, there the focus was on the relationship between the amount of information received by primitive introspection and the amount of information received by other introspective states or pieces of propositional introspective knowledge. Here, instead, the focus is on the relationship between the amount of

¹⁶ In §4.3 I pointed out that, plausibly, information about φ 's *relational* phenomenal properties is not acquired by primitive introspection. *Prima facie*, this seems to defeat from the start the idea that all information about φ 's phenomenology generated by φ 's occurrence is acquired by primitive introspection. However, the idea remains *secunda facie* undefeated if you consider that information about φ 's relational phenomenal properties is *not only* information about φ , but also about the phenomenal state(s) φ is relevantly related to. So, to be more precise, the *complete and perfect grasp* thesis should be framed as follows: primitively introspecting φ provides you with *all and only* the information about φ 's phenomenology, *which is not also information about any other phenomenal state's phenomenology*, which is generated by φ 's occurrence.

information received by primitive introspection and the amount of information generated by φ . The idea is that not only some, but *all the information* about φ 's phenomenology which is *generated* by φ 's occurrence is acquired by the state of primitive introspection ι directed to φ . By using the terminology explained in §2, there is *no equivocation* in the transmission of information from φ to ι .

As noted, reflective introspection does not provide the subject with maximal information amount. *A fortiori*, it cannot provide the subject with all the information generated by the occurrence of the target phenomenal state—it cannot be maximally comprehensive. Therefore, the transmission of information from a target phenomenal state to a state of reflective introspection always implies a certain amount of equivocation.

By 'perfect' I mean 'perfectly reliable' or 'perfectly accurate'. Reflective introspection may sometimes be mistaken: one might introspect that one has a pain sensation when instead what one has is a coldness sensation. This may happen because the state of reflective introspection carries *extra information*, that is, information that is *not generated* by the occurrence of the target phenomenal state (in our example, the information that the target phenomenal state is a pain sensation). In other words, information transmission between a target phenomenal state and a state of reflective introspection may involve a certain amount of *noise* (using §2's terminology). This, of course, does not prevent reflective introspection from being reliable: it may well be that such mistakes are rare enough not to compromise reflective introspection's reliability. However, it does prevent it from being *perfectly* reliable. Primitive introspection, instead, is perfectly reliable. If one primitively introspects a certain phenomenal state φ , all the information carried by the primitive introspective state ι comes from φ : only the information generated by φ 's occurrence is transmitted to ι . In other words, the transmission of information from φ to ι never involves noise.

The information-transmission situation we have in primitive introspection, therefore, is the one described by *Figure 3* in §2: the information transmitted is identical to the information generated which is identical to the information received—all and only the information generated at the source is transmitted to the receiver. There is neither noise nor equivocation. This, together with the following remark by Dretske, squares very well with what has been said in Chapter 4 about the metaphysics of primitive introspection:

E and *N* [that is, *equivocation* and *noise*] are measures of the amount of *independence* between the events occurring at the source and at the receiver. [...] If there is zero independence (maximum dependence), then *E* and *N* will

be zero and the amount of transmitted information $I_s(r)$ will be optimal: $I_s(r) = I(s)$. (Dretske 1981: 23)

The idea is that, if the events occurring at the receiver are fully dependent on the events occurring at the source, the information transmission is optimal—there is neither noise nor equivocation. As noted in Chapter 4, on the integration account (as well as on all versions of the acquaintance view) the state of primitive introspection depends on the phenomenal state it is directed at for its existence: it fully depends on the target phenomenal state. This explains why primitive introspection features optimal information transmission. The absence of noise and equivocation in primitive introspection is therefore grounded in the metaphysics of the primitive introspective state.

In sum, when one primitively introspects a phenomenal state ϕ , one has a perfect and complete grasp of ϕ 's phenomenology: (i) one gets information about ϕ 's phenomenology (grasp), (ii) one gets *all* the information about ϕ 's phenomenology which is generated by ϕ 's occurrence (maximal comprehensiveness or exhaustiveness), and (iii) one gets *only* the information about ϕ 's phenomenology which is generated by ϕ 's occurrence (perfect reliability or accuracy). This epistemic property, which is distinctive of primitive introspection, is grounded in the metaphysical structure of the primitive introspective state.

Incidentally, providing a complete and perfect grasp entails being infallible*. If, necessarily, all the information one gets through primitive introspection comes from the target experience, it cannot be the case that one primitively introspects an experience which is not present. Therefore, there is no hallucination in primitive introspection.

*

To sum up, here is the interim conclusion of this long section. I have defended the following four increasingly strong theses:

1. Knowing ϕ by acquaintance provides you with information about ϕ 's phenomenology.
2. Knowing ϕ by acquaintance provides you with the relative *maximal* amount of information about ϕ 's phenomenology—no piece of propositional knowledge provides you with a greater amount of

information (although there may be pieces of propositional knowledge that provide you with the same amount of information).

3. *Only* knowing φ by acquaintance is a *ground* to acquire the maximal amount of information about φ 's phenomenology—special pieces of propositional knowledge (such as those involving direct phenomenal concepts) may provide you with the same amount of information, but only because such information is transmitted to them by a state of primitive introspection.
4. Knowing φ by acquaintance provides you with *all and only the information* about φ 's phenomenology which is generated by φ 's occurrence—it gives you a complete and perfect grasp of φ 's phenomenology.

For the strict purpose of the present chapter, it is enough that you be persuaded about the truth of 1. However, I hope to have convinced you of the interest of the other three theses. Thesis 2 gives thesis 1 more substance: it tells us that knowledge by acquaintance provides us with more interesting information than that a phenomenal state whatsoever is present. Thesis 3 tells us about knowledge by acquaintance's *epistemic priority*: if 3 is true, then knowledge by acquaintance has a distinctive epistemic status with respect to propositional introspective knowledge—at least in a certain respect, it is epistemically prior. Thesis 4, finally, tells us about knowledge by acquaintance's *special epistemic properties*: if 4 and 3 are both true, knowledge by acquaintance has a distinctive epistemic property—one that no piece of propositional introspective knowledge can have.

5. Epistemic evaluability

In the study of human action and behavior three distinct dimensions are often distinguished (see e.g. Kvanvig 2014): a *descriptive* dimension, which concerns what we typically do, a *normative* dimension, which concerns the acts we ought and ought not to perform, and an *evaluative* dimension, which concerns what objects or states of affairs are good or valuable and what are not. Whether these dimensions are merely conceptually distinct or also metaphysically distinct, and what their interrelations are, are matter of debate in metaethics. One question, for example, concerns whether the normative dimension reduces to the descriptive dimension—Hume's Is/Ought Argument ('ought' cannot be derived from 'is') and Moore's Open Question Argument (the question whether good is constituted by any natural property cannot be settled on conceptual grounds) are famous contributions to this debate. Another question concerns the relationship between the normative dimension

and the evaluative dimension—this is related to the so-called ‘Euthyphro predicament’. Arguably, we ought to pursue the good and ought not to pursue the bad. Yet, ought we to pursue the good because it is good for us to pursue it, or is what we ought to pursue good because we ought to pursue it?

Of course, discussing these questions falls far beyond the scope of the present work. The point is that the same tripartite distinction can be drawn in the more specific domain of human *epistemic practice*. Here, we can distinguish between the descriptive dimension concerning how the epistemic inquiry *is* conducted, the normative dimension concerning how the epistemic inquiry *ought to be* conducted, and the evaluative dimension concerning what makes epistemic inquiry *good* or valuable. The distinction can be applied to specific aspects of the epistemic inquiry. Take *belief*, for example. We can distinguish between what people typically believe (descriptive dimension), how beliefs ought and ought not to be formed (normative dimension), and what makes a belief good or valuable (evaluative dimension). As before, exploring the interrelations between these dimensions is part of the meta-epistemologist’s job and is strictly not of our present concern.

What *is* of our concern is the fact that *the epistemic domain displays an evaluative and a normative dimension*. Not only that: normativity and evaluability are typically considered as *essential* or *intrinsic* aspects of epistemic practice (Kim 1988).

Let us consider the analogy between epistemology and metaethics more closely. In metaethics, we distinguish the (moral) *good*, or *value* and the (moral) *right*, or *norm*. What is susceptible to be *good* are states of affairs. What is susceptible to be *right* are human actions that produce determinate states of affairs. The evaluative dimension concerns what makes a state of affairs good. The normative dimension concerns what actions are right. As noted, the relationship between the right and the good is matter of debate. It is plausible to suppose, however, that there must be a connection between right actions and good states of affairs. Analogously, in epistemology, we can distinguish the *epistemic good* (or epistemic value) and the *epistemic right*. What is susceptible to be good is the epistemic inquiry, or aspects of the epistemic inquiry such as beliefs. What is susceptible to be right is the way in which the epistemic inquiry is conducted—for example, the way in which beliefs are formed. The evaluative dimension concerns what makes an epistemic inquiry (or an aspect of it such as belief) good. The normative

dimension concerns how the epistemic inquiry ought to be conducted—e.g. how beliefs ought to be formed.¹⁷

There is no general agreement about what the epistemic *good* (or epistemic value) is. *Truth* has traditionally been considered as the fundamental epistemic value, but recent epistemological debate has revolved around the existence of alternative epistemic values, such as *knowledge* (Williamson 2000), *understanding* (Kvanvig 2003), or *rationality*, which are sometimes taken to be higher than truth and/or distinct in kind (rather than just greater in quantity). Regardless of what the epistemic value(s) ultimately is (or are), such a value will constitute the *standard* for your *evaluation* of epistemic inquiries. For instance, if truth is your epistemic value and belief your object of epistemic evaluation, you will deem a certain belief good if that belief is true, bad if it is false. An epistemic practice, in this framework, is good if it is truth-conducive, that is, if it maximizes true beliefs and minimizes false beliefs.

The epistemic *norm* (that is, the way epistemic inquiry ought to be conducted) is typically captured by the notions of *epistemic justification* and *rationality*. In this framework, we ought to conduct inquiry *rationally* (rationality is what regulates our epistemic inquiry): we ought to conduct our epistemic inquiry rationally and ought not to conduct our inquiry irrationally. By the same token, we ought to form our beliefs justifiedly (that is, our beliefs must be justifiedly formed) and we ought not to form our beliefs unjustifiedly.

As before, it is plausible to think that there must be a connection between the epistemic value and the epistemic norm. Arguably, a consideration of the epistemic value(s) settles the epistemic goal(s) to be pursued in epistemic inquiry and a right epistemic inquiry somehow aims to pursue such goals. If, for example, truth is our epistemic value and justification is our epistemic norm, then the relevant question concerns the connection between truth and justification. Spelling out such a connection may turn out to be a hard task (Cohen 1984), one which exceeds our present concern. Besides, things may turn out yet more complex since some epistemologists seem to consider justification and rationality themselves as valuable.

Of course, a fundamental epistemological question regards what justification and rationality consist in, that is, what makes a belief justified or an epistemic inquiry rational. Some epistemologists, for instance, think

¹⁷ Thanks to Uriah Kriegel for explaining to me the relationship between good and right in metaethics and epistemology.

that one's belief that *p* is justified only if one's forming that belief is grounded in the consideration of adequate evidence speaking in favor of *p*'s truth. Or, in a similar vein, some maintain that one's belief that *p* is justified only if one has adequate internally accessible reasons for thinking that *p* is true. (These are versions of *internalism* about justification. See e.g. Feldman and Conee (2001), BonJour (1985).) Others, in contrast, argue that epistemic justification does not depend on internally accessible reasons but rather on factors that are external to the subject. For example, one's belief that *p* is justified only if the belief that *p* is reliably formed, or if there is an appropriate causal chain leading from the fact that *p* to the belief that *p*. (These are versions of *externalism* about justification. See e.g. Goldman (1979)) Regardless of how the dispute about the nature of justification should be settled, the relevant point here is that justification works as belief-formation regulatory factor: it sets the conditions a belief-forming process ought to satisfy. From an internalist standpoint, one ought to believe that *p* only if one has adequate evidence in favor of *p*, or if one has good reasons for thinking that *p* is true. From an externalist standpoint, one ought to believe that *p* only if the belief that *p* is reliably formed, or if there is an appropriate causal chain leading from the fact that *p* to the belief that *p*.¹⁸

A last clarification is in order. In thinking about epistemic evaluability, we should distinguish between the *epistemic* value of an epistemic inquiry and the *overall* value of that inquiry—the two may not coincide. Besides its epistemic value, a true belief (say) may also have a moral value, or a practical value.¹⁹ (I keep taking true belief as my main example, but what I say can be applied to knowledge, understanding, and other aspects of the epistemic inquiry.) True belief may be valuable because having true beliefs enables us to *successfully do* things (e.g. having a true belief about the way to the airport enables you to successfully get to the airport and catch your flight) or because it enables us to achieve our moral goals, or conduct morally valuable lives (Zagzebski 2003). Some true beliefs, though, may be practically or morally value-neutral, or even *disvaluable*: being the witness of a mafia-related crime provides you with a true belief, but also with chances that someone will come and try to kill you. Less dramatically, if someone tells you the end of a movie you have not watched yet, your movie-experience will be spoiled by your true belief (Kelly 2003). Moreover, false or evidentially unsupported beliefs may

¹⁸ I am obviously oversimplifying the issue here—there are many more versions of internalism and externalism on the market.

¹⁹ Whether the eminently epistemic value is fundamental, or it reduces to some other kind of value, is matter of debate. At any rate, that we can identify some value which is characteristic of epistemic states or inquiry is at least intuitively plausible.

have a positive moral or practical value. Believing in God may help one live a happier life, or render one better disposed to help and respect others; believing that she will win, notwithstanding the evidence that she will not, may significantly boost the athlete's performance (Rinard 2018).

Epistemic value and moral or practical value, then, at least in some circumstances, and at least *prima facie*, seem to come apart. Regardless of whether they are also *ultima facie* distinct, it is important to keep in mind that what we are concerned with here is the value which is (putatively) distinctive of epistemic practice, i.e. *epistemic value*.

6. Epistemic evaluability and epistemic significance

How does epistemic evaluability help us figure out which phenomena are epistemically significant and which ones are not? I see at least two ways to go here.

First, we may take *being epistemically evaluable* to be a mark of epistemic significance. Accordingly, if a phenomenon is susceptible to being evaluated with respect to an epistemic standard, then there are good reasons for thinking that that phenomenon is epistemically significant. This view accommodates our intuitions about epistemically significant phenomena such as propositional knowledge, knowledge of a subject matter, and understanding (in its different forms). It also accommodates the intuitions of those who think that true belief and justified belief are epistemically significant too. All these phenomena, arguably, are epistemically evaluable—after all, a huge portion of meta-epistemological literature revolves around these phenomena's epistemic value(s) (e.g. whether they differ in epistemic value and, if so, whether the difference is merely quantitative or in kind). Indeed, if there is to be epistemic evaluability at all, it will concern, at the very least, the phenomena just listed.

Perhaps surprisingly, however, epistemic evaluability is also displayed by some phenomena which some may not consider right off the bat as epistemically significant. As noted in §5, mere *belief* may be the target of epistemic evaluation: we can and indeed do evaluate beliefs with respect to some epistemic standard(s). A belief may be deemed as good or bad depending on whether it is true or false, justified or unjustified, a piece of knowledge or not, and so on. *Doubt* may be epistemically evaluable too, as well as *suspicion*. You may deem your friend's doubt about his partner's faithfulness reasonable or unreasonable, your colleague's suspicion that your boss is going to fire her as rational or irrational. Doubt and suspicion can thus be evaluated with

respect to whether they are justified, reasonable, or rational. They can also be evaluated with respect to whether they are or tend to be truth-conducive, as when, during a crime investigation, the police officer comes to doubt that the butler is really guilty or suspect that the victim's wife is implicated in her husband's murder.

Taking epistemic evaluability as a mark of epistemic significance, thus, has the consequence that mental states such as mere belief, doubt, and suspicion can be considered as epistemically significant. I do *not* argue that this consequence is bad—I am just highlighting the fact that this characterization of the mark *does* have that consequence. *If* we want to avoid this consequence, we need to choose a different mark.

Rather than being epistemically *evaluabile*, we might want to explore the idea that being epistemically *valuable* is a mark of epistemic significance instead. On this view, the mark is not constituted by the property of being a potential object of epistemic evaluation, but rather by the property of *being the bearer of epistemic value*. In this framework, mere belief is *not* to be considered as epistemically significant, for *mere* belief is not in itself epistemically valuable (at least on most conceptions of epistemic value): only duly qualified beliefs, such as true or justified ones, are.²⁰ By the same token, mere doubt and mere suspicion (as psychological states of a subject involving attribution of a degree of credence to a certain proposition), are not epistemically valuable, although properly qualified ones (e.g. reasonable or justified doubt or suspicion) may be.

This second way to frame the epistemic evaluability mark, on the other hand, keeps our intuitions about plausible epistemically significant phenomena firm. *Propositional knowledge*, of course, is epistemically valuable (if it is not, what else would be?). If truth is (one of) our epistemic value(s), then *true belief* is also to be considered as epistemically significant in this framework. As long as justification is truth-conducive, *justified belief* is also epistemically (instrumentally) valuable (if justification is itself taken to be an epistemic value, justified belief may also be considered as epistemically *finally* valuable). *Understanding* is (quite obviously) epistemically valuable as well. For one thing, as noted, understanding (in all its forms—understanding-that, understanding-why, and objectual understanding) involves true belief(s). For another thing, as some have argued, understanding may also have a *sui*

²⁰ Some beliefs may be valuable even if false or unjustified (e.g. the believer's belief in God, or the athlete's belief that she will win). However, first, their value would depend *not* on their being *beliefs*, but rather on their specific *content*, and, second, their value would perhaps be practical or moral, but certainly not eminently *epistemic*.

generis epistemic value (Kvanvig 2003). Similar considerations apply to knowledge of a subject matter.

Veridical *perception* and *introspection*, arguably, are epistemically valuable too. At the very least, they may be considered as *instrumentally* epistemically valuable, as long as they tend to produce or be evidence for the formation of true beliefs (on the assumption that truth is an epistemic value).

Knowledge-how is, here too, the trickiest case. On the one hand, it does not straightforwardly seem to be truth-directed, truth-conducive, or evaluable in terms of rationality. Knowing how to play *Suzanne* on the guitar, or knowing how to make a proper pizza, are more likely to be evaluated in terms of performance success (successfully playing the song) or specific-aim achievement (producing a good pizza), rather than in terms of truth, truth-conduciveness, or rationality (what is rational or irrational about knowing how to make a proper pizza?). On the other hand, it may be argued that, even if it does not *reduce* to propositional knowledge, knowledge-how may nonetheless include or presuppose some pieces of propositional knowledge or understanding. Even if knowing how to make a proper pizza is mostly a matter of practical ability and irreducible to any set of propositional knowledge, it is still reasonable to think that it nonetheless involves some pieces of propositional knowledge, say, concerning the exact quantity of flour and water required by the recipe. Similarly, knowing how to play *Suzanne* may partly involve some understanding of music, that is, a certain amount of objectual understanding. If this is right, then knowledge-how possesses some epistemic value, namely, at the very least, the epistemic value it inherits from its propositional-knowledge or understanding component.

This is far from settling the question about the putative epistemic value of knowledge-how, though, and indeed the reader is fully justified in keeping a skeptical stance about this matter notwithstanding the previous paragraph's considerations. To be sure, settling this question is not my purpose (the fate of knowing-how is not my present concern after all and my argument would be untouched by its turning out not to be epistemically significant). I will nonetheless make a last remark for sake of completeness. Knowing how to do things is admirable: we praise the musician who knows how to play all sorts of complex pieces on the guitar, or the pizzaiolo who knows how to make a good Napoli-style pizza. And this, arguably, not only because of the moral, practical, or aesthetical value that knowing-how produces. Of course, we appreciate the musician's knowledge-how partly because of the aesthetic experience it produces in us listeners. Yet, it seems that we also appreciate their knowledge-how, the ability or skill they display, independently of the effects such knowledge has on us. If this is right, then

knowledge-how has a value that exceeds the moral, practical, or aesthetic value. To see the point more crisply, consider the feeling of admiration you may sometimes have had while reading of a perfectly-performed robbery, notwithstanding the moral repulsion that may also have produced in you.²¹ The robbers' knowledge-how, although morally *disvaluable*, seems to possess, at the same time, some *positive* value—something that makes you deem their knowledge-how as *good* in a certain respect. If this is right, then knowledge-how possesses a value that (at least *prima facie*) exceeds any moral, practical, or aesthetic value it may also have. Knowledge-how's displaying a value that (*prima facie*) exceeds other values (moral, practical, and aesthetic) does not entail that such a value is eminently epistemic—it may be, for example, a *sui generis*, non-epistemic kind of value. However, that it *is* eminently epistemic is a legitimate hypothesis, yet one which, unfortunately, I cannot explore further here.

Let us sum this up. The epistemic realm displays an evaluative and a normative dimension, which promise to be the ground for a potential mark of epistemic significance. In this section, I have considered two possible ways to go: first, characterizing the relevant mark in terms of epistemic *evaluability*; second, characterizing it in terms of possession of epistemic *value*. Both features—*being epistemically evaluable* and *being the bearer of epistemic value*—are possessed by almost all potential bearers of epistemic significance listed in §1.1 (the only exception being, perhaps, knowledge-how). The former is also possessed by some phenomena (e.g. belief, doubt, and suspicion) which are not included in our preliminary list. Whether this constitutes a problem for the epistemic-evaluability characterization of the mark (thereby pushing for the epistemic-value characterization) exceeds our present concern. Either way, we have reason to think that epistemic evaluability or epistemic value is a mark of epistemic significance.

7. Epistemic evaluability and knowledge by acquaintance

In the previous section I have argued that either epistemic evaluability or epistemic value is a mark of epistemic significance. In this section, I am going to argue that knowledge by acquaintance displays both features: it is both epistemically *evaluable* (§7.1) and epistemically *valuable* (§7.2).

²¹ I am adapting here an example suggested to me by Uriah Kriegel in personal communication.

7.1. Knowledge by acquaintance is epistemically evaluable

We have seen in §5 that an epistemic value sets a standard of evaluation that constitutes (one of) our epistemic goal(s). The epistemic norm, on the other hand, defines how epistemic inquiry ought to be conducted. If, for instance, we take belief as our object of epistemic evaluation, we may have truth as our epistemic value and justification as our epistemic norm. Accordingly, *true* beliefs are *good* beliefs (evaluative dimension) and *justified* beliefs are *right* beliefs (normative dimension).²²

Now, primitive introspection is certainly not evaluable in terms of truth—given its non-classificatory nature, it cannot aim at truth. If you think that truth is the *only* possible epistemic value (this view is called ‘veritism’; see e.g. Goldman 2015), then the epistemic evaluability of primitive introspection seems to be a nonstarter. However, there may be reasons not to think so. Although primitive introspection is not evaluable in terms of truth (it cannot itself be true or false), it may be evaluated in terms of *truth-conduciveness*. If veritism is true (that is, if truth is the only epistemic value), then primitive introspection cannot be *intrinsically* valuable. However, it can still be *instrumentally* valuable. Consider perceptual experience. Although it may not be itself evaluable in terms of truth or falseness (if, suppose, its content is nonconceptual and non-propositional), it may nonetheless be epistemically evaluated in terms of its constituting an adequate ground for the acquisition of true perceptual beliefs or perceptual knowledge. Similarly, primitive introspection (by its nature nonconceptual and non-propositional) may be evaluated in terms of its constituting an adequate ground for the acquisition of true introspective beliefs or introspective knowledge. This, of course, is much easier said than proved: arguably, such an argumentative line would have to deal (at the very least) with long-standing myth-of-the-Given kind of objections.

If, however, you are open to considering a wider range of potential epistemic values (epistemic values other than truth), a (perhaps) less insidious route suggests itself, one that may lead to the result that primitive introspection has (also) an *intrinsic* epistemic value. There are at least two potential epistemic values in the vicinity of *truth*. The first is *accuracy*, intended here in terms of *absence of reference failure*. On the somewhat technical notion of accuracy at work here, an epistemic inquiry directed to *s* (where *s* may be a state of affairs, an event, a particular, and so on) is accurate iff *s* is present. Accuracy may apply both to classificatory and non-

²² This, of course, is just one possible way to spell out the relationship between the evaluative and the normative dimension in the epistemic domain.

classificatory epistemic states: your belief that the tablecloth is scarlet is accurate iff a tablecloth is present; your perceptual experience as of a scarlet tablecloth is accurate iff a tablecloth is present before you. Accuracy is entailed by truth: if your belief that the tablecloth is scarlet is true, then it is also accurate (it cannot be true if a tablecloth is not present).

The second potential epistemic value I want to consider is what may be called *exhaustiveness* or *comprehensiveness*. An epistemic inquiry concerning *s* is *exhaustive* iff it provides you with all the available information about *s*. That *exhaustiveness* is an epistemic value sounds intuitively true: an exhaustive epistemic inquiry (i.e. one which achieves the acquisition of all available information about its subject matter) is better than a non-exhaustive one. That *exhaustiveness* is irreducible to *truth* is less straightforward. For you may think that an exhaustive epistemic inquiry is just one which achieves all available truths about its subject matter. However, if you agree with me that not all information is propositional in form—as articulated in §4.2—then exhaustiveness may not reduce to truth, for there may be some pieces of information that are not truth bearers.

Classificatory epistemic states are evaluable in terms of exhaustiveness. Your belief that the tablecloth before you is *scarlet* is more exhaustive than your belief that the tablecloth before you is *colored*: it carries more information. I will argue that non-classificatory epistemic states are also evaluable in terms of exhaustiveness—I am going to show how it is so in due course.

Primitive introspection surely possesses the first of the two epistemic values outlined above, namely, *accuracy*. Given that there cannot be hallucination in primitive introspection, it cannot be the case that you have a state of primitive introspection directed to a certain phenomenal state, but no phenomenal state be there. However, that it is *evaluable* in terms accuracy is not straightforward. As just noted, primitive introspection *cannot* be inaccurate: necessarily, it is accurate. (Recall, the no-introspective-hallucination assumption I made in Chapter 4 is a necessity claim: there cannot be hallucination in primitive introspection.) On this assumption, how can primitive introspection be *evaluable* in term of accuracy? Compare: if all beliefs were true, how could we *evaluate* them in terms of truth? They would perhaps still *possess* the epistemic *value* of being true,²³ but they would not be *evaluable* in terms of truth.

²³ This might not be obvious: if a property is such that all the entities susceptible of possessing that property do in fact possess that property, can that property still constitute a value? At

One may reply that, in the latter case, non-evaluability of beliefs in terms of truth would be due not to the fact that all beliefs are true, but to the conjunction of this fact and the specific nature of the value of truth. For truth is an all-or-nothing property—there are no degrees of truth. If we consider a *gradable* property, instead, things seem to change. Let us assume that justification is an epistemic value. Even if all beliefs were justified, they would still be evaluable in terms of justification. For justification is a gradable property and beliefs can thereby be evaluated by their degree of justification. Some beliefs would still be more justified, and therefore better, than others.

This consideration, however, is of no help when it comes to primitive introspection and accuracy. Like truth, accuracy (in the technical sense at work here) is not a gradable, but an on-off property: something is either present or not—there are no *degrees of presence*. Accordingly, you cannot *compare* states of primitive introspection in terms of their accuracy: if all states of primitive introspection are accurate, they cannot be evaluated in terms of accuracy.

A further consideration may be invoked in defense of primitive introspection's evaluability in terms of accuracy, though. The (apparent) non-evaluability is due to the fact that all states of primitive introspection are accurate. This because any state of primitive introspection is by its nature accurate: necessarily, any state of primitive introspection is accurate. Now, imagine you believe that God exists. On your view, God is good, and He is good by His nature: necessarily, God is good. However, this does not prevent you from valuing Him as good: that is, it does not prevent you from *evaluating* Him in terms of goodness. It is not clear whether this analogy goes through all the way, though. For one thing, *valuing* something or someone is not the same as *evaluating* them: the former involves recognizing something as the bearer of a certain value, the latter involves attributing a value or a value-degree.

Regardless of how the issue concerning the relationship between accuracy and the epistemic evaluability of primitive introspection is to be settled, there is *another* epistemic value such that primitive introspection is evaluable in its terms, namely, *exhaustiveness*. Before explaining how it is so, a clarification is in order.

In §4.4, I argued that primitively introspecting a phenomenal state φ provides you with a complete and perfect grasp of φ 's phenomenology. Getting

any rate, it is plausible to think that even if all beliefs were true, there could still be false truth-apt entities (entities different from belief).

a *complete grasp* of φ , recall, entails acquiring all the information φ generates. It therefore entails *exhaustiveness*. I now want to qualify the claim I made in §4.4. Although all states of primitive introspection provide you with a *perfect grasp* of their target phenomenal state (they are therefore all accurate), not all of them provide you with a *complete grasp*—not all of them are exhaustive. States of primitive introspection may be *more or less* exhaustive depending on the amount of information they carry. Now, recall, all states of primitive introspection require that attention be directed toward the target phenomenal state. Attention, however, is a *gradable* phenomenon: a subject may devote more or less attentional resources to a given object. In the case of primitive introspection, one may devote more or less attentional resources to the target phenomenal state. The more attentional resources are deployed, the more information is acquired; the less attentional resources, the less information. Therefore, the more attentional resources are deployed in a given state of primitive introspection, the more that primitive introspective state is exhaustive. A state of primitive introspection directed at φ provides you with a complete grasp of φ 's phenomenology—that is, it is (fully) *exhaustive*—only if you attend to φ *fully* (only if you give to φ *full attention*), that is, only when *all your attentional resources* are directed toward φ .

Now, the idea is that if *exhaustiveness* is (one of) our epistemic *goal(s)* in primitive introspection, then *attention* is what gives primitive introspection its epistemic *norm*. Accordingly, we have that:

- a state of primitive introspection directed toward φ is *good* only if it provides you with a complete grasp of φ 's phenomenology, that is, only if it is *exhaustive*;
- a state of primitive introspection ι_1 directed toward φ is *better than* a state of primitive introspection ι_2 directed toward φ if ι_1 provides you with more information about φ 's phenomenology than ι_2 , that is, if ι_1 is *more exhaustive than* ι_2 ;
- a state of primitive introspection directed toward φ *ought to involve full attention* toward φ .

The first two points concern the *evaluative* dimension of primitive introspection, the last point concerns its *normative* dimension.

Therefore, primitive introspection is epistemically evaluable. It is evaluable in terms of exhaustiveness and its epistemic norm is determined by attentional resources deployment.

7.2. Knowledge by acquaintance is epistemically valuable

If you accept the previous subsection's suggestion that *accuracy* and *exhaustiveness* are epistemic values, as well as the claim I made in §4.4 that primitive introspection involves a perfect and (under the right conditions) complete grasp of its target's phenomenology, then you should also accept that knowledge by acquaintance is epistemically valuable. Indeed, at the very least, if you accept the very existence of primitive introspection, you should also accept that it is accurate (in the technical sense explained above), for, by its nature, primitive introspection cannot be hallucinatory. If you agree that accuracy is an epistemic value (one which, besides, is *entailed* by truth), then you should also believe that primitive introspection, and thereby knowledge by acquaintance, is epistemically valuable.

*

In this section, I have argued that primitive introspection displays the second mark of epistemic significance, namely the mark concerning epistemically significant phenomena's epistemic evaluability. I have pointed out that the mark may be characterized in two different ways: *being epistemically evaluable*, on the one hand, and *being epistemically valuable*, on the other hand. I have argued that primitive introspection displays such a mark either way.

Conclusion

I hope to have shown that information acquisition and epistemic evaluability (being epistemically evaluable or being epistemically valuable) can be reasonably considered as marks of epistemic significance. There seem to be *prima facie* reasons (reasons coming from intuitions) for thinking that information acquisition and epistemic evaluability are individually sufficient for epistemic significance. However, I haven't provided *conclusive* reasons in favor of this strong thesis. For some may think that information acquisition on the one hand, and epistemic evaluability on the other hand, must be supplemented with a further condition. Accordingly, we might say that, at the very least, information acquisition and epistemic evaluability are jointly sufficient for epistemic significance. For those who think that even this thesis is too strong, I am ready to retreat to the thesis that information acquisition and epistemic evaluability are, *jointly*, a mark of epistemic significance.

Arguably, something's displaying the *conjunction* of information acquisition and epistemic evaluability gives us *very strong* reasons for thinking that such a thing is epistemically significant.

I have argued that primitive introspection possesses both features. On the one hand, it involves information acquisition: at the very least, primitively introspecting a phenomenal state ϕ enables one to rule out the possible worlds in which $\neg\phi$. On the other hand, it is both epistemically evaluable and epistemically valuable: it is epistemically evaluable in terms of accuracy and exhaustiveness and it is epistemically valuable in the sense that it *is* accurate and (if full attention is devoted to the target phenomenal state) exhaustive.

Since it displays both marks of epistemic significance, primitive introspection is epistemically significant. Moreover, I have argued, it has a *special* epistemic significance: it possesses a distinctive epistemic property which I have called *complete and perfect grasp* of the phenomenology of the target phenomenal state: it provides the subject with all and only the information about the target phenomenal state's phenomenology which is generated by the occurrence of such phenomenal state.

CONCLUSION

In this dissertation, I have investigated an introspective phenomenon that I have called *primitive introspection*. I have characterized primitive introspection as non-classificatory introspection of phenomenal states, i.e. as a distinctively first-personal method through which one can non-descriptively acquire knowledge of the phenomenology of one's current conscious experience, where by 'non-descriptively' I mean without recognizing the experience as an instance of a previously encountered experience type.

I have defended three theses about primitive introspection concerning, respectively, its *existence* (Part I), *nature* (Part II), and *epistemology* (Part III). As for its *existence*, I have argued that primitive introspection is a psychologically real phenomenon. First, I have provided some preliminary motivation from ordinary-life cases for the idea that primitive introspection does feature in our inner lives (Chapter 1). Second, I have developed a sustained argument—the argument from phenomenal-concept acquisition—for the existence of primitive introspection (Chapter 2): rejecting the existence of primitive introspection entails a highly implausible version of nativism about phenomenal concepts. As for the *nature* of primitive introspection, I have developed a new version of the acquaintance account, what I have called the *integration account* of the metaphysics of primitive introspection. First, I have laid out the general structure of the primitive introspective process by distinguishing act, target, and state of primitive introspection; I have characterized primitive introspection in terms of attentive apprehension and elucidated the relevant notions of attention and apprehension involved (Chapter 3). Second, I have developed my own account of the metaphysics of primitive introspection, the *integration account*, which, differently from extant acquaintance accounts, fulfills what I called the *phenomenal modification desideratum*: it explains why some change in phenomenology occurs at every instance of primitive introspection and how primitive introspection is nonetheless possible notwithstanding such a phenomenal modification (Chapter 4). As for the *epistemology* of primitive introspection, I have argued that knowledge by acquaintance, i.e. knowledge that is *constituted* by primitive introspection, is epistemically significant even though it is irreducible to propositional knowledge. First, I have argued that knowledge by acquaintance cannot be reduced to propositional knowledge or to any other kind of knowledge (Chapter 5). Second, I have argued that, given that knowledge by acquaintance displays what I take to be two marks of epistemic significance (i.e. information acquisition and epistemic evaluability), there are

good reasons for thinking that knowledge by acquaintance is epistemically significant (Chapter 6). Moreover, primitive introspection possesses a special epistemic property: it provides the subject with a complete and perfect grasp of the phenomenology of their current experience.

Here are what I take to be the main contributions of this work. First, I have introduced a notion of introspection that is fundamentally different from the judgment-like, more reflective kind of introspection which is the topic of many philosophical debates about self-knowledge. To be sure (as already noted), I am in good company. Other philosophers, especially some contemporary acquaintance theorists, have investigated introspective phenomena that are very close to what I call ‘primitive introspection’: most of them have argued that there are special introspective beliefs which possess some of the features I attribute to primitive introspection. Notwithstanding the similarities, there is an important sense in which my notion is fundamentally different and perhaps a bit more radical than theirs. For, differently from those special introspective beliefs, primitive introspection is nonconceptual and non-propositional, therefore *not* belief-like.

Second, I have proposed a new version of the acquaintance account, one that promises to solve the tension between the possibility of introspective acquaintance and phenomenal modification associated with introspective attention. As I have insisted, the integration account may not be the sole satisfactory account of the metaphysics of primitive introspection. If anything, I have tried to pinpoint the requirements an account of the metaphysics of primitive introspection must satisfy and put the phenomenal modification issue back on acquaintance theorists’ agenda.

Third, I have attempted to resurrect the now largely discarded Russellian idea that knowledge by acquaintance is a *sui generis* kind of knowledge. I have tried to distill the gist of the Russellian idea by arguing that knowledge by acquaintance, although irreducible to propositional knowledge, is nonetheless epistemically significant. This thesis might be more digestible to those who would reject the Russellian thesis a priori—i.e. those on whose view knowledge is by definition propositional.

Finally, against the current fallibilist orthodoxy and the growing introspective skepticism, I have argued for a quasi-Cartesian thesis about the epistemology of primitive introspection—i.e. that primitive introspection provides the subject with a complete and perfect grasp of their experience. Although the critical stance about introspection has insightfully identified Cartesianism’s overstatements, finer-grained distinctions should be drawn not to throw out the baby with the bath water. For there is nonetheless a grain

of truth in the Cartesian intuition, one which should not be neglected. By restricting the scope of epistemic privilege to primitive introspection, I have tried to vindicate the Cartesian idea, while acknowledging the limits of introspection highlighted in recent critical literature.

I believe that primitive introspection is of non-negligible philosophical significance and that taking it into account could help shed light on various philosophical issues. In the last few paragraphs, I sketch some further developments and potential philosophical applications of the ideas developed in the present work.

I have argued that primitive introspection is prior to the formation of the special introspective beliefs (e.g. direct phenomenal beliefs) to which some contemporary acquaintance theorists attribute epistemic privilege (e.g. infallibility). Further work needs to be done, however, to spell out how primitive introspection grounds or enables the formation of such beliefs. More specifically, a thorough investigation of the psychological, metaphysical, and epistemic relationship between primitive introspection and special introspective beliefs might shed light on whether and how the specialness and epistemic privilege of these beliefs is grounded in or flows out from the specialness and epistemic privilege of primitive introspection.

A related topic deserving further exploration is the relationship between primitive introspection and what I have called *reflective introspection*, that is, the kind of introspection of phenomenal states that does involve classifying what is introspected as an instance of a previously encountered experience type. The envisaged study of the relationship between primitive introspection and special introspective beliefs, mentioned in the previous paragraph, may be a first pass at an analysis of the relationship between primitive introspection and reflective introspection. On the one hand, there is one respect in which those special introspective beliefs are significantly similar to primitive introspection and fundamentally different from reflective introspection: they do not involve recognition or classification. On the other hand, there is another feature of those beliefs that makes them essentially different from primitive introspection and structurally similar to reflective introspection: they have a propositional content and involve the deployment of some concepts. If those beliefs really exist, they may be considered as a third kind of phenomenal-state introspection, besides primitive and reflective introspection. Given their ambivalent nature, those beliefs may constitute the 'bridge' between primitive introspection and reflective introspection. The result would be a sort of foundationalist picture of phenomenal-state introspection, one on which reflective introspection is grounded in special introspective beliefs, which are grounded in primitive introspection. If so,

primitive introspection would be the most fundamental kind of introspection of phenomenal states.

As noted in my introductory discussion of primitive introspection, I have not taken a stance on the exact scope of primitive introspection: throughout the dissertation, I have remained neutral about which conscious states have phenomenology, and thus about which conscious states are the potential target of primitive introspection. Although some conscious states are almost uncontroversially phenomenal (e.g. perceptual and bodily states), there is disagreement about whether, for instance, *cognitive* states like thoughts and *conative* states like desires have phenomenology. Further research could be devoted to developing a more committal account of the *scope* of primitive introspection, partly by engaging in current debates about the nature and phenomenology of different kinds of mental states (cognitive states, emotions, moods, imagination, and so on). One possible upshot is that, if it turns out that all conscious states have phenomenology, this will potentially constitute a strong motivation for the claim that primitive introspection provides a substantial ground of *all* introspective self-knowledge.

Indeed, whether and how primitive introspection constitutes a ground for self-knowledge is a further issue worth delving into. For one thing, the role of phenomenal-state introspection in the acquisition of self-knowledge needs to be clarified. It might turn out that, independently of whether cognitive and conative states have phenomenology, introspection of the phenomenology of our experiences constitutes one of the bases on which knowledge of our (standing or occurrent) beliefs, desires, hopes, and intentions is formed. More generally, it would be interesting to explore how phenomenal-state introspection interacts and cooperates with other knowledge-acquisition methods or processes (memory, testimony, behavior observation, inference, reasoning, and so on) in the achievement of self-knowledge. If primitive introspection is the most fundamental kind of phenomenal-state introspection and phenomenal-state introspection is one of the grounds of self-knowledge, then primitive introspection is itself a ground of self-knowledge. Indeed, my hypothesis is that, although not the *sole* ground, primitive introspection is a fundamental or necessary ground of self-knowledge. Not only it is what self-knowledge is partly grounded in: without primitive introspection self-knowledge would not be possible. In other words, on this picture, primitive introspection is a necessary, though not sufficient, ground of self-knowledge. If so, then primitive introspection possesses not only the *intrinsic* epistemic value for which I argued in Chapter 6 (it is epistemically valuable *in itself* because it *constitutes* knowledge by

acquaintance, which is intrinsically epistemically significant—i.e. it is epistemically significant but not in virtue of making something else epistemically significant), but also an *instrumental* epistemic value: it *grounds* and contributes to the acquisition of (propositional) self-knowledge.

The distinction between primitive introspection and reflective introspection may help assess skepticism about the justification of introspective beliefs. As mentioned, it is sometimes argued that, since we are very often mistaken or uncertain about the phenomenology of our experiences, introspection is unreliable. Skeptical arguments rely the following reasoning: introspective beliefs are beliefs that are grounded in introspection; but introspection is unreliable; therefore, introspective beliefs are not reliably formed and thus not justified. The distinction between primitive and reflective introspection may shed light on the scope of this kind of skeptical arguments. In particular, it may help draw a diagnosis of why and how error or uncertainty occurs. Very roughly, the idea would be that error and uncertainty, if they occur, are not due to a flaw in primitive introspection, but rather to misclassification, or misapplication of concepts on the part of the subject. If so, then the real target of skeptical arguments is only reflective introspection: the reliability of primitive introspection is left untouched.

Moreover, suppose error and uncertainty really are as widespread as some skeptics imply. Then if error and uncertainty are due to misapplication of concepts, rather than to a flaw in primitive introspection, this leaves open the possibility of their being *reduced*. The idea is that error and uncertainty *can* be avoided if our categorizing abilities are *improved*. Consider the following analogy. If the visual system of a subject does not work correctly, then not only the subject *does not*, but they also *cannot*, form correct visual judgements. On the other hand, it might occur that a subject typically forms wrong visual judgements not because of a malfunction in their visual system, but because they fail to correctly apply visual concepts to what they (reliably) see. In the latter case error can be corrected without correcting the subject's visual apparatus. If the subject had learned how to correctly apply visual concepts, they would be in a position to typically form correct visual judgements. Analogously, in the case of introspection, a subject may not typically form correct judgements about their conscious experience either because the process of primitive introspection is unreliable, or because their ability to recognize and classify the type of experience they are primitively introspecting is not sufficiently developed. Whereas in the former case the subject's introspective judgements are difficult to correct, in the latter case they are corrigible upon training, i.e. if the subject learns how to correctly apply phenomenal concepts. In other words, the idea is that the cognitive

process underlying the formation of introspective beliefs is such that, although it may be flawed in naïve introspectors, it can be *improved* and thereby become much more reliable in expert introspectors.

Further work may be devoted to the metaphysics of primitive introspection. As noted, the integration account developed in this dissertation may not be the sole satisfactory account of the nature of primitive introspection: other accounts may be explored. One alternative was sketched in Chapter 4's *Appendix 2*. That option is still underdeveloped, though, and would deserve further scrutiny. Another option that comes to my mind is a sort of idiosyncratic hybrid between the inner sense account and the acquaintance account, what may be called the *naïve realist account of primitive introspection*. The idea would be to explain the nature of primitive introspection on the model of perceptual experience, as the inner sense view does, but by presupposing a *naïve realist* view of perceptual experience, according to which—very roughly—perceptual experiences are partly constituted by the objects perceived. Naïve realism thus explains perception in terms of constitution, rather than in terms of representation and causation. In a sense, thus, the naïve realist version of the inner sense account of primitive introspection may somewhat collapse into the acquaintance account. Yet, developing an acquaintance account from this particular angle might still be useful to elucidate the nature and epistemology of primitive introspection. While inheriting its virtues, naïve realism about introspection may avoid the difficulties of its perceptual counterpart. Exploring the naïve realist account of introspection may help shed light on the aspects of the acquaintance account which may strike some as mysterious.

I have argued that primitive introspection possesses an eminently *epistemic* value. But does it also possess other kinds of value? One important question to be explored, for instance, is whether it has any *moral* or *practical* value. Can primitive introspection contribute to our living a *better* life? If, for instance, as suggested in Chapter 1, primitive introspection is involved in some types of meditation practice, it may have the value associated with the self-improvement such meditation practices bring-about. Or consider the phenomenal modification that has been topic of long discussion in Chapter 4. Primitively introspecting modifies the phenomenology of the target phenomenal state: it always modifies its salience, but it may also modify other phenomenal aspects, such as intensity and clarity. If one learns how primitive introspection may affect the phenomenology of one's experience, one may exploit this knowledge to actively alter, transform, or shape one's experience

by primitively introspecting it.¹ If drawing your attention toward your anger experience changes it in the way that so much worried Brentano in the passage quoted in Chapter 4, knowing this may be for you a useful tool for *controlling* your anger. A sub-question is whether, if primitive introspection does have moral or practical value, it possesses it merely instrumentally (e.g. it helps us live a better life, or become better persons), or also intrinsically (it is good in itself). Moreover, one may wonder whether primitive introspection also has an *aesthetic* value: is primitively introspecting something we value (also) from an aesthetic point of view? In Chapter 1 I suggested that, beside meditation, there may be other cases in which, although one *can* recognize the experience one is introspecting, one *chooses not to* exercise the relevant recognitional ability and thereby introspect one's experience without forming any judgment about it. One possible reason one might want to so introspect is *aesthetic*: one may contemplate and enjoy one's experience better by attending to it in a non-classificatory way.

While arguing that primitive introspection is epistemically significant, I have presented what I take to be two marks of epistemic significance: information acquisition and epistemic evaluability. The notion of epistemic significance might have been underexplored and could deserve further scrutiny. Particularly, further research could be devoted to looking for other potential marks of epistemic significance, and to maximally clarifying what makes some state, process, or phenomenon *epistemic*.

A quite unrelated question concerns the relationship between primitive introspection and what has been sometimes called *non-reflective*, or *pre-reflective*, or *peripheral* inner awareness. The latter is characterized as a kind of awareness that, like primitive introspection, is directed toward one's conscious mental states and does not involve any classification of what it is directed at. However, unlike primitive introspection, peripheral inner awareness does not involve drawing one's attention to the target conscious state. On the contrary, peripheral inner awareness is constitutively *non-attentive*. Primitive introspection and peripheral inner awareness may display interesting connections and it would be worth investigating the psychological, metaphysical, and epistemological relationship between the two phenomena.

Finally, there is a question concerning the relationship between primitive introspection and the metaphysics, phenomenology, and epistemology of the *self*. Primitive introspection provides one with knowledge of one's current phenomenal states. But does it thereby provide one with

¹ Thanks to Johannes Brandl for suggesting this practical application of primitive introspection to me.

knowledge of *oneself*? That is, by primitively introspecting, is one only acquainted with one's phenomenal state or, in addition, also with one's *self* (if there is any such thing)? Or is there any other way (different from acquaintance) one can come to know one's self by primitively introspecting? Although at the current stage of my research I remain neutral about the phenomenology, the ontology, and the epistemology of the self, I consider it urgent and important to develop an account which engages with these issues and explains the relationship primitive introspection bears to them. One hypothesis to explore might be that self-awareness (i.e. awareness of one's self) is grounded in the metaphysical structure of primitive introspection. Accordingly, if there is knowledge of one's self at all, it must depend on primitive introspection.

REFERENCES

- Adams, Fred. 2010. "Information and Knowledge À La Floridi." *Metaphilosophy* 41 (3): 331–44.
- Allport, Alan. 1993. "Attention and Control: Have We Been Asking the Wrong Questions? A Critical Review of Twenty-Five Years." In *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*, edited by David E. Meyer and Sylvan Kornblum, 183–218. Cambridge, MA: MIT Press.
- Armstrong, David M. 1963. "Is Introspective Knowledge Incorrigible?" *The Philosophical Review* 72 (4): 417–32.
- . 1968. *A Materialist Theory of the Mind*. New York: Humanities Press.
- Aydede, Murat. 2009. "Is Feeling Pain the Perception of Something?" *The Journal of Philosophy* 106 (10): 531–67.
- Aydede, Murat, and Matthew Fulkerson. 2014. "Affect: Representationalists' Headache." *Philosophical Studies* 170 (2): 175–198.
- Baker, Lynne Rudder. 2003. "Third Person Understanding." In *The Nature and Limits of Human Understanding*, edited by A. J. Sanford and P. N. Johnson-Laird, 185–208. London: T & T Clark.
- Balog, Katalin. 2012. "Acquaintance and the Mind-Body Problem." In *New Perspectives on Type Identity: The Mental and the Physical*, edited by Simone Gozzano and Christopher S. Hill, 16–42. Cambridge: Cambridge University Press.
- Bar-On, Dorit. 2004. *Speaking My Mind: Expression and Self-Knowledge*. Oxford: Oxford University Press.
- Bayne, Tim. 2010. *The Unity of Consciousness*. Oxford: Oxford University Press.
- Bayne, Tim, and David J. Chalmers. 2003. "What Is the Unity of Consciousness?" In *The Unity of Consciousness: Binding, Integration, and Dissociation*, edited by Axel Cleeremans, 23–58. Oxford: Oxford University Press.
- Bengson, John. 2015. "The Intellectual Given." *Mind* 124 (495): 707–60.
- Block, Ned. 1995a. "How Many Concepts of Consciousness?" *Behavioral and Brain Sciences* 18 (2): 272–87.
- . 1995b. "On a Confusion About a Function of Consciousness." *Behavioral and Brain Sciences* 18 (2): 227–247.
- . 1996. "Mental Paint and Mental Latex." *Philosophical Issues* 7: 19–49.
- Boghossian, Paul A., and J. David Velleman. 1989. "Colour as a Secondary Quality." *Mind* 98 (389): 81–103.

- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- . 2000. “Toward a Defense of Empirical Foundationalism.” In *Resurrecting Old-Fashioned Foundationalism*, edited by Michael R. DePaul. Lanham, Ma: Rowman and Littlefield.
- Bradley, Francis Herbert. 1886. “Is There Any Special Activity of Attention?” *Mind* 11 (43): 305–23.
- Brentano, Franz. 1874. *Psychology from an Empirical Standpoint*. Edited by Linda L. McAlister. Translated by Antos C. Rancurello, Dailey B. Terrell, and Linda L. McAlister. London; New York: Routledge.
- Brogaard, Berit. 2005. “I Know. Therefore, I Understand.” Unpublished manuscript. <https://philpapers.org/rec/BROIKT>.
- . , ed. 2014. *Does Perception Have Content?* Oxford, New York: Oxford University Press.
- Brown, Jessica A. 2013. “Knowing-How: Linguistics and Cognitive Science.” *Analysis* 73 (2): 220–27.
- Brownstein, Michael, and Eliot Michaelson. 2016. “Doing without Believing: Intellectualism, Knowledge-How, and Belief-Attribution.” *Synthese* 193 (9): 2815–36.
- Byrne, Alex. 2005. “Introspection.” *Philosophical Topics* 33 (1): 79–104.
- Campbell, John. 2002. *Reference and Consciousness*. Oxford: Clarendon Press.
- Carey, Susan. 2009. *The Origin of Concepts*. Oxford; New York: Oxford University Press.
- Carr, David. 1981. “Knowledge in Practice.” *American Philosophical Quarterly* 18 (1): 53–61.
- Carrasco, Marisa, Sam Ling, and Sarah Read. 2004. “Attention Alters Appearance.” *Nature Neuroscience* 7 (3): 308–13.
- Carruthers, Peter. 2005. *Consciousness: Essays from a Higher-Order Perspective*. Oxford: Oxford University Press.
- Carter, J. Adam, and Jesús Navarro. 2017. “The Defeasibility of Knowledge-How.” *Philosophy and Phenomenological Research* 95 (3): 662–85.
- Carter, J. Adam, and Duncan Pritchard. 2015a. “Knowledge-How and Cognitive Achievement.” *Philosophy and Phenomenological Research* 91 (1): 181–99.
- . 2015b. “Knowledge-How and Epistemic Luck.” *Noûs* 49 (3): 440–53.
- . 2015c. “Knowledge-How and Epistemic Value.” *Australasian Journal of Philosophy* 93 (4): 799–816.
- Cassam, Quassim. 2015. *Self-Knowledge for Humans*. Oxford, New York: Oxford University Press.

- Cath, Yuri. 2011. "Knowing How Without Knowing That." In *Knowing How: Essays on Knowledge, Mind, and Action*, edited by John Bengson and Mark Moffett, 113–35. Oxford: Oxford University Press.
- Chalmers, David J. 2003. "The Content and Epistemology of Phenomenal Belief." In *Consciousness: New Philosophical Perspectives*, edited by Quentin Smith and Aleksandar Jokic, 220–72. Oxford; New York: Oxford University Press.
- Chudnoff, Elijah. 2013. *Intuition*. Oxford: Oxford University Press.
- Clark, Michael. 1963. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." *Analysis* 24 (2): 46–48.
- Cohen, Stewart. 1984. "Justification and Truth." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 46 (3): 279–95.
- Conee, Earl. 1994. "Phenomenal Knowledge." *Australasian Journal of Philosophy* 72 (2).
- Crane, Tim. 2000. "Introspection, Intentionality, and the Transparency of Experience." *Philosophical Topics* 28 (2): 49–67.
- Dainton, Barry. 2000. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London; New York: Routledge.
- . 2014. "Temporal Consciousness." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2014. <http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/>.
- Davidson, Donald. 1969. "The Individuation of Events." In *Essays in Honor of Carl G. Hempel*, edited by Nicholas Rescher. Dordrecht: D. Reidel.
- . 1985. "Reply to Quine on Events." In *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, edited by Ernest Lepore and Brian McLaughlin. Oxford: Basil Blackwell.
- Dennett, Daniel C. 1988. "Quining Qualia." In *Consciousness in Contemporary Science*, edited by Anthony J. Marcel and Edoardo Bisiach. Oxford: Oxford University Press.
- DePoe, John M. 2018. "Knowledge by Acquaintance and Knowledge by Description." *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/knowacq/>.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, Mass: MIT Press.
- . 1993. "Conscious Experience." *Mind* 102 (406): 263–283.
- . 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Ducasse, Curt John. 1953. *A Philosophical Scrutiny of Religion*. New York: Ronald Press Co.

- Fantl, Jeremy. 2017. "Knowledge How." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2017. <https://plato.stanford.edu/archives/fall2017/entries/knowledge-how/>.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, N.J: Prentice Hall.
- Feldman, Richard, and Earl Conee. 2001. "Internalism Defended." *American Philosophical Quarterly* 38 (1): 1–18.
- Fine, Kit. 1994. "Essence and Modality: The Second Philosophical Perspectives Lecture." *Philosophical Perspectives* 8: 1–16.
- Floridi, Luciano. 2006. "The Logic of Being Informed." *Logique et Analyse* 49 (196): 433–60.
- Foley, Richard. 2012. *When Is True Belief Knowledge?* Princeton: Princeton University Press.
- Frankish, Keith. 2016. "Illusionism as a Theory of Consciousness." *Journal of Consciousness Studies* 23 (11–12): 11–39.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, Md.: Rowman & Littlefield.
- . 2009. "Luminous Enough for a Cognitive Home." *Philosophical Studies* 142 (1): 67–76.
- Gertler, Brie. 2001. "Introspecting Phenomenal States." *Philosophy and Phenomenological Research* 63 (2): 305–28.
- . 2011. *Self-Knowledge*. London; New York: Routledge.
- . 2012. "Renewed Acquaintance." In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar, 89–123. Oxford: Oxford University Press.
- Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–23.
- Giustina, Anna. 2017. "Conscious Unity from the Top Down: A Brentanian Approach." *The Monist* 100 (1): 15–36.
- Giustina, Anna, and Uriah Kriegel. 2017. "Fact-Introspection, Thing-Introspection, and Inner Awareness." *Review of Philosophy and Psychology* 8 (1): 143–64.
- Goldman, Alvin I. 1967. "A Causal Theory of Knowing." *The Journal of Philosophy* 64 (12): 357–72.
- . 1979. "'What Is Justified Belief?'" In *Justification and Knowledge*, edited by George Pappas, 1–23. Dordrecht, Holland: D. Reidel Publishing Company.
- . 2006. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.

- . 2015. “Reliabilism, Veritism, and Epistemic Consequentialism.” *Episteme* 12 (2): 131–143.
- Gordon, Emma C. 2018. “Understanding in Epistemology.” In *The Internet Encyclopedia of Philosophy*.
- Grimm, Stephen R. 2011. “Understanding.” In *The Routledge Companion to Epistemology*, edited by Sven Berneker and Duncan Pritchard. New York: Routledge.
- Hannon, Michael. 2014. “Is Knowledge True Belief Plus Adequate Information?” *Erkenntnis* 79 (5): 1069–76.
- Harman, Gilbert. 1990. “The Intrinsic Quality of Experience.” *Philosophical Perspectives* 4: 31–52.
- Hartland-Swann, John. 1956. “The Logical Status of ‘Knowing That.’” *Analysis* 16 (5): 111–15.
- Hasan, Ali, and Richard Fumerton. 2017. “Knowledge by Acquaintance vs. Description.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2017. <https://plato.stanford.edu/archives/fall2017/entries/knowledge-acquaindescrip/>.
- Hayner, Paul. 1969. “Knowledge by Acquaintance.” *Philosophy and Phenomenological Research* 29 (3): 423–31.
- Heck, Richard G. 2000. “Nonconceptual Content and the ‘Space of Reasons.’” *The Philosophical Review* 109 (4): 483–523.
- Hill, Christopher S. 1991. *Sensations: A Defense of Type Materialism*. Cambridge: Cambridge University Press.
- . 2009. *Consciousness*. Cambridge: Cambridge University Press.
- Horgan, Terry, and Uriah Kriegel. 2007. “Phenomenal Epistemology: What Is Consciousness That We May Know It so Well?” *Philosophical Issues* 17 (1): 123–144.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham, Ma: Rowman & Littlefield.
- Jackson, Frank. 1982. “Epiphenomenal Qualia.” *The Philosophical Quarterly* 32 (127): 127–36.
- James, William. 1890. *The Principles of Psychology*. London: Macmillan.
- Johnston, Mark. 1992. “How to Speak of the Colors.” *Philosophical Studies* 68 (3): 221–63.
- Kammerer, François. 2016. “The Hardest Aspect of the Illusion Problem--and How to Solve It.” *Journal of Consciousness Studies* 23 (11–12): 124–39.
- Kelly, Thomas. 2003. “Epistemic Rationality as Instrumental Rationality: A Critique.” *Philosophy and Phenomenological Research* 66 (3): 612–40.

- Kentridge, R. W., C. A. Heywood, and L. Weiskrantz. 1999. "Attention without Awareness in Blindsight." *Proceedings of the Royal Society of London B: Biological Sciences* 266 (1430): 1805–11.
- Kim, Jaegwon. 1988. "What Is 'Naturalized Epistemology?'" *Philosophical Perspectives* 2: 381–405.
- . 1993. *Supervenience and Mind: Selected Philosophical Essays*. Cambridge Studies in Philosophy. Cambridge: Cambridge University Press.
- Kind, Amy. 2003. "What's so Transparent about Transparency?" *Philosophical Studies* 115 (3): 225–44.
- . 2014. "The Case against Representationalism about Moods." In *Current Controversies in Philosophy of Mind*, edited by Uriah Kriegel, 113–34. London; New York: Routledge.
- Koethe, John. 2002. "Stanley and Williamson on Knowing How." *The Journal of Philosophy* 99 (6): 325–328.
- Kornblith, Hilary. 1998. "What Is It like to Be Me?" *Australasian Journal of Philosophy* 76 (1): 48–60.
- Kriegel, Uriah. 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford: Oxford University Press.
- . 2012. "Personal-Level Representation." *Protosociology* 28: 77–114.
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- . 2014. "Epistemic Normativity." In *Epistemic Norms: New Essays on Action, Belief, and Assertion*, edited by Clayton Littlejohn and John Turri, 115–34. Oxford: Oxford University Press.
- Lehrer, Keith, and Thomas Paxson. 1969. "Knowledge: Undefeated Justified True Belief." *The Journal of Philosophy* 66 (8): 225–37.
- Levin, Janet. 2000. "Dispositional Theories of Color and the Claims of Common Sense." *Philosophical Studies* 100 (2): 151–74.
- Lewis, David. 1990. "What Experience Teaches." In *Mind and Cognition*, edited by William G. Lycan, 499–519. Oxford: Blackwell.
- Loar, Brian. 2003. "Transparent Experience and the Availability of Qualia." In *Consciousness: New Philosophical Perspectives*, edited by Quentin Smith and Aleksandar Jokic, 77–96. Oxford; New York: Oxford University Press.
- Lycan, William G. 1996. *Consciousness and Experience*. Cambridge, MA: MIT Press.
- Mack, Arien, and Irvin Rock. 1998. *Inattentional Blindness*. Cambridge, MA: MIT Press.

- Massin, Olivier. 2013. "The Intentionality of Pleasures and Other Feelings, a Brentanian Approach." In *Themes from Brentano*, edited by Denis Fisette and Guillaume Fréchette, 307–337. Amsterdam: Rodopi.
- Millikan, Ruth Garrett. 2012. "Are There Mental Indexicals and Demonstratives?" *Philosophical Perspectives* 26 (1): 217–34.
- Mole, Christopher. 2008. "Attention and Consciousness." *Journal of Consciousness Studies* 15 (4): 86–104.
- . 2011a. *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford; New York: Oxford University Press.
- . 2011b. "The Metaphysics of Attention." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 60–77. Oxford; New York: Oxford University Press.
- Moran, Richard. 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton, NJ: Princeton University Press.
- Müller, Georg Elias. 1911. *Zur Analyse Der Gedächtnistätigkeit Und Des Vorstellungsverlaufes*. Vol. 1. Leipzig: J. A. Barth.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *The Philosophical Review* 83 (4): 435–50.
- Nanay, Bence, ed. 2010. *Perceiving the World*. Oxford: Oxford University Press.
- Nemirow, Laurence. 1990. "Physicalism and the Cognitive Role of Acquaintance." In *Mind and Cognition*, edited by William G. Lycan, 490–99. Oxford: Blackwell.
- Peacocke, Christopher. 1984. "Colour Concepts and Colour Experience." *Synthese* 58 (3): 365–81.
- . 2001. "Does Perception Have a Nonconceptual Content?" *The Journal of Philosophy* 98 (5): 239–64.
- Peterson, John. 2008. *Aquinas: A New Introduction*. Lanham, Md.: University Press of America.
- Petitmengin, Claire. 2006. "Describing One's Subjective Experience in the Second Person: An Interview Method for the Science of Consciousness." *Phenomenology and the Cognitive Sciences* 5 (3–4): 229–69.
- Pitt, David. 2004. "The Phenomenology of Cognition Or What Is It Like to Think That P?" *Philosophy and Phenomenological Research* 69 (1): 1–36.
- Prinz, Jesse. 2004. *Gut Reactions: A Perceptual Theory of Emotion*. Oxford: Oxford University Press.
- . 2011. "Is Attention Necessary and Sufficient for Consciousness?" In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 174–203. Oxford; New York: Oxford University Press.

- Pryor, James. 2005. "There Is Immediate Justification." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 181–202. Oxford: Blackwell.
- Ren, Huiming. 2012. "The Distinction between Knowledge-That and Knowledge-How." *Philosophia* 40 (4): 857–75.
- Rinard, Susanna. 2018. "Equal Treatment for Belief." *Philosophical Studies* (forthcoming).
- Rosenthal, David M. 1997. "A Theory of Consciousness." In *The Nature of Consciousness: Philosophical Debates*, edited by Ned Block, Owen J. Flanagan, and Güven Güzeldere, 729–53. Cambridge, MA: MIT Press.
- . 2005. *Consciousness and Mind*. Oxford: Clarendon Press.
- Roskies, Adina L. 2008. "A New Argument for Nonconceptual Content." *Philosophy and Phenomenological Research* 76 (3): 633–59.
- Rumfitt, Ian. 2003. "Savoir Faire." *The Journal of Philosophy* 100 (3): 158–66.
- Russell, Bertrand. 1910. "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society* 11: 108–28.
- . 1912. *The Problems of Philosophy*. New York: H. Holt and Company.
- Ryle, Gilbert. 1949. *The Concept of Mind*. Chicago: The University of Chicago Press.
- Schaffer, Jonathan. 2007. "Knowing the Answer." *Philosophy and Phenomenological Research* 75 (2): 383–403.
- Schneider, Susan. 2016. "Events." In *The Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/events/>.
- Schwitzgebel, Eric. 2008. "The Unreliability of Naive Introspection." *Philosophical Review* 117 (2): 245–273.
- . 2012. "Introspection, What?" In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar, 29–48. Oxford: Oxford University Press.
- Searle, John. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge; New York: Cambridge University Press.
- Shannon, Claude E. 1948. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27 (3): 379–423.
- Siewert, Charles. 2004. "Is Experience Transparent?" *Philosophical Studies* 117 (1–2): 15–41.
- . 2012. "On the Phenomenology of Introspection." In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar, 129. Oxford: Oxford University Press.
- Simons, Daniel J., and Christopher F. Chabris. 1999. "Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events." *Perception* 28 (9): 1059–74.

- Smith, Arthur D. 2008. "Translucent Experiences." *Philosophical Studies* 140 (2): 197–212.
- Smithies, Declan. 2011. "Attention Is Rational-Access Consciousness." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 247–73. Oxford; New York: Oxford University Press.
- Solomon, Robert C. 1976. *The Passions*. New York: Anchor Press/Doubleday.
- Stanley, Jason, and Timothy Williamson. 2001. "Knowing How." *The Journal of Philosophy* 98 (8): 411–44.
- Stazicker, James. 2011. "Attention, Visual Consciousness and Indeterminacy." *Mind and Language* 26 (2): 156–84.
- Titchener, Edward B. 1912. "The Schema of Introspection." *The American Journal of Psychology* 23 (4): 485–508.
- Travis, Charles. 2004. "The Silence of the Senses." *Mind* 113 (449): 57–94.
- Tye, Michael. 1995a. "Blindsight, Orgasm, and Representational Overlap." *Behavioral and Brain Sciences* 18 (2): 268–69.
- . 1995b. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.
- . 1997. "A Representational Theory of Pains and Their Phenomenal Character." In *The Nature of Consciousness: Philosophical Debates*, edited by Ned Block, Owen J. Flanagan, and Güven Güzeldere. Cambridge, MA: MIT Press.
- . 2003. *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.
- Watzl, Sebastian. 2011a. "Attention as Structuring of the Stream of Consciousness." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 145–73. Oxford; New York: Oxford University Press.
- . 2011b. "The Nature of Attention." *Philosophy Compass* 6 (11): 842–53.
- . 2017. *Structuring Mind. The Nature of Attention and How It Shapes Consciousness*. Oxford: Oxford University Press.
- White, Alan Richard. 1964. *Attention*. Oxford: Basil Blackwell.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski, Linda. 2001. "Recovering Understanding." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup. Oxford: Oxford University Press.
- . 2003. "The Search for the Source of Epistemic Good." *Metaphilosophy* 34 (1/2): 12–28.

RÉSUMÉ

Cette thèse se concentre sur un phénomène introspectif que j'appelle *introspection primitive*. L'introspection primitive est un type d'introspection d'états phénoménaux qui est non classificatoire : quand on introspecte primitivement, on ne reconnaît pas l'état phénoménaux introspecté comme un exemple de type d'expérience précédemment rencontrée. Je défends principalement trois thèses sur l'introspection primitive. Premièrement, elle *existe* : il y a un phénomène mental qui a les caractéristiques que j'attribue à l'introspection primitive et un tel phénomène est un processus introspectif réel. Deuxièmement, sa *nature* est plus adéquatement expliquée par une version de la théorie de l'*accointance* – ce que j'appelle la *théorie de l'intégration*. Troisièmement, elle a une *valeur épistémique* : elle permet au sujet d'acquérir une connaissance de ses états phénoménaux. Cette connaissance est un type de connaissance *sui generis* : la *connaissance par accointance*. La connaissance par accointance a une propriété épistémique spéciale qui est, pourrait-on dire, analogue à l'*infaillibilité* : elle donne au sujet un *saisi complet et parfait* de la phénoménologie de son expérience.

MOTS CLÉS

Connaissance de soi, Introspection, Concepts phénoménaux, Connaissance par accointance, Accointance, Non conceptuel.

ABSTRACT

This dissertation focuses on an introspective phenomenon that I call *primitive introspection*. Primitive introspection is a non-classificatory kind of phenomenal-state introspection: it is a kind of phenomenal-state introspection that does not involve recognizing the introspected phenomenal state as an instance of any experience type. I defend three main claims about primitive introspection. First, it *exists*: there is a mental phenomenon that has the features I attribute to primitive introspection and such a phenomenon is a full-fledged introspective process. Second, its *nature* is best accounted for by a version of the *acquaintance* theory—what I call the *integration account*. Third, it has a *distinct epistemic value*: it provides the subject with knowledge of their phenomenal states. Such knowledge is a *sui generis* kind of knowledge: knowledge by acquaintance. Knowledge by acquaintance displays a distinctive epistemic property which is somewhat analogous to *infallibility*: it provides the subject with a *complete and perfect grasp* of the phenomenology of their experience.

KEYWORDS

Self-Knowledge, Introspection, Phenomenal concepts, Knowledge by acquaintance, Acquaintance, Non-conceptual.