



La place de la valeur dans le langage naturel : une approche non-factuelle à la signification des termes évaluatives et des jugements de valeur

Andrés Soria Ruiz

► To cite this version:

Andrés Soria Ruiz. La place de la valeur dans le langage naturel : une approche non-factuelle à la signification des termes évaluatives et des jugements de valeur. Philosophie. Université Paris sciences et lettres, 2019. Français. NNT : 2019PSLEE059 . tel-03117265

HAL Id: tel-03117265

<https://theses.hal.science/tel-03117265>

Submitted on 21 Jan 2021

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

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



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à Institut Jean Nicod, DEC,
École Normale Supérieure

The Place of Value in Natural Language

Soutenue par

Andrés SORIA RUIZ

Le 17 Juin 2019

Ecole doctorale n° 540

**Ecole doctorale Lettres, Arts,
Sciences humaines et
sociales.**

Spécialité

Philosophie



Composition du jury :

Paul, ÉGRÉ Directeur de recherche, Professeur; Institut Jean Nicod-ENS-CNRS	<i>Président</i>
Elisabeth CAMP Professor, Rutgers University	<i>Rapporteur</i>
Neftalí , VILLANUEVA FERNÁNDEZ Profesor titular, Universidad de Granada	<i>Rapporteur</i>
Carla UMBACH Wissenschaftliche Mitarbeiterin, Leibniz-Zentrum Allgemeine Sprachwissenschaft (ZAS)	<i>Examineur</i>
Chris BARKER Professor, New York University	<i>Examineur</i>
Isidora STOJANOVIC Directrice de recherche Institut Jean Nicod-ENS-CNRS	<i>Directrice</i>

The Place of Value in Natural Language

A non-factualist approach to evaluative terms and judgment

Andrés Soria Ruiz

A dissertation submitted in partial fulfilment
of the requirements for the degree of

Doctor of Philosophy

Supervisor: Isidora Stojanovic

Institut Jean Nicod
École Normale Supérieure
Paris Sciences et Lettres / CNRS
Awarded unanimously on June 17th 2019

A la memoria de Arcadio Serrano Molina, *Tatín*.

“Todo lo intentábamos ver desde ese prisma mágico que adquieren
las historias cuando uno las escribe, por más que no las escribiéramos”
Ramiro González Coppari

Acknowledgments

There are far too many people to thank for this work, ranging from bureaucrats and bouchers to friends and family. I want to thank *all* of them, even if I don't know their names or I forget to mention them. Those who know me are aware of my bad memory. There is a positive side to it, though: it prevents me from holding grudges. So if you can't find your name here (and expected to), be sure that it is not due to bad faith.

I would like to start by thanking the members of my committee: Paul Égré, Liz Camp, Neftalí Villanueva, Chris Barker, Carla Umbach and Isidora Stojanovic. The beneficial effect of Paul's presence at Institut Jean Nicod cannot be overstated. He has helped me at key points throughout my thesis, and his combination of humbleness, wisdom and curiosity is something for every student of logic and language to aspire to. Liz was incredibly kind in writing a wonderful *pré-rapport* which was the first "external" opinion I heard about my thesis, and it was so positive that it propelled me to prepare the defense with a smile on my face. I can hardly express my gratitude to Neftalí. In a nutshell: I would not be doing analytic philosophy if it wasn't for him. He was first my teacher but he quickly became one of my best friends (and occasional bandmate), and in the process he turned me into a colleague. I have felt his advice and support at every step of the way. It wasn't always easy. NYU was occasionally overwhelming and disorienting, but Chris helped me find my way into formal semantics. He combines scholarship and cool in a way that no one else does. And he encouraged me—wisely—to come to Paris and work with Isidora. Finally, Carla has been of tremendous help writing the final version of this work, as she took the time and trouble to help me through the most formal bits of the dissertation. I am eternally thankful for that. However scared and nervous I was, my defense was literally the most exciting and gratifying moment of my academic career so far. And this was thanks to the incredibly warm, constructive tone and atmosphere in which it took place. It was a fantastic experience and I could not have wished for a better group of scholars to read and assess my work.

My supervisor, Isidora Stojanovic, deserves special thanks. I am honoured that she took me as her student. She relied on my project and fought for it from the very start, even when she barely knew me. During these years, she has guided me gently to the topics that she found most exciting for me to work on. Ultimately, I was free to choose the focus, angle and breadth of my thesis, which I am extremely grateful for. But if I managed to produce a serious work of philosophy, it is absolutely thanks to her.

Many people have contributed to this project with ideas, criticism and support. I want to start by thanking the wonderful people that I met at NYU Philosophy, Linguistics and elsewhere: Kyle Blumberg, David Chalmers, Hartry Field, Camil Golub, Rob Hopkins, Paloma Jeretič, Arden Kohler, Andrew Lee, Sam Lee, Matt Moss, Yining Nie, Jim Pryor, Alex Rigas, Mary Katherine Robinson, Lorenzo Rossi, David Storrs-Fox, Ildi Szabo, Anna Szabolcsi and Andy Yu. Throughout these years, I have met and had exchanges with great people in academic visits and conferences, such as Daisuke Bekki, James Brown, Ramiro Caso, Elena Castroviejo, Matthew Chrisman, Cleo Condoravdi, Alexander Dinges, Berta Grimau, Katie Fraser, Javier González de Prado, Matt Jope, Max Kölbel, Marcin Lewinski, Carol Rose Little, Manolo Martínez, Adam Marushak, Louise McNally, Elin McReady, Marina Ortega-Andrés, Erich

Rast, Craige Roberts, Andrés Saab, Glenda Satne, Uli Sauerland, Giorgio Sbardolini, Carlota Serrahima, Stephanie Solt, Agustín Vicente, Hans Wilke and Dan Zeman.

The ever-increasing group of philosophers working at or stemming from the University of Granada deserve special mention: Juan José Acero, Manuel Almagro, David Bordonaba, Francesco Consiglio, Víctor Fernández (briefly kidnapped to Paris), José Ferrer, María José Frápolli, Daniel Galdeano, Pedro García, Nemesio García, Amalia Haro, Manuel Heras, José Luis Liñán, Alba Moreno, Llanos Navarro, Xabier Osorio, José Palma, Manuel de Pinedo, Eduardo Pérez (*this station is non-operational!*), Mirco Sambrotta y José Ramón Torices. Gracias por hacer de Granada un sitio tan especial para hacer filosofía.

I could not have wished for a better place than Paris to spend my PhD. The relatively small size of IJN, embedded in the wonderful Département d'Études Cognitives of the ENS, makes it an ideal place to do philosophy. Among professors and students (including visitors), administrators and other Parisian characters, I want to thank Nathalie Abitbol-Evin, Mar Alloza, Mireille Babineau, Manuel Beiran, Romain Bourdoncle, Brian Buccola, Roberto Casati, Filippo Contesi, Inés Crespo, Alice Doreille, Rachel Dudley, Émile Enguehard, Andreas Falck, Pablo Fernández, Clothilde Facon, Nils Franzén, Luca Gasparri, Vincent Gaudefroy, Anna Giustina, Helena Hachmann, Naomi Havron, Andreas Heise, Benjamin Icard, Marco Inchingo, François Kammerer, Uriah Kriegel, Manfred Krifka, Manuel Križ, Jeremy Kuhn, Palle Leth, Guido Lohr, Philippe Lusson, John MacFarlane, Paula Marchesini, Salvador Mascarenhas, Olivier Massin, Robert May, Gerda Melnik, Michae Murez, Takuya Niikawa, Boian Nikiforoff-Kharisanoff, Michele Orrù, Eleni Palaologou, Joanna Patriarca, Agostino Pinna-Pintor, David Plunkett, Paul Portner, Ignacio Rebollo, François Recanati, Philippe Schlenker, Nura Sidarus, Benjamin Spector, Brent Strickland, Chloe Tahar, Tristan Thommen, Juliette Vazard, Daria Vitasovic, Christopher Vogel, Yaron Wolf, Hualin Xiao and Julia Zakkou. I owe a special mention to the fantastic philosophers in my (loose) cohort: Géraldine Carranante, Víctor Carranza, Maryam Ebrahimi Dinani, Sandra Lasry, Armando Lavallo, Tricia Magalotti and Louis Rouillé. I've been told a couple of times that our arrival at IJN improved the atmosphere and conviviality. That makes me very proud!

The PhD was not only bearable, but actually a fun and exciting ride thanks to Alex Martin and Katie Prosser, who hosted me in Edinburgh right before submitting and witnessed my various mental breakdowns; to Bianca Cepollaro, who first convinced me to come to IJN and then introduced me to everyone when I arrived; Natalia Karczewska and Michal Wójtowicz, for their friendship, support and hospitality from NYC to Warsaw; Dan Hoek, who I befriended at NYU and then re-befriended at IJN to live many Parisian adventures; Federico Faroldi, who is the best co-author, logician and friend one could wish for; Matheus Valente, without whose occasional dose of calm and cheer I don't think I'd survived this; Camila Scaff, who even made panicking before submitting fun; Slawa Loev, my first Parisian friend and co-explorer of bars, parks and quartiers; Diego Feinmann, who made me laugh till I literally fell off the chair; and Milica Denic, with whom I shared the most therapeutical lunches, tea pauses and hugs. I want to particularly thank Amir Anvari, my Iranian-Chomskyan roommate with a soft spot for Andalusian prog. Thanks for the rollies, the whiskies and the wisdom. Y gracias infinitas a Mora Maldonado, por tantos momentos de complicidad absoluta, de cariños en la frontera; gracias por confiar en mí y hacer mi trabajo más científico entre *mantecatura* y *Te Lo Resumo*. You made Paris feel like the left side of this meme (even though the material conditions were definitely more like the right side):



Por último quiero agradecer a mis amigos Ramis, Hugo, Mariano, Marcos, Alberto, Dani, Llodra, Nacho, Nolo, Caba, Obi, Migue, María, Belinda, Agu, Vero, Joaquín, Lucía, David, Finn, Andy, Agustín y Alber; por oír mis varios coñazos y sacarme (o dejarse ser sacados por mí) de marcha; a mi familia (en especial a mi abuela Asunción, que tuvo clarísimo que me tenía que ir a París); a mis padres Guada y Chris, Andrés y Lauri, por quererme y apoyarme siempre, compartir mi histeria pre-soutenance y hacer sangría para (literalmente) varios pots de thèse; a mi hermanita Emma, que acertó en pedirme que no la masfastidie más y desde entonces hemos sido los mejores hermanos posibles; y a mi abuelo Manolo, por el talento.

Abstract

This is a dissertation about words like *good*, *better*, *bad*, *beautiful*, *fun* or *cruel*. More abstractly, this dissertation is concerned with the question of how to accommodate evaluative expressions, which the former words exemplify, within a theory of natural language meaning. To do that, I recruit tools from both metaethics and formal linguistics. From the point of view of metaethics, this dissertation defends a form of non-factualism, since it defends that evaluatives and the sentences in which they paradigmatically appear play a different role in communication from the standard role assigned to declarative sentences, namely, describing reality and imparting information about it. From the point of view of formal linguistics on the other hand, this work relies on advances in the semantics of gradable adjectives, which evaluatives generally belong to, to gain insight into their semantic properties. Non-factualism is a promising account of the meaning of evaluative sentences in spite of the well-known Frege-Geach problem. Chapter 2 of this dissertation presses the idea that, if non-factualism is to be defended as a linguistically serious proposal, it needs to be implemented compositionally. And since evaluative adjectives are gradable expressions, this is challenging. In particular, non-factualism is most often formulated in such a way that captures most accurately the meaning of adjectives in their positive form. The main challenge, then, is to use such formulation to derive a scalar system. With that objective in mind, the central chapters (3) and (4) of the dissertation are devoted to offering a compositional semantics for a propositional language enriched with a pair of positive and negative unary sentential operators, together with binary operators representing various comparative relations. Following a delineation approach, the semantics of the comparative is derived from that of the positive form. This preserves the non-factualist insight that evaluative language is primarily used to express simple attitudes of support and rejection. Subsequently, the semantic model associated with this formal language is applied to a set of natural language evaluative sentences. Further properties of evaluative adjectives are studied, such as the distinction between relative and absolute-standard adjectives and their scale type. Chapter 5 discusses the notions of judge-dependency and ordering subjectivity. In recent experimental work, Stephanie Solt has argued that the spectrum-like nature of the phenomenon of ordering subjectivity calls for a distinction, within the class of subjective adjectives, between multi-dimensional and judge-dependent ones. Her conclusions are supplemented by arguing that extending her experimental paradigm to moral adjectives suggests a further distinction, within the class of judge-dependent adjectives, between (at least) experiential and evaluative adjectives. Finally, Chapter 6 tackles the distinction between evaluative and descriptive uses of evaluative sentences. It is argued that a brand of dynamic non-factualism offers the best account of such double uses of evaluative sentences.

Resumé

Cette thèse porte sur des mots comme *bon*, *mieux*, *mauvais*, *beau*, *beau*, *amusant* ou *cruel*. Plus abstraitement, cette thèse s'intéresse à la question de savoir comment compter pour les expressions évaluatives, que ces mots illustrent, dans une théorie de la signification du langage naturel. Pour ce faire, je recrute des outils à la fois en métaéthique et en linguistique formelle. Du point de vue de la métaéthique, cette thèse défend une forme de non-factualisme. En effet, l'usage des expressions évaluatives est différent de l'usage standard assigné aux phrases déclaratives, à savoir, décrire la réalité. Du point de vue de la linguistique formelle, en revanche, ce travail s'appuie sur les progrès de la sémantique des adjectifs gradables, auxquels sont généralement associés les évaluatifs, afin de mieux comprendre leurs propriétés sémantiques. Le non-factualisme est une théorie prometteuse de la sémantique des phrases évaluatives en dépit du problème bien connu de Frege-Geach. Le chapitre 2 de cette thèse insiste sur l'idée que, si l'on veut que le non-factualisme soit défendu comme une proposition sérieuse sur le plan linguistique, il doit être mis en œuvre d'une manière compositionnelle. Dans la mesure où les adjectifs évaluatifs sont des expressions gradables, cette tâche est difficile. En particulier, le non-factualisme est souvent formulé de telle manière qu'il s'applique bien à la signification des adjectifs dans leur forme positive uniquement. Le principal défi est donc de partir d'une telle formulation pour obtenir un système scalaire. Dans cette perspective, les chapitres 3 et 4 de cette thèse sont consacrés à présenter une sémantique compositionnelle pour un langage propositionnel enrichi de deux d'opérateurs unaire et binaire représentant les diverses relations de comparaison. Suivant une approche de délimitation sémantique, la forme comparative est dérivée de celle de la forme positive. Cela préserve l'idée non-factuelle selon laquelle le langage évaluatif est principalement utilisé pour exprimer des attitudes simples d'adhésion et de rejet. Ensuite, le modèle sémantique associé à ce langage formel est appliqué à un ensemble d'expressions évaluatives en langage naturel. D'autres propriétés des adjectifs évaluatifs sont étudiées, comme la distinction entre adjectifs relatif et absolu, et leur type d'échelle. Le chapitre 5 traite de la dépendance au jugement et de la subjectivité scalaire. Dans des travaux expérimentaux récents, Stephanie Solt a soutenu que la nature spectrale du phénomène de la subjectivité scalaire exige une distinction, dans la catégorie des adjectifs subjectifs, entre des adjectifs multidimensionnels et dépendants d'un jugement. Ses conclusions sont complétées par l'argument selon lequel l'extension de son paradigme expérimental aux adjectifs moraux suggère une distinction supplémentaire, dans la classe des adjectifs dépendant d'un jugement, entre les adjectifs expérientiels et évaluatifs. Enfin, le chapitre 6 traite de la distinction entre les utilisations évaluatives et descriptives des expressions évaluatives. On fait valoir qu'un type de nonfactualisme dynamique est celui qui rend le mieux compte de cette double utilisation des expressions évaluatives.

Contents

1	Introduction: an overview about evaluative language	1
1.1	Introduction	1
1.2	The problem of value in language	2
1.2.1	The distinction between fact and value	2
1.2.2	The difference between describing and evaluating	3
1.3	Evaluative expressions	5
1.3.1	The empirical landscape: evaluatives in natural language	6
1.3.2	Topics in the study of evaluatives: a small tour	6
1.4	Outline of the chapters in this dissertation	9
2	Evaluativity in linguistics and philosophy: gradability & non-factualism	11
2.1	Introduction	11
2.2	Evaluatives from a linguistic point of view	13
2.2.1	Gradability & subjectivity	13
2.2.2	Evaluatives in degree semantics	16
2.2.3	The notion of evaluativity in linguistics	22
2.3	Evaluatives from a metanormative point of view	23
2.3.1	Action-guidance	23
2.3.2	Non-factualism about evaluative language	26
2.4	The (supra-sentential) Frege-Geach problem	32
2.5	The sub-sentential Frege-Geach Problem	37
2.5.1	Non-factualism and the positive form: POS-ATT	37
2.5.2	From degree to delineation semantics	39
2.6	Delineation semantics & non-factualism	42
2.7	Recap & Conclusion	45
3	Hyperplan semantics for gradable evaluative adjectives	47
3.1	Introduction	47
3.2	Semantically hardwiring action-guidance	49
3.2.1	Semantic pluralism vs. monism about evaluatives	49
3.2.2	Semantic contextualism vs. relativism about evaluatives	51
3.2.3	Action-guidance in plan-semantics	54
3.3	Hyperplans	56
3.3.1	Gibbard 2003’s proposal	56
3.3.2	Yalcin 2012’s intensional implementation	61
3.3.3	Hyperplans as functions over alternatives	65
3.4	Evaluating alternatives: the language of \mathcal{L}_V	69
3.4.1	Syntax	69

3.4.2	Semantics	70
3.5	A toy fragment & structure	75
3.6	Order and consistency	77
3.6.1	Inducing order	77
3.6.2	Preserving consistency	79
3.7	Conclusion	82
4	From \mathcal{L}_V to natural language	83
4.1	Introduction	83
4.2	Thin evaluatives as sentential operators	86
4.2.1	Alternatives, focus and background information	88
4.2.2	Semantics: introducing the alternative-generating function	89
4.2.3	<i>Good, bad, better & worse</i>	91
4.3	Beyond evaluative operators	97
4.3.1	Other bearers of value	99
4.3.2	Specific flavors of evaluation	104
4.3.3	Thick evaluatives	106
4.4	Solving the Frege-Geach Problem	108
4.4.1	Solving the supra-sentential Frege-Geach Problem	109
4.4.2	Solving the sub-sentential Frege-Geach Problem	110
4.5	Relative- vs. Absolute-standard	113
4.5.1	Entailment patterns	114
4.5.2	Comparison class sensitivity	115
4.5.3	Scale bounds	117
4.5.4	Discussion	118
4.6	Scale structure	120
4.6.1	Ordinal	121
4.6.2	Ratio	123
4.6.3	Interval	125
4.6.4	Discussion	127
4.7	Recap & conclusion	130
5	Subjectivity & evaluativity: an experimental approach to ordering subjectivity	132
5.1	Introduction	132
5.2	Solt’s 2018 study	134
5.2.1	Experimental design	136
5.2.2	Results	137
5.2.3	Solt’s account	137
5.3	Our study	140
5.3.1	Participants	141
5.3.2	Materials	141
5.3.3	Procedure	142
5.3.4	Results	143
5.4	Breaking down judge-dependency	145
5.5	Recap & conclusion	147
	Appendix: experiment materials	147
6	The Dynamics of value judgment	150
6.1	Introduction	150

6.2	Evaluative and descriptive uses of evaluative sentences	153
6.2.1	Epistemic status	155
6.2.2	Deceit potential	156
6.2.3	Subjective disagreement	157
6.2.4	Action-guidance	159
6.3	Evaluative <i>versus</i> metalinguistic use of language	161
6.4	Dynamic semantics: a précis	165
6.4.1	Assertion, commands and other animals (in dynamic semantics)	166
6.4.2	Hyperplans and common ground	169
6.5	Evaluative & descriptive update (1st attempt)	170
6.6	Evaluative & descriptive update (2nd attempt)	173
6.6.1	Basic dynamic entry	176
6.6.2	Evaluative use	179
6.6.3	Descriptive use	180
6.6.4	Mixed use	182
6.7	Conclusion	184
7	Concluding remarks	185
	References	194

Chapter 1

Introduction: an overview about evaluative language

Summary

This introductory chapter surveys contemporary debates in metaethics about the nature of evaluative concepts as well as some key issues regarding the meaning of evaluative expressions in the philosophy of language and formal linguistics. These are all issues that to a smaller or larger extent, are dealt with in this dissertation. The last section contains an outline of the following chapters.

1.1 Introduction

This is a dissertation about words like *good*, *better*, *bad*, *beautiful*, *fun* or *cruel*. More abstractly, this dissertation is concerned with the question of how to accommodate evaluative expressions (evaluatives for short), which the former words exemplify, within a theory of natural language meaning. To do that, I recruit tools from both metaethics and formal linguistics. From the point of view of metaethics, this dissertation defends a form of NON-FACTUALISM, since it defends that evaluatives and the sentences in which they paradigmatically appear play a different role in communication from the standard role assigned to declarative sentences, namely, describing reality and imparting information about it. From the point of view of formal linguistics, this work relies on advances in the semantics of gradable adjectives, which evaluatives belong to, to gain insight into their semantic properties.

This introductory chapter has a double purpose: the first is to argue that there is something both philosophically and linguistically special about evaluatives, which sets them apart from other expressions of natural language. The second is to situate this dissertation's contribution within the current debate about evaluatives. I aim to tackle both issues by considering the distinctive features of evaluative language and talk, and seeing how the study of these properties touches upon a number of long-standing research topics in philosophy and linguistics. The chapter is structured as follows: §1.2 motivates the specific challenges raised by a linguistic and philosophical inquiry into evaluative language, by briefly discussing the traditional philosophical distinction between fact and value (§1.2.1) as well as the intuitive distinction between using language to describe the world *versus* using it to evaluate it (§1.2.2). §1.3 introduces various issues in connection with the properties of evaluative expressions. We start in §1.3.1 by characterizing—extensionally—the expressions of natural language that are conventionally

taken to carry evaluative meaning. In §1.3.2 we turn to a quick survey of a few important topics in the semantics of evaluatives: what “flavor” of value they can convey (§1.3.2), what type of arguments they admit (§1.3.2), what is special about their gradability (§1.3.2) and the distinction between thin and thick evaluatives (§1.3.2). To conclude, §1.4 contains a brief outline of each chapter of the dissertation.

1.2 The problem of value in language

The intuition that evaluatives have special and intriguing features is, of course, not new: evaluatives have caught the attention of philosophers throughout history. But those features were not studied as pertaining to words like *good* or *beautiful*, but rather as properties of the very idea or concept of goodness or beauty. Indeed, some of the most recalcitrant problems of philosophy have to do with figuring out the place of value within certain world-views: the problem of evil, the nature of moral and aesthetic judgment, the characterization of the virtuous life or the definition of knowledge are all topics that in one way or other touch upon this issue.

With the turn of the XXth century, philosophers became preoccupied with language, and so a concern for value became a concern for the way in which value is represented in natural language, that is, a concern for evaluative expressions. Simultaneously, the development of formal linguistics and philosophy of language offered new and sophisticated tools to explore natural language meaning. As it turns out, evaluative expressions are interesting not only insofar as they vehicle traditional problems of philosophy, but are also intriguing in their own—linguistic—right. In this section, we first look at the distinction between fact and value (§1.2.1) and then we turn to the difference between describing and evaluating (§1.2.2).

1.2.1 The distinction between fact and value

Throughout the history of philosophy, one finds time and again the intuition that there exists an important difference between facts and values, and furthermore, that a gap separates the realms of the factual and the evaluative. Hume, Moore or Wittgenstein are famous defenders of this idea. Hume’s *Enquiry* (1739/1999) contains his notorious denial that OUGHT can be inferred from IS—what came to be known as ‘Hume’s law’, or the ‘is-ought problem’. In his *Principia Ethica* (1903/1993), Moore held the view that the property of being good cannot be reducible to any natural property, and to argue for that view he coined the famous ‘open question argument’: for any set of properties that one might offer as a definition of the word *good*, it seems to always remain an open question whether something that falls under that definition is indeed *good*. Finally, Wittgenstein’s *Lecture on Ethics* (1929/1965) contains a powerful statement of the apparently insurmountable gap between the factual and the evaluative or normative.

What makes these realms, the realm of the factual and the realm of the evaluative, so different? Here is where considerations about language have appeared to offer a different angle towards answering that question. If instead of considering the difference between the *realms* of facts and values we look at the difference between factual and evaluative *language*, there is a double hope: first, that attention to evaluative language will tell us something about evaluative concepts—that by studying the word *good* we might learn something about “the” good. And

more specifically, following the first Wittgenstein, one might hope that attention to evaluative language might reveal some issues in ethics as pseudo-problems.

Regardless of whether the latter hope was in fact fulfilled, philosophers working in the analytic tradition started to distinguish, chiefly after Moore and following the development of ordinary language philosophy, the domains of NORMATIVE ETHICS and METAETHICS as two distinct sub-fields of philosophical enquiry (Darwall et al. 1992). The first remained preoccupied with the traditional questions of ethics, such as whether it is ever permissible to lie, while the second turned to the question about the meaning of ethical terms, such as *permissible*.¹ Questions about the meaning of moral terms are questions about moral terms, ontology, psychology and knowledge.

In recent decades, the scope of the distinction between normative ethics and metaethics has broadened to capture a similar contrast between enquiry into other evaluative domains *versus* enquiry into the language in which we conduct the former kind of enquiry. Thus, METAEPISTEMOLOGY has been distinguished from epistemology; and METAAESTHETICS from aesthetics (e.g. Alston 1978; Schellekens 2008). Authors like Chrisman (2016) speak more generally of METANORMATIVE THEORY. Moreover, by turning their attention to evaluative language, the metanormative theorist can (or rather *must*) take advantage of the developments of contemporary formal linguistics, which offers increasingly sophisticated tools to study the semantics of natural language expressions. What is really appealing is that evaluative language is not just of interest to the metanormative theorist. Much to the contrary, evaluatives are important case studies for central topics in formal semantics and pragmatics and the philosophy of language in general, such as the semantics of modal, gradable and vague expressions, the interaction between context and content or the performative aspects of language.

1.2.2 The difference between describing and evaluating

We will explore some of those topics both in this introduction and in the dissertation at large, but before that: what is—intuitively—the difference between using language to describe and using language to evaluate? I say *intuitively*, because this is a distinction that we make in everyday speech, albeit maybe under other names or disguises. For example: we speak of stating the “bare” facts *vs.* giving one’s opinion. We expect the judges in a synchronized swimming exhibition to evaluate the participants; but we expect the TV commentators (at least some of them) to describe what is happening. We read the first page of the newspaper to be informed about the latest events, but we turn to the op-eds in the latter pages for assessment of those events. In all those cases, we seem to be exploiting language for two different purposes.

Importantly, this pre-theoretical distinction between describing and evaluating already poses a challenge for the philosophy of language. Since its foundational days, philosophers of language have preoccupied themselves with providing a satisfactory picture of how human communication works. And the paradigmatic model for communication has traditionally been that by which speakers exchange information about their surroundings *via* declarative sentences with verbs in the indicative mood, such as *it’s raining*, *the cat is 2 months old* or *I had a couple of beers last night*. Those sentences represent the world as being a certain way, and our purpose in uttering them is to share the information that is represented in them.

¹As Smith 1994, p. 2 puts the distinction, ethics is concerned with certain questions, while metaethics is concerned with questions about those questions. See Miller 2003 for an introduction to metaethics.

The problem is that when we look at evaluative uses of language we find sentences that look superficially similar to the sentences above, but which are deployed with very different communicative intentions. To fix ideas, let us speak of making two different kinds of JUDGMENT or ASCRIPTIONS, of fact and value.² At first sight, judgments of value are no different from judgments of fact. We express both types of judgment by uttering declarative sentences with verbs in the indicative mood; syntactically, judgments of value and judgments of fact are *prima facie* on equal footing. However, many share the intuition that they are importantly different. Most relevantly, value judgments do not convey factual information in the way that the sentences in the previous paragraphs usually do. Let us stop for a moment in two aspects that set sentences expressing judgments of fact and value apart from each other.

First, if I utter the sentence *the cat is 2 months old*, I seem to be informing you of a state of affairs that is true of the world in which we live; it concerns an object of the world (the cat) and a property of it, namely her age. By contrast, if I say *the cat is annoying*, I am still talking about the cat, but this time it is not clear that I am attributing a property to it—being annoying—that is on a par with her age. For one thing, being annoying is not a property that the cat possesses in and out of itself, so to speak, but in relation to me. It annoys me, but you might find it adorable. Furthermore, most often when I say *the cat is annoying* I will have no obvious intention of informing you of specifics of the cat’s behavior or character, since if those were my intentions, it would be much more useful to say *the cat does such-and-such*. Rather, when I say *the cat is annoying* I am voicing my subjective opinion or appreciation of the cat, not merely describing its behavior. One could perhaps go as far as to say that I am providing you with as much information about the cat as about myself, since I am letting you know something about my taste and values, about the things that I esteem and despise.

Another sense in which judgments of fact and judgments of value come apart is that only the latter reveal an intimate connection to action. The metaethics literature has long expanded on this feature of evaluatives, sometimes called ACTION-GUIDANCE or the constraint of PRACTICALITY (see e.g. Björklund et al. 2012; Darwall et al. 1992; Dreier 1990; Hare 1952). If I say *the cat is two months old*, there is no obvious way in which saying that commits me to adopting one course of action or other. I might then go see the cat, play with it or never pay attention to it again. All of those actions seem compatible with, and independent of, my claim about the cat’s age. By contrast, by saying *the cat is annoying* I seem to be incurring some sort of practical commitment. In particular, my audience might reasonably expect me to avoid the cat’s company. If I said that the cat is annoying and yet I spent all my time caressing and playing with the cat, anyone who heard me would have a reason to challenge my sincerity.

This appears to hold of evaluative claims across the board, but not of factual claims. Ascribing a positive value to something, describing it as *advisable*, *beautiful* or *justified*, commits one to orienting one’s actions towards the promotion, support or obtainment of those things. And similarly, describing things as *bad*, *impractical*, *disgusting* or *unwarranted* commits oneself to avoiding the things so called. Evaluative language appears to be connected to action in a way that factual language is not (an early reference of this observation is Stevenson 1937, who talked about the ‘magnetism’ of value: the good attracts; the bad repels).

In general, when we use evaluative language we seem to be aiming at eliciting some kind of practical response: op-eds seek to produce a certain attitude on the reader, not simply to

² As I will use these terms, a fact [value] ASCRIPTION or ATTRIBUTION is a communicative act by which one attributes factual [evaluative] properties to some object(s). It is the outward, public expression of a factual [value] JUDGMENT, which is a mental act.

provide her with information; similarly, the judges of the swimming competition are not there to say what the participants are doing, but to assess the quality of their performance so that a ranking is established and prizes are won. But how does this happen? Is it *via* the use of certain expressions, phrases, words? Is it in virtue of the context in which a certain sentence is used that it turns out to describe or evaluate? These are the kind of questions that this dissertation tackles.

1.3 Evaluative expressions

It is most natural to give an affirmative answer to the first of the two questions we just posed: yes, when speakers use language to evaluate they use certain expressions. To see this, consider the contrast within each pair of sentences:

- (1.1) a. Stealing is common.
b. Stealing is unethical.
- (1.2) a. Anna Karenina is Russian.
b. Anna Karenina is brave.
- (1.3) a. Sushi is cold.
b. Sushi is tasty.
- (1.4) a. It's possible that she sent that e-mail.
b. It's great that she sent that e-mail.
- (1.5) a. Your suspicions are shared.
b. Your suspicions are unfounded.
- (1.6) a. Elena Greco is a Neapolitan.
b. Elena Greco is a hero.
- (1.7) a. The third movement was executed slowly.
b. The third movement was executed beautifully.
- (1.8) a. Voting costs money.
b. Voting matters.

In most contexts, uttering the *b*-sentences of each of these pairs would be a way for the speaker to express a judgment of value, while uttering the *a*-sentences would most naturally be a way for the speaker to express a judgment of fact. The only difference between the *a*- and *b*-sentences is the presence of terms like *unethical*, *brave*, *great* or *beautifully*.

In what follows, we will call sentences like the *b*-sentences in this list—namely, unembedded declarative sentences containing evaluatives—EVALUATIVE SENTENCES, and the speech acts that those sentences are characteristically used to perform will be called ATTRIBUTIONS OF VALUE. An attribution of value is a type of communicative act whereby a speaker communicates evaluative information, that is, information to the effect that something has a ‘positive or negative standing—merit or demerit, worth or un-worth—relative to a certain kind of standard’ (Väyrynen 2013, p. 29). In this section, we look at what words these are (§1.3.1), and we survey some issues in connection to their semantics (§1.3.2), namely the different flavors of evaluativity, the variety of arguments that evaluatives admit, their gradability and the distinction between thin and thick evaluatives.

1.3.1 The empirical landscape: evaluatives in natural language

Evaluatives belong to a larger class of NORMATIVE expressions, those which in some way or another are connected to norms, rules or prescriptions. Besides evaluatives, normative expressions include the class of DEONTIC expressions, which is not our focus in this work.³

The most common type of evaluatives are adjectives. Here is a (partial) list:

great, good, bad, just, justified, credible, beautiful, ugly, virtuous, ethical, important, tasty, clear, advisable, disgusting, delicious, courageous, chaste, evident, charming, lazy, stupid, lucky, smart, brave, sensible, cowardly, fearful, timid, clever, intelligent, brilliant, idiotic, foolish, pretty, gorgeous, elegant, handsome, hideous, doubtful or precious.

Of course, one can build adverbs out of all those adjectives. Evaluative nouns (*idiot, hero, friend*)⁴ and verbs (*to matter, to justify, to blame*) are less numerous, but equally quotidian. Finally, certain modifiers like *too* or *really* (as in *too / really tall*) arguably carry evaluative meaning as well.

That list is not intended to be exhaustive—expressions of other syntactic categories might be considered evaluative as well. Moreover, current research is exploring the possibility that whether a certain adjective is evaluative is not immediately transparent (such is the case of *normal*, see Bear and Knobe 2017). But the most common examples fall under the categories just indicated. Conversely, there are syntactic categories that for the most part do not feature evaluatives, such as determiners or quantifiers.⁵ This dissertation is focused on adjectives, although some of the claims that I make apply to evaluatives of other syntactic categories. For example, gradability is a property most clearly instantiated by (and studied through considerations about) adjectives, but there are gradable nouns and verbs as well (Lassiter 2017; Sassoon 2013b). In fact, evaluative nouns and verbs are gradable; and most of what is said here about the gradability of evaluative adjectives probably carries over to evaluative nouns and verbs as well. However, nouns and verbs will not be specifically discussed.

1.3.2 Topics in the study of evaluatives: a small tour

In this section, we briefly introduce some important topics in the semantics of evaluatives: the different flavors of evaluativity, the variety of arguments that they can take and their gradability.

³ Deontics include modal verbs like *ought, may* or *have (to)*, and verbs like *allow* or *forbid*. Deontics express the propositions and inferential patterns characteristic of deontic logic, that is, the logic of obligation and permission. Even though there exist clear conceptual ties between deontic and evaluative terms (e.g., one can define what is morally good in terms of what one is obliged to do, and *vice versa*), there are reasons to keep the two realms apart. To mention two: many deontic adjectives, such as *obligatory, permitted* or *forbidden*, are not gradable. And claims about what is obligatory, permitted or forbidden do not seem expressions of judgments of value, or at least not as intuitively as claims about what is good, beautiful or advisable (see Carr 2017 for an introduction to deontic modals and Tappolet 2013 for a more extensive discussion on the difference between evaluative and deontic concepts). Except tangentially in Chapter 3, deontics are not discussed in this dissertation.

⁴ It has been pointed out recently that some nouns that appear to be descriptive might additionally encode a normative or evaluative dimension in their meaning. These have been called ‘dual-character concepts’. For example: a father is, descriptively speaking, a man who has children. But a father is also someone who satisfies certain stereotypes associated with fatherhood. Many of these are normative, such as caring for one’s children, treating them kindly, etc. See Del Pinal and Reuter 2017; Leslie 2015. Dual character concepts are not the object of this dissertation.

⁵ Although see Egré and Cova 2015 for the evaluative implications of *many*.

Flavors of evaluativity

The above list contains evaluatives belonging to what we may call different “flavors” of value (to adopt an expression commonly associated to different kinds of modalities in language, c.f. Kratzer 2012): first and foremost, we find evaluatives of moral praise and blame (*ethical, virtuous, right, wrong, worthy, brave*). These have to do with what is morally good or what is our duty. There are evaluatives of practical commendation and criticism (*advisable, sensible*), which concern how best to orient action given our goals and interests. There are epistemic evaluatives (*credible, evident, doubtful or clear*, as in *it is clear that*, see C. Barker 2009), having to do with knowledge, evidence and the justification of belief. Moreover, there is a large class of evaluatives that can be interpreted either epistemically or morally/practically, depending on whether these predicates are applied to beliefs or actions (*justified, rational, reasonable, warranted*). One finds aesthetic adjectives (*beautiful, ugly, elegant*) and predicates of personal taste (PPTs, e.g. *tasty, fun, disgusting, boring*), the latter of which have received a great deal of attention in recent philosophy of language and linguistics (see Lasersohn 2005; MacFarlane 2014; Stephenson 2007b; Stojanovic 2007; and McNally and Stojanovic 2017 for the connection and contrast between aesthetic adjectives and PPTs). And finally, there are what we may call “all-purpose” evaluatives that can express more than one flavor of evaluativity (*good, bad, valuable*).

Objects of evaluation

Another important issue is: what are the bearers of value? In other words: what type of arguments do evaluatives take? The set of examples at the start of this section reveals a surprising variety. Evaluatives can be applied to individuals (including people, things or abstract concepts), action types, events or propositions. Some evaluatives are selective about their arguments (events cannot be heroines; people cannot be advisable), but all-purpose evaluative terms like *good* are very flexible: they can be predicated of people (*Mary is good*), objects (*this knife is good*), actions (*volunteering is good*) and events (*the Spanish Confiscation was good*). From the viewpoint of semantics, this feature is surprising and challenging, since it is not shared by many other classes of expressions and requires a sophisticated story of how evaluatives combine with such different objects.

Gradability

Finally, most (if not all) evaluatives are GRADABLE. This means that in order to give a semantics for evaluatives we need to employ scales and orderings. In linguistics, gradability in the adjectival domain is attested by the admissibility of adjectival modification. To see this, contrast the following pair of sentences:

- (1.9) a. # Anna Karenina is more married than Elena Greco.
 b. Anna Karenina is braver than Elena Greco.

An adjective like *married* does not admit of modification; people are either married or not. This is why to say that someone is *more married* than another sounds odd. By contrast, an adjective like *brave* does admit of degrees: people can be braver than others. This is true of all the evaluative adjectives on the above list: they all admit of degrees and comparisons.

Gradability has been a topic of much interest in semantics, chiefly because gradable expressions challenge a simple view of semantics on which predicates denote sets of objects (see Burnett 2017; Klein 1991; Solt 2015). Moreover, gradability is connected to traditional topics in philosophy, such as vagueness or context-dependency (Kennedy 2007; Kennedy and McNally 2005). According to that simple view, predicates are functions from individuals to TRUTH VALUES. For instance, *married* is a function that takes an individual, e.g. Anna, as its argument and returns the value true if Anna is married, and false if Anna is not married. Alternatively, the denotation of the predicate *married* can be thought of as the set of married individuals, that is, the set of individuals for which the function *married* gives the value true.

Gradable predicates challenge this picture because many individuals do not fall neatly within (or without) their extension: is a person who is 175cm tall *tall*? Is a 100,000€ apartment *expensive*? For this reason, it is often assumed that gradable predicates denote functions from individuals to points or degrees on a scale, rather than to truth values. For example, an adjective like *tall* is a function that takes an individual, e.g. Anna, as its argument and returns its degree of height as its value, say, 175cm. Calling someone *tall* however, is more than just naming their height. To call someone *tall* is to say that its height has a certain property. Tall people are people whose heights stand out in some way. So the idea that gradable predicates pick out degrees of the relevant property has to be complemented with the idea that when we apply those adjectives to individuals, we are predicating a further property of the relevant degree—that it is significant, maximal or something like that.

These observations apply to evaluative adjectives across the board, although these introduce further complications. Standardly, applying what we said about *tall* to an evaluative adjective like *stupid* would work as follows: classifying any individual as *stupid* involves two operations: first, one needs to figure out an ordering of individuals with respect to their degree of stupidity; and secondly, a threshold degree of stupidity needs to be determined, such that any individual who is more stupid than the threshold counts as *stupid* (Kennedy 2007). There are two difficulties with this picture: the first is that, in contrast to adjectives like *tall*, it is difficult to imagine precise and measurable degrees of stupidity. This makes it difficult both to establish an ordering of individuals as well as to determine any threshold degree on that scale. Secondly, the very criteria of stupidity can and do vary from speaker to speaker, thereby resulting in different orderings. In assessing people's stupidity, you might give great weight to lack of intelligence; while I may consider lack of politeness and etiquette to be more important. Thus, if we are asked to order a set of individuals by their stupidity, our orderings may not coincide (Bylinina 2017; Solt 2018).

Thick and thin evaluative adjectives

An important issue that cuts across the previous debates about evaluatives is the distinction between THICK and THIN evaluative adjectives. Thick evaluative adjectives have evaluative and descriptive (or non-evaluative) meaning. The thin evaluative adjectives, by contrast, are those that have only evaluative meaning. Examples of the first kind are words like *cruel*, *brave*, *generous* or *chaste*. Thin adjectives include all-purpose evaluative terms like *good*, *bad*.

The intuitive contrast lies in the fact that to call someone *cruel*, for example, is to convey something negative about that person: it is to communicate that s/he is bad. But not all bad things are cruel. To say that someone is cruel is to say something more specific, something along the lines of *inflicts unnecessary suffering*. That more specific content however, is not evaluative.

So the intuitive idea is that thick terms like *cruel* have evaluative and descriptive components of their meaning, whereas thin terms like *bad* only have the evaluative component. The evaluative component is somewhat general, while the descriptive component is more specific, thereby “thickening” the meaning of the term. Arguably, one can distinguish thin and thick evaluatives within some of the flavors of evaluatives that we discussed before. For example: *ethical* is a thin moral adjective; while *brave* or *generous* are thick; *beautiful* is a thin aesthetic adjective, while *elegant* is arguably thick. And *advisable* is a thin practical adjective; whereas *dangerous* could be considered a thick practical adjective.

The descriptive element—which we call (*D*) below—is usually taken to be a value-neutral description (for *cruel*, it is something like *inflicts unnecessary suffering*; for *selfish*, it is something like *giving priority to oneself over others*; etc.). The evaluative component—(*E*) below—can be expressed with the aid of thin evaluative terms, as the claim that whatever falls under the descriptive component is good/bad in virtue of it. Thus, the meaning of a sentence like (1.10) can be factored out in the following two components:

- (1.10) John is cruel
 (*D*) John inflicts unnecessary suffering.
 (*E*) Anyone that inflicts unnecessary suffering is bad in virtue of doing so.

From the viewpoint of contemporary philosophy of language and linguistics, the debate surrounding thick terms largely assumes that descriptive and evaluative elements are separable, and then goes on to consider how these elements combine with each other in the meaning of a thick term. More precisely, the debate assumes that the descriptive component (*D*) contributes to the truth-conditions of sentences like (1.10), and then considers the question of how the evaluative component (*E*) is communicated. In particular, the debate revolves around the question of whether the evaluative component is part of the truth-conditions as well, or whether it is conveyed through some other semantic or pragmatic mechanism, such as presupposition, implicature or something else entirely (see Cepollaro and Stojanovic 2016; Kyle 2013; Roberts 2013; Väyrynen 2013, 2017, for some key references).

1.4 Outline of the chapters in this dissertation

This chapter has surveyed a number of contemporary debates in philosophy of language and linguistics surrounding the nature of evaluative language. We started out by tracking the intuitive distinction between facts and values, or between descriptive and evaluative language. After a few general considerations about the words that tend to appear in the latter, namely evaluative adjectives, we turned to consider some of the semantically interesting features of evaluative adjectives. Hopefully, this survey succeeds in providing a bird’s-eye view of the current debate about this class of expressions and can help the reader situate the present dissertation.

It is time now to actually give an overview of what is to come. This dissertation defends a form of non-factualism about evaluatives, which is the view that these expressions do not describe reality. Rather, we use evaluatives primarily to update and coordinate our practical commitments. But non-factualism has not, for the most part, been defended as a view that takes into account the gradability of evaluatives. My purpose is to amend this. Along the way, further related topics are touched upon, such as the phenomenon of subjective or faultless

disagreement or the possibility of using evaluative sentences to both express judgments of value and judgments of fact. Here is an outline of the chapters in this dissertation.

Chapter 2 explores evaluatives from a linguistic as well as metanormative point of view. Non-factualism is presented as a promising account of the meaning of evaluative sentences in spite of the well-known Frege-Geach problem. If non-factualism is to be defended as a linguistically serious proposal however, it needs to be implemented compositionally. This task faces special difficulties in light of the fact that evaluative adjectives are gradable. I argue that the gradability of evaluatives suggests a “sub-sentential” version of the Frege-Geach problem, and that in order to solve this problem one must move away from a standard degree semantics and adopt a delineation semantics for evaluative adjectives.

In **Chapter 3**, a compositional semantics is given for a propositional language introducing a pair of positive and negative unary sentential operators representing positive and negative value, together with binary sentential operators representing equative, comparative and equal value. The semantics is based on Gibbard’s notion of a hyperplan, as this framework offers the most intuitive account of the connection between evaluative concepts and action/motivation (argued for in the previous chapter). And following a delineation approach, the semantics of the binary operators is derived from that of the unary operators. In this way, we preserve the non-factualist insight that evaluative language is primarily used to express simple positive and negative attitudes towards objects of evaluation.

Chapter 4 aims to apply the scalar hyperplan framework presented in Chapter 3 to evaluative adjectives across the board. The chapter is divided in three parts. The first explores how to map the propositional language explored in chapter 3 onto all-purpose evaluatives like *good*, *bad*, and *better*, when these take propositional complements. The second part is devoted to both applying this semantics to uses of those adjectives with other types of arguments. The last part of the chapter is dedicated to exploring further semantic properties of evaluative adjectives, such as the relative-absolute distinction and the internal properties of their scales.

Chapter 5 discusses two notions that have been the focus of much discussion in recent linguistic and philosophical literature, namely the notion of *judge-dependency* and *ordering subjectivity*. In recent experimental work, Stephanie Solt has argued that the spectrum-like nature of the phenomenon of ordering subjectivity calls for a distinction, within the class of subjective adjectives, between multi-dimensional and judge-dependent ones. We supplement her conclusions by arguing that applying her experimental paradigm to moral adjectives suggests a further distinction, within the class of judge-dependent adjectives, between experiential and evaluative adjectives.

Finally, evaluative sentences are by default used to make attributions of value. In the appropriate setting however, certain types of evaluative sentences can also be used to express judgments of fact. In **Chapter 6**, I argue for a *dynamic* implementation of the non-factualist framework presented in previous chapters as offering the best account of such double uses of evaluative sentences. I have two aims: first, I want to diagnose such seldom noticed contrast between descriptive and evaluative uses of evaluative sentences. Secondly, I offer a dynamic model of communication involving evaluative language that aims to capture that contrast.

Chapter 2

Evaluativity in linguistics and philosophy: gradability & non-factualism

Summary

This chapter explores evaluative expressions from a linguistic as well as metanormative point of view. Non-factualism is presented as a promising account of the meaning of evaluative sentences—sentences containing evaluative adjectives—in spite of the well-known Frege-Geach problem. If non-factualism is to be defended as a linguistically serious proposal however, it needs to be implemented compositionally. This task faces special difficulties in light of the fact that evaluative adjectives are gradable. I argue that the gradability of evaluatives posits a “sub-sentential” version of the Frege-Geach problem, and that in order to solve this problem one must move away from a standard degree semantics and adopt a delineation semantics for evaluative adjectives.

2.1 Introduction

As we tried to convey in the introduction, evaluative adjectives (or evaluatives, for short) are perplexing for a number of reasons. Most crucially, they show a particular sensitivity to contextual features, and they are conceptually connected to action and motivation in a way that sets them apart from other expressions of natural language. Their context-sensitivity has been discussed in philosophy and linguistics, insofar as these adjectives belong to superordinate *geni*, such as gradable (Kennedy 2013b), multi-dimensional (McNally and Stojanovic 2017; Sassoon 2013a) or judge-dependent expressions (Bylinina 2017; Lasersohn 2005; Stephenson 2007a). By contrast, their connection to action has been long discussed in metaethics (Björklund et al. 2012; Björnsson et al. 2015), but less so in linguistics and philosophy of language.

From a linguistic point of view, the interaction between gradability and subjectivity is the most salient property that evaluative adjectives share. It is natural, then, to subsume evaluative adjectives under standard accounts of gradable adjectives, the most standard of which is degree semantics. From a metanormative point of view, the most salient feature of evaluatives is their action-guidance. The theories that have proved most promising at dealing with this feature of evaluatives are those that fall broadly under the *non-cognitivist* or *non-factualist* banner.

Interestingly though, non-factualists have for the most part not paid attention to the gradability of evaluatives. I want to argue that this is a shortcoming. These authors tend to cash out the

meaning of evaluatives in terms of the expression of *non-gradable* attitudes of outright support or rejection, approval or disapproval, *PRO* or *CON*, etc. And it is not obvious what to say about sentences like *stealing for food is better than murdering for food*, which on the face of it do not convey such simple positive and negative attitudes.

Regardless, it seems that these two theoretical frameworks, degree semantics and non-factualism, could complement each other: degree semantics might employ conceptual tools from non-factualism in order to figure out the scalar particularities of evaluatives, and non-factualists should take into account the gradability of evaluatives (and all that comes with it) thanks to degree semantics. In this chapter I argue that things are not so easy—that there is an important tension between those frameworks.

To see this, note that the main obstacle for non-factualism about evaluatives is not their gradability—it is, of course, the Frege-Geach problem: the problem of determining the meaning of complex constructions containing evaluative sentences in light of the revisionary semantics that non-factualists assign to those sentences. The tension between non-factualism and degree semantics appears when we appreciate that the gradability of evaluatives gives rise to a different “incarnation” of the Frege-Geach problem. Whereas the traditional problem—the “supra-sentential” Frege-Geach problem—is a problem concerning how evaluative sentences embed in larger constructions, this version of the problem concerns how evaluative sentences are semantically composed in the first place—the “sub-sentential” Frege-Geach problem. Put simply, any view about the semantics of gradable expressions is committed to the fact that gradability arises from a combination of semantic factors. Thus, the meaning of any gradable adjective ought to be traced back to such factors and the way in which they are combined. The problem is that the characteristic semantic features of evaluatives (which prompt non-factualists to propose their revisionary semantics) cannot be straightforwardly traced back to such factors in the way that degree semantics describes them. In light of this, I defend that non-factualism is more amicable to a different framework for scalar expressions, namely Delineation semantics. That is the framework that the dissertation will be couched upon.

The chapter is organized in 6 sections. §2.2 tackles evaluatives from a linguistic point of view. §2.2.1 presents the basic scalar properties of evaluatives and how these properties are cashed out in standard degree semantics for gradable adjectives (§2.2.2). §2.2.3 presents the various conceptions of evaluativity discussed in linguistics, which will be relevant later on. §2.3 tackles evaluatives from a philosophical, or more specifically metanormative perspective. §2.3.1 discusses the action-guiding property of evaluatives and the way in which non-factualism account for it (§2.3.2). Section §2.4 focuses on the single most important problem that non-factualists face, namely the Frege-Geach Problem. We will discuss how this objection was developed historically, and then we will briefly present how hybrid and pure non-factualists proposals aim to overcome that problem. Section §2.5 presents what we will dub the “sub-sentential” Frege-Geach problem. The comparative form of evaluative adjectives are presented as a challenge for non-factualism, and this leads to a further generalization, namely a correlation between the presence of the positive form of evaluative adjectives and the outright *PRO*- and *CON*-attitudes that non-factualists assign as the meaning of evaluative adjectives (§2.5.1). We will argue that accounting for that generalization causes problems for degree semantics (§2.5.2). Finally, section §2.6 presents the framework that the proposal defended in this dissertation will be couched upon, namely Delineation Semantics. §2.7 will recapitulate and conclude.

2.2 Evaluatives from a linguistic point of view

In this section, we present the properties that make evaluative adjectives special from a linguistic perspective. These have to do with the interaction of their scalar and their subjective properties (§2.2.1). Then, in section §2.2.2 we will see how to accommodate those properties in a standard degree semantics. In §2.2.3 we will say something about the way in which the linguistics literature has understood the notion of *evaluativity* and how it differs from the usage of that term in philosophy.

2.2.1 Gradability & subjectivity

From a linguistic standpoint, evaluatives belong to a class of gradable predicates that give rise to (what we will call) “subjectivity effects” both in positive and comparative form.

Gradability

First, evaluatives are GRADABLE adjectives. Gradability is attested by the admissibility of adjectival modifiers, such as adverbs (*very*, *a little*) or measure phrases. Consider the contrast in acceptability between the first and the second set of examples:

- (2.1) a. The theatre was a little empty.
- b. My tailbone was completely bent inwards.
- c. That building is very tall.
- d. The soup was too salty.
- (2.2) a. All weapons of mass destruction are (?? a little) nuclear.
- b. It’s (?? completely) freezing outside.
- c. That man is (?? very) dead.
- d. They made the square (?? too) hexagonal.

Adjectives like *empty* or *tall* denote properties that can be had to various degrees and quantities, which is why we can say things like *a little empty* or *very tall*. By contrast, adjectives like *dead* or *hexagonal* denote properties that things either have or lack.¹ Evaluative adjectives admit adjectival modifiers (with the notable exception of measure phrases, but more on that later), and they belong to the class of gradable predicates:

- (2.3) Isn’t that a little cruel to treat your dog like some kind of drooling Swiss Army knife?
- (2.4) This is a completely unethical company.
- (2.5) That building is very beautiful.
- (2.6) The acting wasn’t too good.

¹This is not to say that sentences like (2.2) are absolutely unintelligible; but in order to recover a meaning we need to do some interpretative work. For instance, someone could mean by *they made the square too hexagonal* that they gave it a hexagonal shape that was too perfectly delineated or too regular.

Subjectivity (I): the positive form

Many gradable adjective are SUBJECTIVE. This means that whether they apply to an individual seems to be intuitively a matter of subjective opinion rather than a matter of objective fact. Two standard linguistic tests for subjectivity are whether adjectives can figure in so-called subjective or faultless disagreements and whether they are embeddable under subjective attitude verbs, most notably *find*.

A SUBJECTIVE or FAULTLESS DISAGREEMENT is a dialogue in which two speakers hold divergent opinions about a certain matter and yet both of them appear equally justified or entitled to that opinion.² To see this, contrast the following dialogues:

- (2.7) a. Natalia: Fede is tall. *≈ subjective*
b. Matheus: No, he is not.
- (2.8) a. Natalia: The tank is empty. *≈ not subjective*
b. Matheus: No, it is not!

Whereas in (2.7) it seems perfectly possible for Natalia and Matheus to both be entitled to having different opinions about what counts as *tall*, that does not seem to be true of a dialogue like (2.8), where it is clear that, strictly speaking, one of them must be mistaken.³

Evaluative adjectives behave like *tall* in this respect:

- (2.9) a. Natalia: Eating animals is unethical. *≈ subjective*
b. Matheus: I disagree, I see nothing wrong with it.

It seems clear that this is a matter about which Natalia and Matheus can disagree and at the same time being entitled to having divergent opinions.

² The tags ‘faultless’ and ‘subjective’ are both problematic for different reasons. Regarding ‘faultless’: note, first, that the faultlessness of any such disagreement is not for the participants in the disagreement to perceive: if you think that the wine we just had was good, I may very well think that you are mistaken in your appreciation (Neftalí Villanueva, *p.c.*). Faultlessness is a property to be perceived from a bird’s eye perspective, so to speak (Stojanovic 2007; Umbach 2016). Secondly, the conceptual possibility that any genuine disagreement could really be *faultless* has been challenged (Glanzberg 2007; Stojanovic 2007). We do not need to take a stance with respect to that debate; rather, we are happy to observe that, among dialogues that sound reasonably natural and look like disagreements, some of them give the impression that one of the participants has somehow gotten the facts wrong, and some of them give the impression that speakers are entitled to having different opinions (see Kölbel 2003; Lasersohn 2005; MacFarlane 2014; Stephenson 2007a for classic discussions, as well as the papers in García-Carpintero and Max Kölbel 2008. For recent experimental work on the scope of this phenomenon see Solt 2018 and Chapter 5 of this dissertation; see also Stojanovic 2017a for a recap of the debate between contextualism and relativism and the role played by the notion of faultless disagreement). Regarding ‘subjective’: characterizing certain predicates as ‘subjective’ or as giving rise to ‘subjectivity effects’ should not be conflated with proposing a *subjectivist* semantics for those predicates. ‘Subjectivity’ here refers to an observable set of linguistic properties; ‘subjectivism’, by contrast, is a particular family of metaethical proposals. In one standard construal, subjectivism is the view that moral predicates describe the moral feelings of the speaker who uses them. That is not the view defended here. Furthermore, the family of views known as *expressivist* or *non-factualist* (under which this dissertation does fall), was from the very beginning characterized in opposition to subjectivism: while subjectivists claim that moral terms describe the moral feelings of the speaker, expressivists claim that moral terms express such feelings (see e.g., Ayer 1946, pp. 107–8; Gibbard 2003, p. 85 for discussion of this point). More on this in §2.3.2.

³ Note that, in dialogue (2.8), speakers could be arguing about whether the tank is empty relative to some implicit purpose. For example, they might be going somewhere very close, and there might still be a little gas in the tank, and they are discussing whether they have *enough* to reach their destination. This interpretation is possible, but it is a loose use of *empty* that we want to abstract away from.

The second test for subjectivity that we mentioned is embeddability under SUBJECTIVE ATTITUDE VERBS, like *find*.⁴ Note that some gradable adjectives are fine when embedded under *find*, while others are comparatively much less natural:

- (2.10) I find that suitcase heavy.
- (2.11) I find the exam difficult.
- (2.12) ?? I find the antenna straight.
- (2.13) ?? I find the door open.

Evaluatives are in general admissible under *find*—although there are differences: evaluatives that clearly involve some kind of appreciation of sensorial qualities, such as predicates of personal taste (PPTs), are more natural under *find* than, say, moral adjectives (or even aesthetic adjectives, as argued by McNally and Stojanovic 2017, p. 29; (2.15) is their example (18-a)).⁵

- (2.14) I find the coffee here tasty.
- (2.15) ? I find Miró's mosaic on the Rambles mediocre.
- (2.16) ? I find that custom unethical.

Subjectivity (II): the comparative form

A further distinction within subjective adjectives is the following: some of them give rise to subjectivity effects both in the comparative form and in the positive, unmodified form; while some of them are subjective only in the positive form. Consider the contrast between *tall* and *difficult* in this respect:

- (2.7) a. Natalia: Fede is tall. ≈ subjective
b. Matheus: No, he is not.
- (2.17) a. Natalia: Fede is taller than Stephen. ≈ not subjective
b. Matheus: No, he is not.

As we said before, speakers can subjectively disagree about whether someone is tall. But they cannot subjectively disagree about whether someone is tall-*er* than someone else. Things are different with *difficult*, where both positive and comparative form can give rise to subjective disagreement.

- (2.18) a. Natalia: The homework was difficult. ≈ subjective
b. Matheus: Not at all!
- (2.19) a. Natalia: Today's homework was more difficult than yesterday's. ≈ subjective
b. Matheus: Not at all!

⁴ See Fleisher 2013; Franzén 2018b; Kennedy 2013b; Sæbø 2009; Stephenson 2007b; Umbach 2016, 2019. See also McNally and Stojanovic 2017 for some data on other subjective attitude verbs, such as *look* or *sound*; and Kennedy and Willer 2016 for extensive data on *find* vs. *consider*.

⁵ McNally and Stojanovic 2017 actually claim that *find*-embeddings *anti-select* for evaluativity. They argue that *find* introduces an experiencer requirement, and that that requirement ends up coercing an experiencer reading of otherwise evaluative adjectives. It is not clear to me in what sense the experiencer reading trumps the evaluative one, however. But we will not go into that discussion.

Embeddability under *find* follows suit: *taller* cannot embed under *find*, while *more difficult* can.⁶

(2.20) ?? I find Fede taller than Stephen.

(2.21) I find today's homework more difficult than yesterday's.

Evaluative adjectives behave in this respect like *difficult* and unlike *tall*:

(2.22) a. Natalia: Eating animals is less ethical than wearing them. \approx subjective
b. Matheus: That's not true.

(2.23) I find eating animals less ethical than wearing them.

Bylinina (2017) calls this property *scalar variation*, whereas Solt (2018) calls it *ordering subjectivity*; we will follow Solt's terminology. As we will see, ordering subjectivity suggests that evaluative adjectives give rise to variable orderings.

In sum: evaluative adjectives are a class of gradable adjectives that give rise to subjectivity effects both in positive and comparative form. We turn now to the standard way of accommodating these observations in degree semantics.

2.2.2 Evaluatives in degree semantics

In degree semantics, gradable adjectives are assigned measure functions as their lexical meaning. A measure function is a function μ from some semantic type (most commonly individuals) to degrees. DEGREES are abstract representations of measurement that are introduced as a primitive semantic type (along with individuals, truth-values, etc.).⁷ The adjective *tall* for instance, maps individuals to degrees of height.⁸ Let us assume an extensional and context-sensitive semantics. Then, using lambda notation, we can write the meaning of *tall* as follows:

$$(2.24) \quad [[\text{tall}]]^c = \lambda x. \mu_{\text{height}}(x)$$

This is a function that takes an individual, say Ann, and returns its height, say 165cm, at any context.

The default assumption in this framework is that all gradable expressions have this type of lexical meaning, and hence that evaluatives do so as well. An evaluative adjective like *good*, then, would have as its denotation a function from whatever its semantic argument is (let us say individuals for simplicity, although we will eventually argue that evaluative adjectives are, at core, propositional operators) to degrees of the relevant property.

$$(2.25) \quad [[\text{good}]]^c = \lambda x. \mu_{\text{goodness}}(x)$$

⁶ Kennedy 2013b points out that dimensional adjectives such as *heavy* or *long* have “qualitative” interpretations under which they give rise to subjectivity effects also in the comparative. One can say *I find this bag heavier than that one*, even though there is a fact of the matter about which bag really is heavier.

⁷ See Cresswell 1976; Kennedy 2007; Kennedy and McNally 2005 for classic references.

⁸ Alternatively, one may take adjectives to denote relations between individuals and degrees (Heim 2000; Kennedy 2013a) but for our purposes the difference between these two approaches will not matter.

However, we have seen that gradable adjectives differ with respect to their subjectivity: some give rise to subjectivity effects and some do not; and some give rise to subjectivity effects only in the positive form, while some do so both in positive and comparative form. In general, subjectivity is assumed to arise as a result of the presence of a contextual parameter (Kennedy 2013b). But to be more precise we need to say something about the difference between the comparative and the positive form.

Positive form subjectivity

Take an adjective like *tall*. *Tall* is subjective in its positive, but not in its comparative form. The semantics of the comparative is straightforward: comparatives simply state a relation between degrees. A sentence like (2.17a), for example, says that Fede’s degree of height is superior to Stephen’s:

$$(2.17a) \quad [[\text{Fede is taller than Stephen}]]^c = \mu_{\text{height}}(\text{Fede}) > \mu_{\text{height}}(\text{Stephen})$$

The reason why the comparative *taller* does not give rise to subjectivity effects is accounted in this framework by the fact that its semantics is context-invariant.

When we turn to the positive form things get more complicated. Notice that the meaning of *Ann is tall* cannot simply be the result of feeding *Ann* as the argument of $\lambda x. \mu_{\text{height}}(x)$, because the semantic value of that sentence at a context should be a truth-value and not a degree of height. In other words, the predicate *is tall* cannot have as its meaning a measure function—it has to be a predicate, that is, a function from individuals to truth-values. In degree semantics, the way to “transform” a measure function into a regular predicate is to combine it with a function that predicates some property of the degree that the measure function maps its argument to. And depending on what property that is, the positive form of an adjective is going to turn out to be subjective or not.

We saw that certain gradable adjectives are subjective in the positive form (*tall*, *difficult*) while some are not (*empty*, *straight*). With exceptions, this contrast correlates with the distinction between relative- and absolute-standard gradable adjectives:⁹ RELATIVE-STANDARD adjectives are gradable adjectives that, in the positive form, attribute to their arguments a degree (on the relevant scale) *equal to or greater than a contextually determined threshold*. By contrast, ABSOLUTE-STANDARD adjectives like *empty* or *straight* are those whose positive form predicates that their argument *possesses a degree that is equal to or greater than the maximum [or minimum] endpoint degree* (of the relevant scale). The reason why relative- but not absolute-standard adjectives give rise to subjectivity effects is because the property associated with the positive form is contextually determined only in the former case.

When Nathalia and Matheus argue about whether Fede is tall, they might have different threshold for tallness in mind. This offers an intuitive explanation for why we take them to be having a subjective disagreement. But when they argue about whether the gas tank is empty, given that *empty* means having no volume at all, only one of them can be right—the tank is either empty or not, and there is no context-sensitive issue that they could be discussing or negotiating (barring loose talk).¹⁰

⁹ The *locus classicus* of this distinction is Kennedy and McNally 2005. We come back to it in §4.5.

¹⁰ To avoid distraction from the main argument, we illustrate here the rest of the composition process for *tall*: the property denoted by the positive form is assumed to be semantically contributed by a silent morpheme, *POS*,

Find-embeddability should work in a similar way—although the jury is still out.¹¹ A natural view is to think that *find* is semantically sensitive to a contextual parameter like the one that the positive form of relative-standard adjectives have, and that a *find*-sentence is acceptable only if the predicate that *find* combines with is appropriately context-sensitive.¹² However, *find*-embeddability is notably unstable cross-linguistically (Umbach 2019); and for this reason we will not rely on it as a criterion for subjectivity in what follows.

Ordering subjectivity

What about adjectives that give rise to subjectivity effects also in the comparative form? So far, we have seen that certain gradable adjectives, namely relative-standard adjectives, have a contextual parameter in the positive form whose value speakers can subjectively disagree over, and which licenses embedding under *find*. Comparatives lack this contextual parameter, and so are predicted to not be subjective. This prediction is borne out for *tall*, but fails for *difficult*. Why? The standard answer is that the lexical meaning of adjectives like *difficult*, that is, the measure functions that they denote are themselves context-sensitive. For this reason, in what follows we will call this class of adjectives LEXICALLY or ORDERING-SUBJECTIVE.¹³

Formally, the idea is that there is a contrast between the lexical meaning of *tall* and the lexical meaning of *difficult*: whereas the former returns the same degree at any context of use, the latter does not. We can represent this by indexing the measure function to the context:

$$(2.24) \quad \llbracket \text{tall} \rrbracket^c = \lambda x. \mu_{\text{height}}(x)$$

$$(2.26) \quad \llbracket \text{difficult} \rrbracket^c = \lambda x. \mu_{\text{difficulty}}^c(x)$$

The degree of height of an individual is invariant across contexts, while the degree of difficulty of something can change from context to context. More precisely, it changes depending on

with the following meaning:

$$(2.1) \quad \llbracket POS \rrbracket^c = \lambda g. \lambda x. g(x) \geq d_t^c$$

POS takes as its first argument a measure function (written with variable *g*) and predicates a property of whatever degree is assigned to the individual in its second argument, namely that such degree is greater than or equal to a *contextually determined threshold* on the relevant scale (that is what d_t^c stands for). However, there is more to be the meaning of *POS* than an arbitrary specification of a degree on the relevant scale; see Graff 2000, who characterizes *POS* as representing the property of *being significantly greater than*—rather than *at least as great as*—the degree that is normal, or typical, for the relevant property. The regular predicate *is tall* results from applying the lexical meaning of *tall* as the argument of *POS*:

$$(i) \quad \llbracket \text{is tall} \rrbracket^c = \llbracket POS_{rel} \rrbracket^c(\llbracket \text{tall} \rrbracket^c) = \lambda g. \lambda x. g(x) \geq d_t^c(\lambda x. \mu_{\text{height}}(x)) = \lambda x. \mu_{\text{height}}(x) \geq d_t^c$$

When an individual is supplied to this function, the result is a sentence with the following truth conditions:

$$(ii) \quad \llbracket \text{Ann is tall} \rrbracket^c = \mu_{\text{height}}(\text{Ann}) \geq d_t^c$$

That sentence is true just in case Ann's height surpasses the contextually determined threshold of height.

Now the question arises: do things work in the same way with absolute adjectives? Because if they do, then one should expect their positive form to be context-sensitive and subjective too. Empirically however, that is not the case. Kennedy (2007, §4) treats extensively this issue, so we refer the reader to that discussion for further reference. See also Burnett 2017, 69 and ff.

¹¹ See n.4.

¹² This is, very roughly, the proposal in Sæbø 2009. But compare McNally and Stojanovic 2017, who argue that *find*-embeddability might depend on more factors beyond semantic type.

¹³ Klein 1980, §3.3 calls them *non-linear*.

who is assessing the difficulty at a given context. In turn, this allows for context-sensitive comparisons:

$$(2.19a) \quad \llbracket \text{Today's homework was more difficult than yesterday's} \rrbracket^c = \mu_{\text{difficulty}}^c(\text{today's exam}) > \mu_{\text{difficulty}}^c(\text{yesterday's exam})$$

In one context, the two exams might be measured by a scale of difficulty that assigns to today's homework a higher degree than yesterday's; but a different context may determine a scale of difficulty that assigns completely different degrees of difficulty to each homework. Suppose, for instance, that today's exam was heavier on history but lighter on math. The disagreement between Natalia and Matheus (2.19) could seem subjective due to the fact that, whereas Natalia is better at math, Matheus is better at history, and each of them may order the two exams differently by their difficulty. By contrast, note that Natalia and Matheus cannot each have their own height scale, which is why the disagreement in (2.18) is not subjective.

As we saw, evaluatives largely fare like *difficult*, in that they give rise to subjectivity effects both in the positive and the comparative form. The standard way to account for this would be, then, to assign to them context-sensitive measure functions. Just as what we have said about *difficult*, the subjective character of disagreements like (2.22) (repeated here) would arise as a result of the fact that the scale of *ethical* is context-sensitive. A similar account might be given for *find*-embedding, i.e., (2.23).

- (2.22) a. Natalia: Eating animals is less ethical than wearing them. \approx subjective
 b. Matheus: That's not true.

- (2.23) I find eating animals less ethical than wearing them.

Multidimensionality & experiencer-sensitivity

If we wanted to subsume evaluative adjectives under this model, evaluatives would be classified along with any adjective that can give rise to subjectivity effects in positive and comparative form. But note that, whereas positive-form subjectivity is generally taken to be uniformly caused by the positive form morpheme *POS* (see n.10 in this chapter), ordering subjectivity can arise for a variety of reasons.

Some ordering-subjective adjectives, for example, are multidimensional (McNally and Stojanovic 2017; Sassoon 2013a, 2016). MULTIDIMENSIONAL adjectives are adjectives that integrate different respects or dimensions. The pair *<healthy, sick>* is a paradigmatic example: one can be healthy or sick with respect to various dimensions or respects, such as blood pressure, cholesterol or blood sugar level. By contrast, there is but one dimension associated with *tall* (i.e., height).

Following Sassoon, the main available linguistic tests for multi-dimensionality are the admissibility of “dimension-accessing” operators and modifiers, such as the PPs *with respect to ...* and *in some/most/every respect(s)*. Contrast *healthy* with *tall*:

- (2.27) a. Chris is healthy in {every/most/three/some} respect/way(s).
 b. In what respect/way(s) is Chris healthy?
 c. Chris is healthy except for blood pressure.

- (2.28) a. ?? Chris is tall in {every/most/three/some} respect(s)
b. ?? In what respect way(s) is Chris tall?
c. ?? Chris is tall except for...

Might evaluative adjectives just be a class of multidimensional adjectives? It is not clear, and authors are not of one mind about this. For example, McNally and Stojanovic 2017 hypothesize that aesthetic adjectives are multidimensional; but other evaluative adjectives certainly do not pass these tests with merits—even though they fare better than dimensional adjectives (the first set of examples is taken from Solt 2018, p. 72):

- (2.29) a. The chili was tasty in {every/?most/??three/some} respect/way(s).
b. In what respect/way(s) was the chili tasty?
c. The chili was tasty except for...
- (2.30) a. Palm oil is unethical in {every/most/three/some} respect/way(s).
b. In what respect/way(s) is palm oil unethical?
c. ?? Palm oil is unethical with respect to ...?

Depending on the choice of adjective, “dimension-accessing” operators and modifiers are more or less acceptable. And it seems intuitive to think that different dimensions are involved in determining the extension of evaluatives. However, given that evaluatives do not pattern uniformly with respect to these tests, we are hesitant to simply adopt the multidimensional view.

Moreover, some of those sentences can easily be interpreted in ways that do not suggest that the relevant adjective is multidimensional. For example, it is very natural to take the dimension-accessing operators in (2.30) to be accessing *reasons* to think that palm oil unethical—rather than different ethical dimensions.

Alternatively, some ordering-subjective adjectives are clearly experiential or qualitative (*salty*, *soft* or *colorful*). EXPERIENTIAL adjectives are assumed to involve a subjective, first-hand experiencer. For instance, whether a food is *salty* is not just a matter of its salt content, but of how salty it *tastes* to an experiencer or judge.

An informal test for this property is to see whether a sentence featuring these adjectives invites an inference that the speaker has had some kind of first hand experience.¹⁴

- (2.31) a. The soup is salty. \leadsto the speaker tried the soup
 b. The chair is soft. \leadsto the speaker tried the chair
 c. The room was colorful. \leadsto the speaker saw the room

Evaluative adjectives fare quite differently with respect to this test. While some involve a first-hand experience (most clearly PPTs, but arguably also some aesthetic adjectives, c.f. Franzén 2018a *pace* McNally and Stojanovic 2017), some seem not to require any such thing—notably moral adjectives:

- (2.32) a. The entrée was tastier than the main. \leadsto the speaker tried the food
b. Her house is beautiful. \leadsto the speaker saw the house

¹⁴ Compare McNally and Stojanovic 2017, p. 27, for an alternative test for the presence of an experimenter.

- c. What she did was cruel. ↗ the speaker experienced what she did
d. Volunteering is a generous thing to do. ↗ the speaker has volunteered

Another test for the presence of an experiencer is the admissibility of so-called judge PPs (*tasty for/to me/Ann/Bill*; c.f. Glanzberg 2007; McNally and Stojanovic 2017; Stephenson 2007a a.o.), relative to which evaluatives again give mixed results.

We come back to the question of experientiality in Chapter 5, but for now let us conclude that, according to these linguistic tests, even though we can safely describe evaluative adjectives as a class of ordering subjective gradable adjectives (in virtue of the fact that they are subjective in the comparative form), we cannot safely assimilate them to the usual categories that linguists distinguish within these adjectives. As we will argue, evaluative adjectives are a proper subclass of the adjectives that we have been calling lexically or ordering-subjective. The picture that emerges from this discussion is something like this:

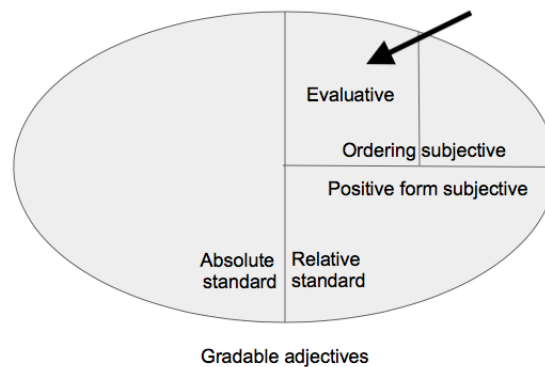


Figure 2.1: The place of evaluative adjectives within gradable adjectives.

Linguistic enquiry into evaluatives, therefore, points to a necessary property of evaluative adjectives, namely ordering subjectivity. In §2.3 we will turn to a necessary property of evaluatives, namely action-guidance.

Before that however, degree semantics faces a general conceptual difficulty when it is applied to adjectives whose scales involve properties that lack units of measurement. This is the case with most ordering-subjective adjectives (including evaluatives). The difficulty is best seen by considering the positive form. Recall that, according to degree semantics, the positive form of an adjective involves a triple operation: first, we assign a degree of the relevant property to an object; secondly, a threshold degree on the same scale is determined; and finally, those two degrees are compared, and the positive form can be predicated of the object just in case the former degree is equal to or superior to the latter.

Given evaluative adjective's lack of measurement units, it is difficult to conceptually isolate those semantic operations. The idea that claiming that a painting is beautiful involves determining two different degrees of beauty (one that is had by the painting, and another that corresponds to the threshold) and then comparing them seems bizarre. But that picture seems much more natural when we think of *tall*, since we have perfectly established ways of determining and comparing degrees of height.

This is not meant as a knock-down argument against degree semantics; my purpose is simply to point out that this is a framework that is much more appropriate for dimensional, measurable

adjectives like *tall*, *heavy* or *wide* than for adjectives like *good*, *beautiful* or *cruel*.¹⁵ Moreover, we have not yet gotten to discussing the property that unifies evaluatives at all, namely their action-guidance. And when we get there, it will be argued that degree semantics does face a difficulty. But before that, let us say something about how the notion of evaluativity has been traditionally understood in linguistics.

2.2.3 The notion of evaluativity in linguistics

Evaluativity means different things in philosophy and in linguistics. As we have seen in Chapter 1, philosophical talk of evaluative terms, concepts or properties is most often associated with metanormative inquiry about the good, the bad or the ugly, or else with the contrast between the realms of the descriptive *versus* the normative. And more specifically, in this work the term *evaluative* is used to pick out a set of gradable and subjective adjectives that feature a special type of connection to action and motivation.

In linguistics however, *evaluativity* is used in various ways. First, *evaluativity* has a more specific, technical usage. This notion comes from Bierwisch 1989, who called it *normbezug* (“*norm-relatedness*”), and has been more recently studied by Rett (2007). In this sense of evaluativity, a sentence containing a gradable adjective is *evaluative* just in case it makes reference to a contextually determined threshold on the relevant scale.

Evaluativity is attested, following Rett, by considering whether a sentence containing any form of a gradable adjective invites an inference to the positive form of the adjective. One reason for this is that the positive form of a gradable adjective involves the morpheme *POS*, which makes reference to a contextually determined threshold (however, evaluativity is not accounted in its entirety by *POS*, as observed in Rett 2007, p. 212).

Under this interpretation, it is not certain gradable adjectives which are evaluative, but rather it is certain forms of *any* gradable adjective which are evaluative. Here are some examples, using *tall* and *short* (these observations are taken from Rett 2007, p. 213; a thorough exploration of this inference pattern can be found in Bierwisch 1989):

- | | | |
|--------|------------------------------------|----------------|
| (2.33) | a. Ann is tall. | ↗ Ann is tall |
| | b. Ann is taller than Bill. | ↗ Ann is tall |
| | c. Ann is as tall as Bill. | ↗ Ann is tall |
| | d. How tall is Ann? | ↗ Ann is tall |
| | e. Ann is too tall for her pants. | ↗ Ann is tall |
| (2.34) | a. Ann is short. | ↗ Ann is short |
| | b. Ann is shorter than Bill. | ↗ Ann is short |
| | c. Ann is as short as Bill. | ↗ Ann is short |
| | d. How short is Ann? | ↗ Ann is short |
| | e. Ann is too short for her pants. | ↗ Ann is short |

At first glance, this and our sense of *evaluativity* might seem completely unrelated, but it is very much to the contrary. In the coming sections, we return to the *evaluative* profile—in Rett’s sense—of *evaluative* adjectives—in our sense.

¹⁵ Compare Cresswell 1976, 280 and ff.

Before we move on, let us note two other uses of *evaluative*. Note that Bierwisch also uses *evaluative* but does not mean *norm-related*. Rather, he uses that term to pick out a class of adjectives, those which are not dimensional (also Umbach 2016; Bylinina 2017, §5). Similarly, Zehr & Égré (2018, p. 35)—claiming to follow Kennedy 2013b, p. 264, although the latter has a slightly different characterization—characterize evaluative adjectives as those that embed under *find* in the comparative. Arguably, for all these authors, *evaluativity* is roughly equivalent to what we have been calling *subjectivity* (that qualification is needed because, as we have seen, some forms of the same adjective may embed under *find* and others may not, c.f., *tall* and *taller*).

These are all broader senses of *evaluative* than what is meant in this dissertation, as we will see in the next section. Before we move on, a closer notion of evaluativity in the linguistics literature to the one that is adopted here has been defended, for example, by McNally & Stojanovic, who write: ‘we consider as evaluative those adjectives that carry with their use an implication of a *positive or negative attitude or evaluation* on the part of the speaker’ (p. 28 2017, my emphasis).

2.3 Evaluatives from a metanormative point of view

By *evaluative* adjectives, we mean a proper subclass of lexically or ordering-subjective adjectives, namely those that are action-guiding. But this notion comes from the philosophical literature, so we turn now to the way in which evaluatives have been studied in philosophy, or more precisely, in the metanormative literature.

2.3.1 Action-guidance

As was said in Chapter 1, evaluative adjectives are characterized by their connection to action and motivation. The metanormative literature has long expanded on this feature of normative vocabulary (see e.g. Björklund et al. 2012; Blackburn 1998; Dreier 1990, 2009; Finlay 2004; Gibbard 1990; Hare 1952, *a.m.o.*). The idea is that, when a rational speaker makes a positive [negative] value attribution about an object, they should *eo ipso* be inclined to act in favor of [against] that object. In other terms, attributing positive [negative] value invites the inference that one will orient one’s action towards [against] the object under evaluation.

For instance, if I judge football to be a *great* sport, I am expected to be willing to watch it, play it or follow it somehow. By contrast, if I judge football to be a *popular* sport, there is no comparative expectation about what my practical attitudes towards football will be. This is what Stevenson 1937, p. 16 called the “magnetism” of evaluative vocabulary—the good is attractive; the bad is repulsive.

Among the lexically subjective adjectives, some are action-guiding and some are not. For example, *difficult*, *soft* or *salty*, although they are lexically subjective (they give rise to subjective disagreement and embed under *find*), they do not invite an inference that the speaker will behave in any particular way. We will represent this inference using \sim (although we remain neutral for now on the nature of the specific inferential mechanism that is in place), where *support/reject* can be taken as placeholders for any positive or negative attitude or stance:

- | | | |
|--------|--------------------------------|---|
| (2.35) | a. Today's exam was good. | ↪ the speaker supports today's exam |
| | b. The floor is still ugly. | ↪ the speaker rejects the floor |
| | c. The soup is awful. | ↪ the speaker rejects the soup |
| (2.36) | a. Today's exam was difficult. | ↪ the speaker supports/rejects today's exam |
| | b. The floor is still soft. | ↪ the speaker supports/rejects the floor |
| | c. The soup is salty. | ↪ the speaker supports/rejects the soup |

Evaluative sentences invite the inferences in (2.35), while adjectives such as *difficult* or *soft* do not. Nonetheless, it is easy to imagine that, with a bit of context, a speaker of any of the sentences in (2.36) could be intending to communicate an attitude of support or rejection. But we cannot know what that attitude is or would be in virtue of the meaning of the sentence alone.

The presence of such inferences is difficult to attest in traditional linguistic terms, but we may rehearse the following two tests.

The *although*-test

First, consider the contrast within each pair of sentences:

- | | |
|--------|---|
| (2.37) | a. Matheus thinks that volunteering is virtuous, although he doesn't have any intention or plan whatsoever of supporting, promoting or doing it. |
| | b. ?? Natalia thinks that volunteering is unpaid although she doesn't have any intention or plan whatsoever of supporting, promoting or doing it. |
| (2.38) | a. Matheus thinks that donating is unethical although he belongs to a charity. |
| | b. ?? Natalia thinks that donating is widespread although she belongs to a charity. |

While the *a*-sentences are acceptable, the *b*-sentences sound marked. Why is this so? We venture that the oddness is due to the connective *although*, which suggests a contrast between its arguments. To describe an individual as having a certain evaluative stance—thinking that an activity is virtuous or unethical—while at the same time attributing to her a practice, or a practical intention, that is markedly incoherent with that evaluative stance constitutes good grounds for such contrast. But regular beliefs about matters of fact—thinking that an activity is unpaid or widespread—do not contrast with practices or practical intentions in any immediate way.

Thus, the connective *although* marks an in-existent contrast in the *b*-sentences, and the sentences are awkward for that reason. But in the *a*-sentences, the contrast between attributing a certain evaluative belief and practical intentions that are somehow incoherent does license the connective.¹⁶

¹⁶ The significance of this test should not be overstated, however: background or contextual information could very well fill the missing pieces that would make the *b*-sentences acceptable: for example, in a context in which it is common ground that Natalia only ever enjoys very exclusive activities, (2.38b) would be a perfectly natural thing to say. In absence of such contextual clues however, (2.38b) is odd.

Substitution *salva validitate* in practical inferences

Secondly: cogent practical inferences (that is, inferences whose conclusions are courses of action, or at least attributions of intentions to engage in some course of action) require at least one premise concerning the practical stance of their agent. For example, we accept chains of inferences like the following:

- (2.39) a. It's raining.
b. You are going out.
∴ Take an umbrella.

But it is clear that, as it stands, that inference is only cogent as an enthymeme whose elided premise is something like *you do not want to get wet*, or something like *it is a bad idea to get wet*:

- | | | | |
|--------|-------------------------------|--------|--------------------------------|
| (2.40) | a. It's raining. | (2.41) | a. It's raining. |
| | b. You are going out. | | b. You are going out. |
| | c. You don't want to get wet. | | c. It's a bad idea to get wet. |
| | ∴ Take an umbrella. | | ∴ Take an umbrella. |

The important thing to notice is that no premise void of some reference *either* to the desires, intentions or plans of the agent, *or* to what is valuable, could do the trick and make that inference cogent. But note that there is a crucial difference between (2.40) and (2.41): the former introduces as its (elided) premise an attribution of a propositional attitude (a desire), while the latter introduces an evaluative sentence, but says nothing about desires, plans or intentions. Regardless, the inference in (2.41) goes through.

In view of this, we may hypothesize that evaluative sentences are the only type of non-attitudinal premise that can make a chain of practical inference cogent.¹⁷ As the examples (2.40)-(2.41) show, in a practical inference an action-oriented propositional attitude ascription can be substituted for an evaluative sentence *salva validitate*.

Before we move on, let us point out two potential objections. First, one might complain that action-guidance is less robust than what we have said here. Restricting ourselves to the moral case for example, do we not constantly behave in ways that go against our moral beliefs? In other words, is Matheus' plight in (2.37a) not a rather common one? This is the classic problem of Weakness of the Will, which is a general problem for accounts of normative thought which claim there to be a tight connection between normative thought and action (see Davidson 1969 for a classic statement).

Discussing that problem would take us too far afield, but let us repeat that we are purposefully remaining agnostic about the nature, and thus about the exact *strength*, of the connection between evaluative thought and action. It could be a semantic entailment (perhaps a *sui generis* "practical entailment"), a presupposition, an implicature or something else entirely. Moreover, it is important to remember that practical inferences are in general *defeasible*. An agent who

¹⁷ Thanks to Salvador Mascarenhas for this suggestion. See Hare 1952; Wright 1972 for classic discussions of these and related points. A similar point can be found in Chrisman (2018, p. 416), who proposes to think of the practical commitments associated with normative language in terms of commitments to *reason practically* in certain ways, that is, to accept certain considerations as reasons *for action*. See also Finlay 2004, p.206 n.3.

goes through the chain of inference in (2.40) or (2.41) might fail to take an umbrella. So the fact that someone, for example, avows an evaluative belief and then fails to behave accordingly can hardly be taken as evidence that the connection between evaluative thought and action is cancellable in the same way that, say, some implicatures are—the defeasible nature of practical reasoning makes it difficult to draw that conclusion.¹⁸

A second potential objection to my characterization of evaluative adjectives as action-guiding has to do with the *scope* of action-guidance. While it may be accepted that moral and practical adjectives are action-guiding, one might resist the view that this is true of aesthetic adjectives as well. In what sense is a sentence like (2.42) connected to action?

(2.42) What a beautiful sunset!

(2.42) seems to be merely an expression of pleasure at a particular perceptual experience.

Nonetheless, note that aesthetic adjectives show the same connections to practical attitudes and inferences to those displayed by other evaluative adjectives:

(2.43) Matheus said that the sunset was stunning although he immediately took out his phone to look at Instagram.

- (2.44) a. It's raining.
b. You are going out.
c. Rain makes your hair ugly.
∴ Take an umbrella.

We might think of it this way: whereas moral and practical adjectives guide us towards the thing *to do*, aesthetic adjectives guide us towards the thing *to experience*. An aesthetic judgment, even if based on pure contemplation like (2.42), involves the idea that the object of evaluation is worth experiencing. It involves a commendation of whatever is experienced, that is, an invitation to share the experience (*mutatis mutandis* for negative aesthetic adjectives).¹⁹

We conclude that all evaluative adjectives are action-guiding, and that it is this property which sets them apart from other, lexically subjective gradable predicates. We turn now to the family of philosophical proposals which claim to have the upper hand *vis-à-vis* an account of action-guidance, namely non-factualism.

2.3.2 Non-factualism about evaluative language

Non-factualism names a broad family of views originally about the meaning of moral vocabulary, but that has been extended to other parcels of natural language, notably perspectival expressions. Non-factualism can be stated as a thesis about the *semantic value* of declarative sentences that belong to a certain fragment of natural language, namely the thesis that those sentences do not describe reality. Thus, non-factualism is first and foremost a negative thesis:

Definition 1 (*Negative non-factualism about F*) *Unembedded declarative sentences of a fragment F of natural language do not describe the way the world is.*

¹⁸For accounts of action-guidance as implicature, see S. J. Barker 2000; Copp 2001, 2009; Finlay 2004, 2005; Strandberg 2011; and see Bar-On and Chrisman 2009 for criticism.

¹⁹ I thank Carla Umbach and James Brown for useful discussion and suggestions in relation to these issues.

There are non-factualists about various domains, such as normative claims (Gibbard 1990, 2003), epistemic (inc. probability) modals (Yalcin 2007, 2011, 2012), the *a priori* (Field 2000), higher-order concepts in general (Frápolli and Villanueva 2012), conditionals (Gibbard 1981), knowledge attributions (Chrisman 2007; Field 2009, 2018), or taste predicates (Bordonaba Plou 2017; Gutzmann 2016).

Non-factualists say that the function of the characteristic type of vocabulary in each of those fragments is ‘non-representational’ (Charlow 2014, 2015), that is, its function is not to describe ‘ways the world is’ (Frápolli and Villanueva 2012, p. 471), or ‘how things stand’ (Rosen 1998, p. 388). Sometimes these views are presented as committed to the idea that the relevant sentences are ‘not straightforwardly factual’ (Field 2009, 2018; Yalcin 2012).^{20,21} We can apply this view to evaluative sentences (sentences containing evaluative adjectives) more generally:

Definition 2 (Negative non-factualism about value) *Unembedded evaluative sentences do not describe the way the world is.*

To take an example, consider sentences (2.45) and (2.46):

(2.45) Eating animals is widespread.

(2.46) Eating animals is bad.

Non-factualists about evaluatives hold that the presence of the evaluative adjective *unethical* in (2.46) marks a profound semantic contrast between those two sentences. In most contexts, an utterance of (2.45) aims at informing an audience of a fact, namely that the practice of eating animals has a certain property (*being widespread*). In contrast to this, non-factualists claim, by uttering a sentence like (2.46), a speaker does not aim at informing her audience of a similar fact, namely that eating animals has the property of *being bad*. Rather, an utterance of (2.46) serves a different communicative function.

Different authors differ widely in their *positive* characterization of non-factualism, so it is difficult to give a positive thesis that everyone would agree with. Most of them, however, emphasize that the meaning of evaluative terms is somehow connected to action and practice. In order to keep our characterization maximally open, we are going to define non-factualism about evaluative sentences as the view that evaluative sentences express practical attitudes of support and rejection.

Definition 3 (Non-factualism about value) *Unembedded evaluative sentences express practical attitudes of support & rejection.*

²⁰ See the opening paragraphs of e.g. Bedke 2017; Blackburn 2006; Boghossian 1990; Camp 2017a; Frápolli and Villanueva 2012; Gibbard 1986; Jackson and Pettit 1998; Pérez Carballo and Santorio 2016; Rothschild 2012; Schroeder 2008a,d for different statements of this negative thesis.

²¹ Treating non-factualism as a semantic view, that is, as a view about the semantic meaning of certain sentences, is one option among others. Non-factualism has been defended as a *metasemantic* view (a view about ‘what it is in virtue of which particular kinds of words have the semantic contents that they do’, Chrisman 2012, p. 325, also Pérez Carballo 2014; Ridge 2014) and as a *pragmatic* view (Yalcin 2012). It is not extremely clear to me what Yalcin means by calling non-factualism/expressivism a *pragmatic* thesis, although in Yalcin 2012, pp. 140–2 the basic thought seems to be the following: by refusing to call non-factualism a semantic thesis, one can remain agnostic about two things: (i) what the correct compositional implementation of these ideas might turn out to be, and (ii) whether the characteristic type of update on the conversational common ground that Yalcin attributes to modals (which are his focus) is identifiable with the semantic value of those lexical items.

Non-factualists tend to say that evaluative vocabulary serves to communicate things like desires, plans, states of (dis-)approval, PRO- and CON-attitudes, intentions or emotional or affective states. For example, a sentence like (2.46) may express a plan to avoid eating animals, or a desire to not eat animals, or a state of disapproval towards eating animals, or a negative or displeasing emotion or affective state produced by the thought of eating animals, etc. These are all action-oriented attitudes against eating animals.²²

In fact, we can be more specific and say that attitudes of support are expressed by evaluative sentences that contain evaluative adjectives of *positive* polarity (*good, generous, beautiful...*) while attitudes of rejection are expressed by evaluative sentences containing evaluative adjectives of negative polarity (*bad, cruel, ugly...*).

Definition 4 (Non-factualism about positive/negative value) *Unembedded evaluative sentences containing positive evaluative adjectives express practical attitudes of support; evaluative sentences containing negative evaluative adjectives express practical attitudes of rejection.*

Furthermore, it is standard in the literature on non-factualism and expressivism to construe those attitudes as mental states (Schroeder 2008c, *a.o.*), but we can remain neutral about that—you can think of having an attitude as being in a particular type of mental state, or you can think of it as adopting some kind of public commitment to think and/or behave in certain ways. To have a practical attitude of support of football, for instance, can be thought of as being disposed or committed to display certain types of behavior with respect to football. It is a further question whether this attitude is cashed out in terms of being in certain types of mental states, such as desiring to watch a game tonight, or planning to play football with one's friends next week, or crying joyful tears when one's team wins, etc. We do not wish to take a stand on what is the right way to construe the mental state counterpart of such practical attitudes.

It is an important part of these views to rely on a relation of EXPRESSION that holds between a speaker, a sentence type and an attitude type that can be attributed to the speaker of that sentence, as a result of using it in paradigmatic occasions of use (Schroeder 2008b). Let » denote that relation (although we omit mention to the speaker). Then, we can represent the different attitudes expressed by sentences (2.45) and (2.46) as follows:

- (2.45) Eating animals is widespread.
» *non-practical/representational attitude about eating animals*
- (2.46) Eating animals is bad.
» *practical/non-representational attitude against eating animals*

As we stressed at the start of this section, often it has been more important for non-factualists to stress the contrast between representational or doxastic attitudes associated with descriptive vocabulary and the practical or non-doxastic attitudes associated with evaluative vocabulary than to specify what those practical attitudes really are. According to this statement of non-factualism about evaluatives, the contrast between non-evaluative and evaluative sentences lies in the fact that the former express attitudes that represent the world as being a certain way; while the latter express attitudes that aim at making the world a certain way, rather than representing it as such. Waxing Marxist, one might say that representational, or doxastic attitudes (like belief and knowledge) aim at *interpreting* the world; while the practical attitudes associated

²² See e.g. Ayer 1946, p. 110; Blackburn 1984, 190 and ff; Blackburn 1993, p. 113; Blackburn 1998, p. 70; Camp 2017a, p. 87; Finlay 2004, p. 206; Gibbard 1986, p. 473; Gibbard 1990, p. 33; Schroeder 2008a, p. 574; Jackson and Pettit 1998, p. 245, as well as references in n.20.

with evaluative vocabulary aim at *changing* it. Again, we need not be more precise than this about what exactly is the attitude in question, or how it may map onto different psychological states.

Another illuminating aspect of non-factualism arises in contrast to so-called metaethical SUBJECTIVISM (see n.2). Note that, if one accepts that evaluative terms are semantically related to practical attitudes, the non-factualist view that evaluative terms *express* those attitudes is not the only possible way of cashing out that relation. An alternative view is subjectivism, according to which evaluative terms express regular beliefs *about* those attitudes. According to this view, a sentence like (2.46) self-ascribes certain practical attitude towards eating animals; it is equivalent to saying something like *I have certain practical attitude towards eating animals*.²³ Early proponents of non-factualism insisted on distinguishing their view from subjectivism²⁴, essentially because subjectivism was prey to Moore's open question argument as well as to the obvious charge that speakers engaging in moral disagreement could hardly be represented as describing their respective moral responses (the problem of lost disagreement). By stressing that, in making a moral claim, one does not *say* that one has a practical attitude but rather *expresses* the attitude itself appeared to overcome the difficulties of subjectivism raised by Moore.

Non-factualism and action-guidance

This simple characterization of non-factualism should already make it clear that this approach is well equipped to account for the action-guiding feature of evaluatives (§2.3.1): simply put, non-factualists hardwire the action-guiding inference that we observed into the semantic meaning of evaluatives through the aforementioned relation of *expression*: they claim that that the content of that inference *just is* what evaluative sentences semantically express.

According to non-factualists about evaluatives, to call something *good* is to express a—perhaps underspecified—positive practical attitude towards that thing; it is to express one's intention of orienting one's actions towards that thing. Thus, for non-factualists, the connection between evaluative language and action-guidance is no mystery; to make an evaluative claim *just is* to express how one intends to direct one's action.

With respect to the *although*-test, a non-factualist would say that the reason why (2.37a) and (2.38a) sound natural is because in those sentences we are ascribing to someone a relatively irrational, or conflictive state of mind: when we describe Matheus as someone who thinks that volunteering is a virtuous thing to do and at the same time lacks any intention of doing it, it is as though we are saying that he takes himself to have certain practical attitudes when he actually does not have them (2.37a), or displays intentions or dispositions to act that are markedly incoherent with the former (2.38a). Marking that contrast with a connective like *although* sounds like a perfectly natural thing to say.

(2.37a) Matheus thinks volunteering is virtuous, although he doesn't have any intention or plan whatsoever of supporting, promoting or doing it.

²³A forceful defense of this view can be found in Dreier 1990.

²⁴ '[W]e reject utilitarianism and subjectivism, not as proposals to replace our existing ethical notions by new ones, but as analyses of our existing ethical notions. Our contention is simply that, in our language, sentences which contain normative ethical symbols are not equivalent to sentences which express psychological propositions, or indeed empirical propositions of any kind' (Ayer 1946, pp. 107–8)

≈ *Matheus thinks he has practical attitude P although he does not have P.*

(2.38a) Matheus thinks donating is unethical although he belongs to a charity.

≈ *Matheus thinks he has practical attitude P although he behaves incoherently with P.*

Similarly, for a non-factualist it is no surprise either that evaluative sentences and practical propositional attitude ascriptions can be substituted *salva validitate* in practical reasoning, given that the former semantically express the same kind of attitudes that the latter, in a standard semantic sense, describe or represent. So a non-factualist would say that the reason why the third premise in (2.40) and (2.41) (repeated here) can either be an evaluative sentence or the attribution of a desire is that evaluative sentence express practical, or action-oriented, attitudes (such as desires).

(2.40) a. It's raining.

b. You are going out.

c. You don't want to get wet.

∴ Take an umbrella.

(2.41) a. It's raining.

b. You are going out.

c. It's a bad idea to get wet.

∴ Take an umbrella.

Non-factualism and gradability

Things start to get more difficult when we consider the question of whether non-factualists can secure the fact that evaluative adjectives are lexically subjective, as we saw in §2.2. At the outset, non-factualism looks promising: given that practical attitudes can vary from speaker to speaker, different speakers can use evaluative adjectives to express different practical attitudes in different contexts of use. And it is sensible to think that those different attitudes give rise both to variable thresholds for the application of the positive form of gradable adjectives (allowing for positive form subjectivity), as well as to variable orderings of individuals along the relevant evaluative dimensions (allowing for ordering- or lexical subjectivity). This subjective element in non-factualism could, in principle, account for the intuition that speakers in disagreement dialogues like (2.22) (repeated here) can have different opinions.

Nonetheless, things are not so straightforward. As we have seen, evaluative sentences containing positive evaluative adjectives like (2.47) express support; while evaluative sentences containing negative evaluative adjectives like (2.48) express rejection. This approach, however, faces a fundamental shortcoming when faced with comparative constructions like (2.49):

(2.47) Volunteering is good.

» *practical attitude for volunteering*

(2.48) Donating is bad.

» *practical attitude against donating*

(2.49) Volunteering is better than donating.

» *practical attitude for / against ??*

When a speaker utters (2.49), she need not endorse nor reject either *relata*. She is merely comparing the two actions; and her uttering (2.49) is compatible with adopting almost any combination of positive and negative attitudes towards either of them (with the exception of being for donating money while being against volunteering). This is a clear counterexample to

Definition 4:

Definition 4 (Non-factualism about positive/negative value) *Unembedded evaluative sentences containing positive evaluative adjectives express practical attitudes of support; evaluative sentences containing negative evaluative adjectives express practical attitudes of rejection.*

The reason why comparatives cannot be easily accommodated by non-factualism, at least the way we have formulated it, is because we have defined the practical attitudes associated with evaluative vocabulary as *non-gradable*, that is, as outright positive or negative practical attitudes. But evaluative adjectives are gradable, and therefore, if non-factualists are right, the practical attitudes that evaluatives express should be gradable as well. Non-factualists need to enrich their favored semantics if they want to accommodate comparative uses of evaluative adjectives.

More specifically, what one is inclined to say about a sentence like (2.49) is not that it expresses outright support/rejection for either *relata*, but that it expresses *support, or preference, or priority of volunteering over donating*. But simply assigning the expression of such a relational attitude as the meaning of *better*, while assigning the simpler, outright attitude to *good* is not enough—for the simple reason that *good* and *better* are semantically related in an obvious way. So non-factualists about evaluatives need to say something about how these more sophisticated relational attitudes that one would be inclined to assign to comparatives are semantically built from the simpler attitudes that non-factualists assign to the positive form of the same adjectives.

Importantly, comparatives are just a case study, but the issue goes well beyond: basically, the type of practical attitudes that non-factualists associate with evaluative adjectives are most appropriate for capturing the meaning of those adjectives *in their positive form*. But as we saw in §2.2, gradable adjectives are extremely flexible. They appear in comparative, equative or superlative form, they admit certain adjectival modifiers and not others, etc. Therefore, non-factualism needs to be refined in order to be able to account for the meaning of the great variety of forms that evaluative adjectives appear in.

Before moving on, it bears pointing out that some proposals in the non-factualist camp are better equipped to meet this challenge than others. Comparatives (and relational constructions involving evaluatives more generally) are a particularly pressing problem for versions of non-factualism that hold that evaluative sentences express attitudes *towards* a single object of evaluation. Comparatives are problematic because it does not make sense to say that the speaker is expressing the relevant practical attitude towards any of the two *relata*. This is true of most non-factualist proposals, but not all of them. In particular, Gibbard's early NORM-EXPRESSIVISM (Gibbard 1986, 1990) was built upon the idea that ascriptions of rationality (*ϕ-ing makes sense*; *ϕ-ing is rational*, etc.) express a non-cognitive attitude of ACCEPTANCE of a set of norms that stands behind such evaluation. For instance, to say that eating vegetables is the rational thing to do is to express one's acceptance of one's system of norms. This is a system of norms that so happens to sanction eating vegetables. But one does not thereby express any positive or negative attitude *towards* eating vegetables; one expresses an attitude that has one's system of norms as its object.

This feature of Gibbard's early view makes it better equipped to deal with (2.49). For Gibbard can also say that to judge that volunteering is better than donating is to express one's acceptance of a system of norms, a system of norms that gives priority to volunteering over donating. Nonetheless, that would not be enough: he would still need to say what it is for something to give priority to something over something else. Gibbard's examples—outright judgments of rationality, sensibleness and the like, were treated in his earlier work as expressing acceptance

of systems of norms such that the objects of evaluation either met, failed to meet or were indifferent to them. The problem remains that neither of those cases applies in a straightforward manner to comparatives.²⁵

To sum up these two last sections, we have arrived at a sort of unexpected crossroads. Walking in the linguist's shoes (§2.2), we arrived at a necessary property of evaluative adjectives, namely ordering subjectivity (§2.2.1), and we saw how a standard degree semantics can accommodate those properties (§2.2.2). But we also argued that this approach has certain shortcomings: it requires thinking of evaluative adjectives in a way that is similar to the way in which we think about dimensional adjectives, which is awkward in light of the fact that evaluative adjectives eschew measurement. More importantly however, degree semantics offers no clear way of making a distinction, among the lexically subjective adjectives, between those that are action-guiding and those that are not. Then, putting on the metanormative theorist's shoes (§2.3), we argued that a sufficient property of evaluatives is precisely their action-guidance (§2.3.1), and we saw how non-factualism, by assigning a specific semantics that connects evaluatives to practical attitudes, can in principle account for the action-guiding properties of evaluatives (§2.3.2). But non-factualism has its own shortcoming, which is that it offers an account best suited to capture the use of evaluative adjectives in their positive form, and it is not obvious how to incorporate gradability into the picture.

The outcome of all this is that the scalar properties of evaluatives need to inform non-factualism, and conversely that the action-guiding properties of evaluatives need to inform the standard degree semantics for evaluatives. Unfortunately however, this is not simply a matter of theories complementing each other: in the following two sections (§2.4-§2.5), I aim to argue that capturing the action-guiding properties of evaluatives in light of their scalar properties causes a problem for degree semantics, and that this should push us to move to a different framework for scalar semantics, namely delineation semantics, as the proper semantic framework in which to couch non-factualism. But in order to see this problem clearly, it is best to first present the most important problem that besieges non-factualism, namely the Frege-Geach problem.

2.4 The (supra-sentential) Frege-Geach problem

If non-factualism about a subset of declarative sentences—let us call it fragment F —is the thesis that elements of F do not describe or represent facts, then it follows that they do not have propositions that can be true or false as their semantic values (at least under standard conceptions of propositions as representations of facts or states of affairs). But *truth-functional* operators such as modals or logical connectives take propositions as their arguments. If non-factualism about F is right, then one should expect that the sentences of F fail to embed under such operators. Now, *all declarative sentences*—including sentences of F —embed under such operators. Therefore, non-factualism about F has a problem.

²⁵ Silk 2015's *ordering expressivism*, and Ridge's view of *good* as meaning 'highly ranked according to any acceptable standard' (2014, p. 26, my emphasis) incorporate explicit scalar elements in their proposals, which makes them better suited for the challenge of building a scalar semantics. Schroeder treats the comparative *better than* as atomic (Schroeder 2008c, p. 7), and says that *a is better than b* expresses a non-doxastic attitude of "being for preferring *a* to *b*" (Schroeder 2008c, pp. 58, 77). All these authors are on the right track, although none of them (including Gibbard) tackle the compositional issue of relating the meaning of the comparative and the positive form of evaluative adjectives.

More generally, the syntactic and semantic behavior of declarative sentences (including evaluative and non-evaluative sentences) is broadly uniform, and therefore the distinction that non-factualists insist on drawing between evaluative and non-evaluative language is difficult to justify on purely linguistic grounds. Schroeder puts the matter bluntly, discussing non-factualism about moral terms: ‘[t]here is no linguistic evidence whatsoever that the meaning of moral terms works differently than that of ordinary descriptive terms. On the contrary, everything that you can do syntactically with a descriptive predicate like ‘green’, you can do with a moral predicate like ‘wrong’, and when you do those things, they have *the same semantic effects*’ (Schroeder 2008c, p. 704, his emphasis).²⁶

It is worth it to go briefly over Geach 1965’s original complaint about early non-factualist proposals. In particular, Geach targeted Hare’s *Language of Morals* (1952).²⁷ Hare aimed to give a semantics for evaluative expressions based on considerations about the kinds of speech acts that competent speakers make with those expressions. Hare claimed that to call something *bad* was to make a speech act whereby one condemns it; conversely, to call something *good* is to make a speech act whereby one commends that thing. The meaning of the pair *good/bad* just turn out to be those very acts of commendation and condemnation which are, in turn, kinds of practical attitudes.

- | | |
|--------------------------------|--|
| (2.46) Eating animals is bad. | » <i>practical attitude against eating animals</i> |
| (2.50) Eating animals is good. | » <i>practical attitude for eating animals</i> |

But Geach was quick to observe that there are many uses of *good/bad* where those words are not used to commend/condemn. For example, when it is under negation, inside a question or in the antecedent of a conditional (let us focus only on *bad*):

- | | |
|--|--|
| (2.51) Eating animals is not bad. | » <i>practical attitude for / against ??</i> |
| (2.52) If eating animals is bad, then I should stop. | » <i>practical attitude for / against ??</i> |
| (2.53) Is eating animals bad? | » <i>practical attitude for / against ??</i> |

²⁶ Even more emphatically:

‘[E]very natural-language construction that admits of descriptive predicates admits of moral predicates, and seems to function in precisely the same way: tense; conditionals; every kind of modal—alethic, epistemic, or deontic; qualifiers like ‘yesterday’; generics and habituals; complement-taking verbs like ‘proved that’ and ‘wonders whether’; infinitive-taking verbs of every class, including ‘expects to’, ‘wants to’, and ‘compels to’; binary quantifiers like ‘many’ and ‘most’; and more. It is crucially important to understand that the embedding problem for non-cognitivism is not simply a problem about the validity of *modus ponens*, or even simply about logic. Every construction in natural languages seems to work equally well no matter whether normative or descriptive language is involved, and to yield complex sentences with the same semantic properties. (Schroeder 2008a, p. 5).

This is not completely true, however. To name a few points of contrast between evaluative and non-evaluative expressions: while almost all evaluative adjectives embed under *find*, many non-evaluative adjectives do not (§2.2.1). As we will see in §4.6, evaluative adjectives do not admit precise ratio modifiers (# 5.4x *crueler* vs. 5.4x *taller*), although they admit round ratio modifiers (*twice as cruel*). Evaluative adjectives also do not admit measure phrases (# 2 “units” *good* vs. 2m *tall*). Moreover, compare Frápolli and Villanueva 2012, p. 475, who point out that ‘there are syntactic moves that can be performed with a second-order predicable [*good*] that cannot be performed using a first-order predicable [*green*]’. An obvious example is the fact that *good* can take sentences as well as nominal phrases as syntactic argument (more on this in §4.3).

²⁷ I partially follow here Schroeder’s (Schroeder 2008c,d) presentation of the problem.

None of these sentences express condemnation. So if non-factualists like Hare wanted to maintain that the meaning of *bad* just is its use to express condemnation, they one would need to explain why the very same word figures in sentences that clearly fail to express that.

An initial reaction to this problem was to assume that evaluative sentences had different meaning when embedded—(2.51)-(2.53)—and when unembedded—(2.46),(2.50). Perhaps the condemnation effect arises when evaluative sentences are asserted but it is otherwise absent. This route, however, faced an even greater hurdle, which is to account for the obvious inferential connections between embedded and unembedded instances of evaluative sentences. Consider the following inference:

- (2.54) a. If eating animals is bad, then eating beef is bad.
 b. Eating animals is bad.
 ∴ Eating beef is bad.

It seems obvious that this chain of inference relies crucially on the antecedent of the first premise and the second premise having the same meaning. If non-factualists denied that, then they would have to find an alternative explanation for why that inference holds.

To be sure, the culprit of this problem is not the conditional operator ‘*if...then...*’ nor *modus ponens*, because the problem generalizes to any inference involving embedded and unembedded uses of evaluative sentences. For an illustration, the following inference relies on disjunction and disjunctive syllogism:

- (2.55) a. Either eating animals is not bad or eating beef is bad.
 b. Eating animals is bad.
 ∴ Eating beef is bad.

As opponents of non-factualism emphasized, there is no satisfactory way of accounting for the cogency of those otherwise unremarkable inferences, unless one assumes that the embedded and unembedded instances of the relevant evaluative sentences have the same meaning. The inference in (2.54) holds because the meaning of the conditional is a function of the meaning of its antecedent and the meaning of its consequent and the way in which they are composed (*mutatis mutandis* for (2.55)).

Nonetheless, Hare (1970) observed that accounting for the compositionality of language is not a special problem for non-factualists, but for *anyone* doing semantics. The special problem that non-factualists face is due to the fact that the semantic values that they were eager to assign to evaluative sentences are, in principle, of a different kind from those that are attributed to declarative sentences in general. To stick with Hare’s view, let us assume that a sentence like (2.46) expresses a practical attitude of condemnation. But how can a practical attitude figure as the argument of a disjunction, or a conditional? Practical attitudes would need to be able to combine with the meaning of sentential operators in order to result in the expected meanings for the complex sentences.

The Frege-Geach problem, at heart, is the question of how to apply the principle of *compositionality* to evaluative sentences in light of the revisionary semantics that non-factualists assign to them: *how is the meaning of complex sentences containing embedded evaluative sentences built from the meaning of its constituents parts* (namely, unembedded evaluative sentences together with various sentential operators) *and the way in which they are combined*? This is the

traditional Frege-Geach problem, or as I will call it here, the “SUPRA-SENTENTIAL” FREGE-GEACH PROBLEM.

Among the non-factualists taking up this challenge, we can distinguish those who have tried to maintain a pure non-factualist account, and those who have opted for some kind of HYBRID VIEW.²⁸ The difference between these two approaches can be summarized as follows: given non-factualism (Definition 3/4) and given (2.54), either one has to be more liberal about what kind of semantic object can go into the antecedent of a conditional (or any sentential operator), or one needs to say that, in addition to whatever unconventional semantic object is assigned to unembedded evaluative sentences, these sentences also have traditional semantic values that can be plugged into sentential operators in a standard way. Let us consider briefly both options.

Hybrid non-factualists choose the latter option and assume that one needs to make room for traditional semantic values—classic propositions²⁹—if one is to avoid the Frege-Geach problem (this move is advocated as early as Gibbard 1986). Hybrid views thus factor out the meaning of evaluative sentences into a traditional descriptive belief representing a state of affairs that can be evaluated for truth and a non-representational, practical component. The non-representational component is meant to account for the connection of evaluative sentences to action and motivation. The descriptive component is what embeds under truth-functional operators and is therefore productive in explaining the meaning of complex constructions and the inferences that they enter into. Hybrid theorists would thus claim that a sentence like (2.46) expresses both a regular, representational attitude as well as a practical attitude. The representational component is most often constructed as a regular belief that the speaker has the relevant practical attitude; while the non-representational component is the practical attitude itself.

(2.46) Eating animals is bad.

» *representational/non-practical attitude* that the speaker has a practical attitude towards eating animals

» *practical/non-representational attitude* towards eating animals

In this type of view, when one sees a sentence like (2.46) embedded in a conditional (cf. (2.52)), it is only the representational attitude component which survives the embedding, and that is sufficient to account for why an inference like (2.54) goes through.

Further details mark important differences between hybrid views. Note, for instance, that we have not specified the content of neither the representational nor the practical attitudes involved. For example, it marks an important difference between hybrid proposals whether the practical attitude is different for different speakers uttering one and the same evaluative sentence (that is, whether anyone who utters (2.46) is expressing the same practical attitude or whether that practical attitude can be different for different speakers). And it also matters whether the same attitude is expressed by different evaluative sentences in the mouth of one and the same speaker

²⁸ The former group includes Schroeder himself (2008a,b,c,d), but arguably also Gibbard 2003 and more recent authors writing in the wake of Schroeder and Gibbard’s later work, such as Yalcin (2012, 2018, 2019), Charlow (2014, 2015) Silk (2015) or Willer (2017). Hybrid theorists include Bar-On & Chrisman (2009), early Gibbard (1986, 1990), Ridge (2006, 2014), Boisvert (2008) or Copp (2001). See also Schroeder 2009 and the papers in Fletcher and Ridge 2014.

²⁹ Throughout this work, by *proposition* I will mean the classical notion of proposition from intensional semantics, that is, a set of possible worlds. As we will see briefly in a few paragraphs, and at large in the following chapters, the semantic values that we propose to assign to declarative sentences are *not* classical propositions, but more complicated semantic objects. We will not call them propositions; we will speak more generally of the semantic values or the denotation of declarative sentences.

(e.g., whether my utterance of (2.46) and my utterance of *it is generous to donate money to charity* express a single, “content-neutral” practical attitude) (see Schroeder 2009, for discussion).

On the other hand, Gibbard’s PLAN-EXPRESSIVISM (2003; see also Charlow 2014, 2015; Yalcin 2012, 2019), is—arguably—a view of the former kind. Gibbard defends that, just as the semantic value of descriptive sentences is identified with sets of possible worlds, that is, (classical) propositions, the semantic value of normative sentences can be identified with sets of HYPERPLANS (written with variable h), which are maximally specific plans of action:

(2.45) ‘Eating animals is widespread’ = $\{w : \text{eating animals is widespread in } w\}$

(2.46) ‘Eating animals is bad’ = $\{h : \text{eating animals is forbidden by } h\}$

Without going into further detail (see Chapter 3, esp. §3.3), the fact that hyperplans are set-theoretical objects just like possible worlds provides Gibbard with the same modelling tools that are conventionally used to model logical connectives in standard possible world semantics, namely the usual Boolean operations over sets: negation is set complementation, conjunction is set intersection, etc. All that we require in order to account for the Frege-Geach problem is a liberal (but not particularly controversial) view of logical connectives according to which they can take arguments of other semantic types besides classical propositions.³⁰

The answer to the question of how can a practical attitude figure as an argument to a disjunction or a conditional is that, if we can model such attitudes in terms of some appropriate model-theoretical object, this is not a problem. And once we make a conditional (or disjunction, conjunction, etc.) flexible enough to have a set of hyperplans in its antecedent, at least this version of the Frege-Geach problem vanishes. Because even though (2.46) expresses a practical attitude and (2.52) does not, we have dropped the assumption that what is fed to the *if*-clause of (2.52) has to be a classic proposition (a set of worlds). So (2.46) and the antecedent of (2.52) can have the same semantic value. It is a special kind of semantic value (a set of hyperplans, in Gibbard’s case), such that, when it is asserted, it expresses that attitude, but when it is inside an *if*-clause, it is considered hypothetically, so to speak—just as a classical proposition. But crucially, one and the same semantic object can figure within and without the *if*-clause, so that inferences like (2.54) are no longer unexplained.³¹

Critics like Schroeder (2008c, see also Charlow 2014; Willer 2017) have pointed out that Gibbard’s view lacks a principled account of the logical relations among the model-theoretical objects that Gibbard brings into the picture; but other authors have rebutted that possible-world semantics is in no better position (Pérez Carballo 2011). Having said that, let us stress that we do not aim to offer any novel solution to the Frege-Geach problem in this chapter. But the Frege-Geach problem helps to illuminate the type of problem that we want to tackle.

To recap: the Frege-Geach problem is the problem of explaining how the alternative semantic values that non-factualists assign to evaluative sentences can combine with logical connectives

³⁰ Such a view is independently needed to account for non-propositional uses of certain connectives, such as *and* and *or* (*She was wearing a new and expensive dress*, Partee & Rooth 1983/2002), as well as for uses of conditionals with non-propositional consequents, such as *if I fall asleep, wake me up!*. See Charlow 2015, 5 and ff for discussion.

³¹ We are simplifying things a little bit for the sake of presentation. Strictly speaking, normative and descriptive sentences *have* to have the same type of semantic value, since sentences of both kinds can figure together as arguments of logical connectives. As we will see at length in §3.3, this is achieved by letting declarative sentences have sets of world and hyperplan *pairs* as their semantic value.

and other truth-functional operators in the appropriate way to give rise to complex expressions. Faced with this difficulty, hybrid non-factualists opt for assuming that evaluative sentences must have traditional propositions as part of their meaning; while pure non-factualists like Gibbard (and others in his wake) have turned to enriching our view of semantic values and logical connectives.

2.5 The sub-sentential Frege-Geach Problem

In §2.3, we raised a challenge for non-factualism: In a nutshell, non-factualists claim that evaluative sentences express practical attitudes of support or rejection, but evaluative sentences with evaluative adjectives in comparative form do not express such attitudes. We argued that this was not a problem about comparatives *per se*—the problem was that the standard non-factualist story was best suited for the positive form of evaluative adjectives. All of this should sound very similar to the Frege-Geach problem, and the aim of this section is in part to argue that the challenge raised by comparatives reveals a new, “sub-sentential” incarnation of the Frege-Geach problem.

2.5.1 Non-factualism and the positive form: POS-ATT

There is a clear correlation between the capacity of an evaluative sentence to express outright support and rejection, on the one hand, and the presence of the positive form of the relevant evaluative adjective, on the other. Put differently, it turns out that there is a connection between the *metanormative* and the *linguistic*—or more precisely, Rett’s—notion of evaluativity (which probably seemed unrelated at the outset).

The connection is the following: in order for an evaluative sentence to express practical support/rejection, it has to be evaluative (in Rett’s sense); conversely, every linguistically evaluative construction that features evaluative adjectives expresses support/rejection.

Observation 1 (POS-ATT) *An utterance U of an unembedded evaluative sentence S containing an evaluative adjective E expresses practical support/rejection if and only if U invites an inference to the positive form of E .*

It follows that any evaluative sentence that does not invite an inference to the positive form of the relevant adjective—any evaluative sentence that is not *linguistically* evaluative—poses a counterexample to our definition of non-factualism:

Definition 4 (Non-factualism about positive/negative value) *Unembedded evaluative sentences containing positive evaluative adjectives express practical attitudes of support; evaluative sentences containing negative evaluative adjectives express practical attitudes of rejection.*

Let us illustrate this with comparatives. The comparative *better* is not linguistically evaluative, as shown by the fact that (2.56) is a perfectly acceptable thing to say:

(2.56) Volunteering is better than donating, and none of them are good.

POS-ATT rightly predicts that *better* should not express practical support or rejection of any *relatum* (which we observed at the end of §2.3.2):

(2.49) Volunteering is better than donating. » *practical attitude for / against ??*

This applies to many other constructions as well. For instance, as we saw in §2.2.3, equatives built using the marked or negative member of a pair of gradable antonyms are linguistically evaluative, while the same construction with the unmarked member is not (Rett 2007):

(2.57) Ann is as tall as Mary. ↗ Ann or Mary are tall

(2.58) Ann is as short as Mary. ↘ Ann and Mary are short

The same contrast appears with pairs of evaluative antonyms:

(2.59) The first movie was as good as the second. ↗ Both movies were good

(2.60) The first movie was as bad as the second. ↘ Both movies were bad

While uttering (2.59) does not commit one to either movie being good, uttering (2.60) suggests that both movies were bad. And in accordance with **POS-ATT**, while (2.59) does not express any positive or negative attitude towards either movie, (2.60) clearly does.

(2.59) The first movie was as good as the second. » *practical attitude for / against ??*

(2.60) The first movie was as bad as the second. » *practical attitude against both movies*

Therefore, equatives built with the positive or unmarked member of a pair of evaluative antonyms also cause trouble for Definition 4.³²

In §2.4, we presented the “supra-sentential” Frege-Geach problem as the problem of *how evaluative sentences semantically embed*. The challenge was to assign a non-factual semantic value to evaluative sentences that can compose semantically with truth-functional operators in a satisfactory manner. Turning our attention to **POS-ATT** reveals an arguably more pressing question, namely the question of *how evaluative sentences semantically compose*. The challenge here is to assign a non-factual meaning to evaluative adjectives in a way that can then compose semantically with other constituents in the sentences where those adjectives figure and yield the appropriate meaning of those sentences. This is as a different “incarnation” of the Frege-Geach problem, which we dub the “SUB-SENTENTIAL” Frege-Geach problem.

This problem is already diagnosed in the previous paragraphs in pretty much the same way that the traditional Frege-Geach problem was: non-factualists claim that evaluative sentences express practical attitudes. But by (**POS-ATT**), this is true only of some evaluative sentences,

³² Two classes of potential counterexamples to (**POS-ATT**) should be mentioned. First, tautological sentences such as *every good action is good* entail the positive form *good* although they don’t seem to convey practical support for anything. There are various ways of tweaking (**POS-ATT**) to avoid this problem. One could restrict (**POS-ATT**) to non-tautological sentences. Alternatively, one could bite the bullet and say that sentences like this do express practical support for *every good action*. Secondly, sentences such as *Your paper is good enough*, while they do not invite an inference to the positive form, do seem to convey an outright practical attitude, in particular, a negative one. It’s plausible to think, however, that this is the result of a pragmatic process: calling something *good enough* suggests that its value does not reach the standard for *good* (*simpliciter*), that is, that it’s *not good*. It’s natural to think of this as some kind of scalar inference. Still, inferring *not good* is not enough to generate a negative attitude—for this we would need an inference to the positive form *bad*. Regardless, it is a rather stable convention of our linguistic practices that calling something *not good* is a polite way of calling it *bad*. If that convention is assumed, then the scalar inference generated by *Your paper is good enough* could actually result in the inference that your paper is bad, which would convey a negative attitude, as expected by (**POS-ATT**).

namely those that are *linguistically* evaluative. However, regardless of whether a given evaluative sentence is linguistically evaluative or not, its meaning has to be derived compositionally from the lexical meaning of the evaluative adjective that it contains. For example, the meaning of the binary predicate *as good as* is derived from the meaning of *good*, and the meaning of *as bad as* is derived from the meaning of *bad*—even though only the latter predicate produces sentences that are linguistically evaluative and express rejection. So non-factualists need to say how and why, even though evaluative adjectives semantically contribute the expression of practical support/rejection, such attitudes somehow vanish when those adjectives appear in certain forms.

2.5.2 From degree to delineation semantics

Here is where the difficulties for degree semantics appear. As we saw, degree semantics holds that the meaning of the positive form of any gradable adjective results from the composition of two semantic constituents: the lexical meaning of the adjective and a threshold-contributing silent morpheme *POS*. The problem is that neither of those components seems to be the right source of the practical attitudes of support/rejection that non-factualist assign to evaluative adjectives.

Let us look at this more closely. As we saw in §2.2.2, in degree semantics gradable adjective are assigned measure functions as their lexical meaning. Recall our entry for *tall*:

$$(2.24) \quad [[\text{tall}]]^c = \lambda x. \mu_{\text{height}}(x)$$

In order to build the predicate *is tall*, the measure function in (2.24) has to combine with the silent morpheme *POS*:³³

$$(2.61) \quad [[POS]]^c = \lambda g. \lambda x. g(x) \geq d_t^c$$

POS predicates a property of whatever degree is assigned to the individual in its second argument, namely of *being equal or greater to a specific degree on the relevant scale*. This degree is called a THRESHOLD or STANDARD, and it is determined contextually (d_t^c stands for the *contextually determined threshold degree*). The predicate *is tall* results from applying the lexical meaning of *tall* as the argument of *POS*:

$$(2.62) \quad \begin{aligned} [[\text{is tall}]]^c &= [[POS]]^c([[\text{tall}]]^c) = \\ &\lambda g. \lambda x. g(x) \geq d_t^c(\lambda x. \mu_{\text{height}}(x)) = \\ &\lambda x. \mu_{\text{height}}(x) \geq d_t^c \end{aligned}$$

When an argument is supplied to the resulting predicate, the result is a sentence with the following truth conditions:

$$(2.63) \quad [[\text{Ann is tall}]] = \mu_{\text{height}}(\text{Ann}) \geq d_t^c$$

Those truth conditions state that Ann's height surpasses the contextually determined standard of height.

³³ See n.10 in this chapter.

A very similar story should apply to evaluative adjectives, such as *good*: the positive form predicate *is good* would attribute to its argument the property of possessing a degree of goodness that surpasses a threshold (the difference with *tall* lies in the fact that *good* is also lexically subjective, so in this framework the very measure function μ_{goodness} should be context-sensitive in addition to the threshold contributed by *POS*):

$$(2.64) \quad [[\text{The movie is good}]]^c = \mu_{\text{goodness}}^c(\text{the movie}) \geq d_t^c$$

As we have seen, according to non-factualists, an utterance of this sentence expresses a practical attitude of support of the movie. Our question, then, is: where does the practical attitude semantically derive from? In other words, what constituent of this sentence contributes the practical attitude of support? Given that, in degree semantics, the predicate *is good* results from the combination of two semantic pieces—and assuming that the practical attitude does not magically appear when you put the pieces together—we have two options: the first is to locate the source of the practical attitudes in the adjective’s lexical meaning, and the other one to locate it in the silent morpheme *POS*.

Prima facie, each option has something going for it: the first option is attractive in virtue of the fact that only evaluative adjectives express practical attitudes. And the second option is attractive in light of **POS-ATT**, that is, in light of the observation that the expression of support/rejection correlates with the presence of the positive form. But neither is satisfactory, (partly) for the reason that makes the other option attractive. If the expression of a practical attitude is contributed to the sentence’s meaning by the adjective’s lexical meaning, then why does it go away in the comparative, for example? After all (according to standard degree semantics), measure functions figure in the comparative’s meaning just as they figure in the positive form. On the other hand, if the outright positive attitude expressed by an utterance of *The movie is good* is contributed by *POS*, then why does not an utterance of *The movie is long* express that positive attitude too, given that *POS* plays the same role in both sentences?

Let us try to refine each possibility: in favor of the lexical option, we could say the following: just as an adjective like *tall* refers to a scale of height, an adjective like *good* refers to a scale of goodness. And the higher one goes on that scale, the more support receives any object that sits up there. This does not mean that ascribing *any* degree of goodness to an object is tantamount to expressing an outright positive attitude towards it—it depends on how good the object is! The higher you go on the scale of goodness, the stronger the expression of support. Expressing support is not a matter of surpassing a certain threshold degree on a scale, that in turn licenses the positive form—it is a matter of having a sufficiently high degree of goodness that the positive form is licensed. The reason why the positive form expresses outright practical support is because the positive form signals that the object is high up in the goodness scale; the reason why the comparative does not express outright support is because the comparative tells us how two objects are related to each other on the goodness scale, but not where they sit on that scale.

This line of argument can capture, I think, part of the contrast between the positive and the comparative form: the comparative relates two things but does not tell us where they stand on the goodness scale; while the positive form tells us that an object sits relatively high up on that scale. Nonetheless, this story still leaves the correlation **POS-ATT** somewhat mysterious. To see this, suppose that the scale of goodness goes from 0 (total practical indifference) to 1 (absolute unconditional support). And suppose that the positive form is predicated of any object that gets more than 0.7. The idea would be that the positive form signals that an object sits on

the 0.7-1 portion of that scale, and this is why an outright attitude of support gets expressed. But why do only objects that sit on the 0.7-1 portion of the scale get actual practical support? In other words, why does the location of the threshold also mark a qualitative “jump” in the practical attitude that gets expressed, from neutral to endorsing? At best, it remains to be explained why the location of the threshold degree somehow determines the location of the support-expressing degrees. In absence of some further story relating *POS* to the expression of outright support, this story is not satisfactory.

On the other hand, a defender of the “*POS* approach” could argue that the reason why the predicate *is good* comes with the expression of practical support while *is tall* does not really is due to *POS*... because there are different *POS*-morphemes for different adjectives. For evaluative adjectives, the *POS*-morpheme marks, not just the degree on the goodness scale from which you can start using the positive form *good*, but also the degree from which you start expressing outright support. By contrast, the *POS* morpheme for a dimensional adjective like *tall* only does the first of those two jobs, namely telling you where on the height scale to start calling individuals *tall*. The *POS*-morpheme for evaluatives is special.

To this argument there can be three lines of rebuttal: the first is that one of the virtues of positing *POS* is precisely that it is meant to be the same semantic component regardless of the adjective with which it is combined. Hence, it is expected that it always produces the same semantic effect. So it remains a mystery why a certain *POS* morpheme is capable of carrying the expression of practical support or endorsement, and another *POS* morpheme is not. The second line of rebuttal is simply that this proposal is *ad hoc*. Positing another *POS*-morpheme that does what we need it to do but is otherwise unmotivated is not explanatory. Finally, positing a specific *POS*-morpheme for evaluatives seems more like “lexical camp” solution—the special function that this new *POS*-morpheme serves (contributing the expression of outright support) can only be due to its combination with evaluative adjectives, as there is no other lexical item with which it combines to produce such an effect.

In sum: non-factualism clashes with degree semantics. Evaluative adjectives in their positive form have a property, namely the capacity to express practical support/rejection, which, as we have argued, cannot be traced back to any of the ingredients that, according to degree semantics, conspire in their semantic composition—it cannot be attributed to the measure function, nor to the silent morpheme *POS*. The culprit, as I see it, is the fact that in degree semantics the positive form of a gradable adjective is decomposed into two semantic constituents.

Fortunately for non-factualism, degree semantics is not the only available view about gradable adjectives: in *DELINEATION SEMANTICS* (C. Barker 2002; Benthem 1982; Burnett 2017; Klein 1980, a.m.o.), the positive form of a gradable adjective is a simple predicate, and comparatives and other constructions are semantically derived from it. By adopting this approach, our problem disappears, for the simple reason that the positive form of an evaluative adjective is not factored out into further semantic components. Hence, we are free to assign the expression of practical support/rejection as the very meaning of the positive form of evaluative adjectives (which are simple predicates). The question, then, becomes how to build a scalar system starting from such simple predicates. This will be our task in the coming chapters.

2.6 Delineation semantics & non-factualism

But before taking on that task, we will finish this chapter by introducing the Delineation framework. We will illustrate how the approach works by giving a semantics for a simple, relative-standard gradable adjective, namely *tall*.

Delineation approaches to gradable adjectives start off from the idea that the positive form is semantically prior to the comparative. This is independently supported by the fact that the comparative is morphosyntactically derived from the positive form, and early proponents of this approach emphasized that this procedure was the only compositionally-respectable way of giving a semantics for comparatives (Benthem 1982; Klein 1980).

In delineation semantics, positive form relative-standard gradable adjectives like *tall* are simple (although vague) predicates determining an extension, an anti-extension and an extension gap, which are sets of individuals, defined relative to a contextually determined comparison class. That much is involved in our use of *tall*. Is Vera the Matrioska *tall*, *not tall* or is it unclear? Depends on who you compare her with: compared to the smaller matrioskas, she is tall; compared to bigger ones, she is not; and compared to the ones around her, it is unclear (see Figure 2.2). Comparison classes vary contextually, which makes these adjectives context-sensitive in the positive form (and in turn allows for the positive form to give rise to the subjectivity effects that we saw in §2.2.1).



Figure 2.2: Natasha, Vera and the other Matrioskas.

The comparative *taller* is defined *via* the positive form *tall* and quantification over comparison classes, roughly as follows: for any two individuals x and y , x is taller than y just in case one can find some comparison class relative to which x falls in the extension of *tall* while y does not (introducing quantification over comparison classes captures the comparative *taller*'s lack of context-sensitivity, and therefore of *subjectivity*—again, in line with what was observed in §2.2.1).

We turn now to a more formal presentation of the Delineation framework, as applied to *tall*.³⁴ Let us define an extensional model as a tuple $M = \langle D, [[\cdot]] \rangle$, where D is a non-empty set of individuals and $[[\cdot]]$ is an interpretation function from pairs of well-formed expressions of English and comparison classes to the extension of those expressions in D . A comparison class is a subset of D that works as a semantic index relative to which the extension of gradable adjectives is determined at any M . So the model defines a domain of individuals, and then the comparison class determines a subset of that domain relative to which the extension, anti-extension and extension gap of *tall* are determined. Comparison classes can vary from context

³⁴ I partially follow Burnett's clear presentation of the Delineation Framework, cf Burnett 2017, 57 and ff.

to context (but we will skip mention of context to avoid clutter). Finally, assume that proper names are constant functions from comparison classes to individuals in D (that is to say, the semantic value of *Vera* in model M relative to a comparison class F , $\llbracket \text{Vera} \rrbracket_F^M$, is an individual in D).

Relative to model M and comparison class F (which is some subset of the domain of individuals D),

$$(2.65) \quad \llbracket \text{Vera is tall} \rrbracket_F^M = \begin{cases} 1 & \text{if } \llbracket \text{Vera} \rrbracket_F^M \in \llbracket \text{tall} \rrbracket_F^M \\ 0 & \text{if } \llbracket \text{Vera} \rrbracket_F^M \notin \llbracket \text{tall} \rrbracket_F^M \end{cases}$$

(We skip mention of the model M in what follows).

Suppose, again, that the domain are all the Matrioskas, but that the comparison class is on one occasion the matrioskas from Vera leftwards (call that comparison class L) and on another the matrioskas Vera rightwards (call it R , see Figure 2.3). Then we get the following result:

$$(2.66) \quad \llbracket \text{Vera is tall} \rrbracket_L = 1$$

$$(2.67) \quad \llbracket \text{Vera is tall} \rrbracket_R = 0$$



Figure 2.3: Natasha and Vera in comparison to some of the other Matrioskas.

The comparative *taller* is defined by considering how *tall* applies to each *relatum* across possible comparison classes. For any comparison class F ,

$$(2.68) \quad \llbracket \text{Vera is taller than Natasha} \rrbracket_F = 1 \text{ iff there is some } X' \subseteq D : \\ \llbracket \text{Vera is tall} \rrbracket_{X'} = 1 \text{ and } \llbracket \text{Natasha is tall} \rrbracket_{X'} = 0$$

That is, Vera is taller than Natasha (relative to comparison class F , although F plays no role in the comparative) just in case there exists a set of individuals that form a comparison class relative to which Vera is tall but Natasha is not tall. In turn, this means that Vera is taller than Natasha just in case there is a comparison class relative to which Vera is in the extension of *tall* and Natasha is not.

$$(2.69) \quad \llbracket \text{Vera is taller than Natasha} \rrbracket_F = 1 \text{ iff there is some } X' \subseteq D : \\ \llbracket \text{Vera} \rrbracket_{X'} \in \llbracket \text{tall} \rrbracket_{X'} \text{ and } \llbracket \text{Natasha} \rrbracket_{X'} \notin \llbracket \text{tall} \rrbracket_{X'}$$

Relative to our Matrioska model, such a comparison class can be found (namely L), thereby making the comparative true in that model.

By defining the comparative with the aid of a simple predicate that represents the positive form, the Delineation approach guarantees that the positive form of relative-standard gradable adjectives is semantically prior to the comparative form.

So far, this framework is maximally liberal with respect to how the extension of *tall* is determined, in the sense that it depends wholly on the choice of comparison class. In addition, such liberality captures the context-sensitivity of dimensional adjectives like *tall* in their positive form, and also their lack of context-sensitivity in the comparative form.

But in fact, as Burnett says (2017, p. 60), the Delineation framework is *too* liberal. Since there are no restrictions on how comparison classes are defined, it is possible to define a comparison class X that results in an undesirable extension assignment to *tall*:




Figure 2.4: A strange comparison class.

This has the bad consequence that, relative to model M , the sentence *Vera is taller than Natasha* is true—in virtue of the fact that we can find a comparison class relative to which *Vera is tall* is true but *Natasha is tall* is not true (for example L); but the sentence *Natasha is taller than Vera* also comes out true, in virtue of the fact that we can also find a comparison class, namely B (see Figure 2.4), relative to which Natasha, but not Vera, counts as tall.

To prevent this, authors in the Delineation tradition invoke a series of *consistency axioms* ensuring that ordering relations are preserved across comparison classes. The axiom that concerns us here is the axiom of **No Reversal**, which rules out comparison classes such as B .

Axiom 1 (No Reversal) *For any model M , any two individuals $x, x' \in D$ and any comparison class $X \subseteq D$ such that $x \in \llbracket \text{tall} \rrbracket_X$ and $x' \notin \llbracket \text{tall} \rrbracket_X$, there is no $X' \subseteq D$ such that $x \notin \llbracket \text{tall} \rrbracket_{X'}$ and $x' \in \llbracket \text{tall} \rrbracket_{X'}$.*

No Reversal ensures that ordering relations are not flipped across comparison classes, thereby preventing a situation like the one depicted in Figure 2.4.³⁵

³⁵ Two other axioms are often invoked in Delineation semantics approaches (Benthem 1982; Burnett 2017; Klein 1980), although we do not need them here. Thus, for any model M , any two individuals $x, x' \in D$ and any comparison class $X \subseteq D$ such that $x \in \llbracket \text{tall} \rrbracket_X$ and $x' \notin \llbracket \text{tall} \rrbracket_X$,

Axiom 2 (Upward Difference) *For all X' such that $X \subseteq X'$, there are individuals $y, y' : y \in \llbracket \text{tall} \rrbracket_{X'}$ and $y' \notin \llbracket \text{tall} \rrbracket_{X'}$;*

Axiom 3 (Downward Difference) *For all X' such that $X' \subseteq X$ and $x, x' \in X'$, there are individuals $y, y' : y \in \llbracket \text{tall} \rrbracket_{X'}$ and $y' \notin \llbracket \text{tall} \rrbracket_{X'}$.*

Upward Difference guarantees that, if the extension and anti-extension of *tall* are non-empty at some comparison class, then these will also be non-empty at any “bigger” comparison class. On the other hand, **Downward Difference** guarantees that, if the extension and anti-extension of *tall* are non-empty at some comparison class C , then *tall* will also have a non-empty extension and anti-extension at any “smaller” comparison class that includes the individuals that fall in the extension and anti-extension of *tall* at C .

This rough characterization suffices to point out the feature of this approach that will be crucial for us: in Delineation semantics, the comparative relation is based on the application of a simple, non-gradable predicate. We begin by defining *tall* as a (partial) function that, relative to a comparison class, takes individuals to $\{1, 0\}$, and then we define *taller* by considering how that very same function applies to other comparison classes, different from the one that is provided by context.

By couching evaluative adjectives in a delineation framework, we can preserve the basic non-factualist insight about their meaning, namely that speakers use them to express practical attitudes of support/rejection, in the definition of the comparative form. To give the reader a preview, we will say that α *is good [bad]* expresses outright support [rejection] for α ; and we will say that α *is better [worse] than* β expresses support [rejection] for α over β given a choice between the two.

2.7 Recap & Conclusion

In this chapter, we considered evaluatives from a linguistic and from a metanormative point of view. From a linguistic point of view, evaluative adjectives are one among other kinds of gradable and lexically subjective predicates (along with, e.g., experiential and multi-dimensional adjectives). Moreover, the linguistic literature has not aimed to single out evaluative adjectives like *good*, *beautiful* or *cruel*, but has rather adopted a different notion of evaluativity—namely, Bierwisch’s notion of *norm-relatedness*. Adopting a metanormative perspective however, we saw that philosophers have long held that there is something that characterizes normative vocabulary and normative vocabulary only (including evaluative adjectives): its action-guiding properties. So we proposed to adopt this criterion as the defining property of evaluative adjectives that singles them out from the larger class of lexically subjective adjectives.

We then turned to presenting non-factualism, which has been defended as the family of views about the meaning of evaluatives that stands a better chance of capturing the action-guiding properties of this class of terms. In its semantic construal, non-factualism is the view that evaluative sentences express practical attitudes of support/rejection. We saw how non-factualists accommodate the action-guiding properties of evaluatives, and then we turned to the problems of non-factualism. First, non-factualism tends to be formulated in a way that is most appropriate for sentences containing evaluative adjectives in positive form. But it has trouble accounting for comparative uses of evaluatives, which fail to express outright attitudes of support/rejection. And secondly, non-factualism is subject to the Frege-Geach problem, which is the problem of explaining how such heterodox semantics can play the role that standard semantic values play.

In fact, paying attention to comparatives revealed two things: (i) first, a new incarnation of the Frege-Geach problem, which we dubbed the “sub-sentential” Frege-Geach problem: where the traditional, “supra-sentential” problem consists in explaining how evaluative sentences embed in larger constructions, the “sub-sentential” problem is the problem of how evaluative sentences are composed in the first place. And (ii), we noticed a correlation between the presence of the positive form of evaluative adjectives and their ability to express outright support/rejection. Accommodating this generalization proved challenging for a standard degree semantics, in virtue of the fact that, in degree semantics, positive form adjectives are semantically constituted by two ingredients, none of which can play the role of carrying the expression of practical support/rejection. This pushed us to abandon degree semantics and adopt a Delineation frame-

work, where the positive form of a gradable adjective is a simple, non-gradable predicate, and comparatives and other, more complex constructions are semantically built from it.

In the following chapter, we will see how to use a Delineation approach to build a non-factualist semantics for evaluative adjectives that can accommodate much of the scalar behavior of this words.

Chapter 3

Hyperplan semantics for gradable evaluative adjectives

Summary¹

In this chapter, a compositional semantics is given for a propositional language introducing a pair of positive and negative unary sentential operators representing positive and negative value, together with binary sentential operators representing equative, comparative and equal value. The semantics is based on Gibbard's notion of a hyperplan, as this framework offers the most intuitive account of the connection between evaluative concepts and action/motivation (argued for in the previous chapter). And following a delineation approach, the semantics of the binary operators is derived from that of the unary operators. In this way, we preserve the non-factualist insight that evaluative language is primarily used to express simple positive and negative attitudes towards objects of evaluation.

3.1 Introduction

This chapter introduces a compositional semantics for a propositional language containing sentential operators that represent the meaning of evaluative adjectives of positive and negative valence in their positive, comparative and equative form. Evaluative adjectives are extremely heterogenous, but they are all characterized, as we saw in Chapter 2, by their connection to action and motivation in an uniform way: unembedded sentences containing positive evaluative adjectives express practical attitudes of support; and unembedded sentences containing negative evaluative adjectives express practical attitudes of rejection. Our semantics will essentially hardwire this connection into the semantics of those adjectives.

We will do this *via* a formal tool, namely, HYPERPLANS. As we saw briefly in §2.4, hyperplans are a model-theoretical tool devised by Gibbard (2003) to represent the connection of normative predicates and deontic modals to action and motivation. Gibbard pushes an illustrative parallelism between hyperplans and possible worlds: just as possible worlds represent the content of doxastic attitudes such as belief or knowledge, hyperplans represent the content of non-doxastic practical attitudes, such as planning or intending to act. My purpose is to take

¹This chapter and the next have benefitted tremendously from extensive discussions with Carla Umbach. Any errors are mine.

this formal tool, enrich it into a full-fledged scalar system and apply it to evaluative adjectives across the board.

As we saw in the previous chapter, non-factualism was characterized by proposing that evaluative sentences, at least in their most common use, are not used by speakers to represent the world as being a certain way, but rather, they are used by speakers to communicate and operate upon their practical attitudes and those of their audience. Non-factualists cash out this intuition by saying that evaluative sentences express non-gradable attitudes of outright support and rejection of the object(s) under evaluation.

We saw, however, that there are many uses of evaluative adjectives that fail to express such attitudes: this is, in general terms, the Frege-Geach problem. But we saw that one can distinguish two incarnations of that problem, the traditional embedding or “supra-sentential” problem and what I proposed to call the “sub-sentential” one. The supra-sentential Frege-Geach problem is the traditional problem of accounting for the meaning of evaluative sentences when they are embedded under logical connectives and other sentential operators; the sub-sentential problem is the problem of accounting for the meaning of evaluative adjectives when they fail to invite an inference to the positive form—as it is the case with comparatives or some equatives, for example.

Recall that the type of non-doxastic attitudes of support and rejection that non-factualists invoke are only expressed by the positive form of evaluative adjectives, so expressing those attitudes crucially correlates with an inference to the positive form of the relevant adjective (we called this observation **POS-ATT**):

Observation 1 (POS-ATT) *An utterance U of an unembedded evaluative sentence S containing an evaluative adjective E expresses practical support/rejection if and only if U invites an inference to the positive form of E .*

For the most part, comparatives do not express the non-gradable attitudes of support and rejection that non-factualists assign as the meaning of evaluative sentences and that are expressed when we use, or somehow imply the use of, those adjectives in their positive form. Semantically though, the meaning of comparatives and the meaning of the positive form has to be related. That is, in a nutshell, the sub-sentential Frege-Geach problem.

Our approach is to assume as primitive the non-factualist insight about evaluative adjectives in their positive form, and use a delineation semantics approach to derive comparatives, equatives and other, more complex sentential constructions from the positive form. To do this, we define a simple propositional language \mathcal{L}_V enriched with the following evaluative sentential operators: a positive value unary operator \Uparrow , its negative counterpart \Downarrow , and binary operators \geq , $>$ and \approx representing equative, comparative and equal value respectively.

Semantically, expressions are evaluated in an semantic model that contains, in addition to the usual domain of possible worlds, a domain of hyperplans that select, among the set of all propositions, the preferred and dis-preferred ones. Well-formed expressions of the language are evaluated for truth at indices of evaluation, which are world-hyperplan-alternatives triplets. As an example, informally: a proposition p has positive value relative to a world-hyperplan-alternative triplet just in case the hyperplan of the index is such that p is among the preferred alternatives selected by it.

The chapter is structured as follows: §3.2 introduces the broad semantic framework on which this dissertation will be couched. §3.3 introduces the Gibbardian notion of hyperplans; it

presents Gibbard’s original proposal (§3.3.1) and introduces the type of semantic implementation of that idea that we will base our account on, namely that of Yalcin 2012, (§3.3.2). In §3.3.3, we will modify Yalcin’s view in order to overcome certain difficulties. In §3.4 we will introduce the language \mathcal{L}_V . We will start by defining its syntax (§3.4.1), and we will give a semantic structure and model, together with the relevant rules of composition (§3.4.2). Then, we will offer a toy fragment (§3.5) and we will also present the consistency axioms that generate the appropriate predictions for equative and comparative constructions, following a Delineation approach (§3.6).

3.2 Semantically hardwiring action-guidance

As we saw in the previous chapter, recent philosophical and linguistic literature has given much attention to evaluatives. Different authors have picked up on different properties of this heterogeneous class of expressions and have provided semantics designed to account for those properties. In this and the following chapter, I aim to do just that. And more specifically, in line with authors in the non-factualist or expressivist tradition, the property of evaluatives that my semantics aims to capture is their ACTION-GUIDANCE, which is the unifying property of this class of expressions (as we argued in Chapter 2).

How am I going to do that? Simply put, the procedure that I favor is to posit what we may call an “enriched” truth-conditional semantics, according to which the truth of declarative sentences at a context depends on what the world is like *together with something else*. Researchers who propose this type of semantics fall along two broad axes with respect to that “something else”: the first axis concerns what that extra semantic ingredient is (§3.2.1); the second axis concerns how to implement it semantically (§3.2.2). Let us tackle these two questions in turn.

3.2.1 Semantic pluralism vs. monism about evaluatives

First, what is the extra semantic ingredient? The foremost principle of truth-conditional semantics is that to know the meaning of a sentence is to know its truth-conditions, that is, ‘what the world would have to be like for it to be true’ (Heim and Kratzer 1998, p. 1). Any enriched semantics starts off from the claim that knowing ‘what the world is like’ is not enough to determine the truth of certain sentences at a context. This has been one of the basic motivations for proposing various types of enriched semantics for (among others) sentences containing temporal information (tenses), vague expressions, modals, certain indexicals (‘now’, ‘here’, ‘I’) and predicates of personal taste (PPTs for short; i.e., ‘tasty’, ‘fun’). With respect to epistemic modals for instance, the idea is that the truth of sentences containing epistemic modals depends, not just on what the world is like, but also on the available information about the world at a context. Thus, the correct semantics for those sentences should build in a sensitivity to states of information (over and above the usual possible world-sensitivity). Similarly, with respect to PPTs the idea is that the truth of sentences containing PPTs depends, not just on what the world is like, but also on someone’s *taste*. Therefore, the correct semantics for PPTs should incorporate a sensitivity to tastes (this is one option among others, as we will see).

Now, since we aim to offer a semantics for evaluatives, the question for us is the following: for any evaluative sentence E , what extra semantic ingredient is required? More precisely, what

does the truth of E depend on, beyond what the world is like? Simplifying somewhat, we can distinguish between MONIST and PLURALIST answers to that question. The monist attempts to find a single semantic ingredient with which to build a unified semantics for evaluatives. The pluralist on the other hand denies that there is a single semantic ingredient that the truth of evaluatives depends on, and proposes a plurality of extra ingredients to account for different sub-classes of evaluatives.

Pluralists about evaluatives include philosophers like MacFarlane (2014). MacFarlane, whose traditional focus has been PPTs, future contingents and various types of modals, does not posit a single semantic parameter to cover all the expressions that require an “enriched” semantics. For example, he holds that the truth of sentences containing PPTs depends on a standard of taste, and that the truth of sentences containing epistemic modals depends on a state of information. Presumably (although he does not develop this), he would claim that the truth of sentences containing *moral* adjectives does not depend on any of those things, but rather on a moral standard. Given that evaluatives include PPTs as well as moral adjectives, MacFarlane would not posit a single semantic parameter for evaluatives, but a plurality of them.

The problem with the pluralist approach is that it loses sight of the commonalities between evaluatives. That is, if one thinks (as I do) that evaluatives deserve a special semantic status in virtue of the linguistic properties that they *share* as a class, namely their action-guidance, then the pluralist’s approach is unnecessarily specific—as it posits different semantic parameters for otherwise similar expressions. That is, the pluralist will say that the meaning of PPTs depends on a standard of taste and that the meaning of moral adjectives depends on a moral standard, even though the property in virtue of which these expressions deserve special status is one and the same.

Monists, by contrast, include linguists like Lasersohn (2005) or Stephenson (2007a), who put forward the notion of a JUDGE as the extra semantic ingredient.² Barring differences in the scope and implementation of their proposals, monists claim that the truth of sentences containing a set of expressions that includes PPTs—and epistemic modals in the case of Stephenson—depends, not just on what the world is like, but also on a judge. The notion of a judge is sufficiently underspecified to encompass taste, moral and aesthetic standards as well as available worldly information, thus making it possible to give a unified semantics for a great variety of expressions.

Now, whereas MacFarlane’s pluralist approach is too specific, Lasersohn and Stephenson’s monist approach is too general: there are expressions that arguably require a judge-dependent semantics even though they do not count as evaluative. To wit: epistemic modals as well as experiencer-sensitive adjectives like *difficult*, *salty* or *soft* (§2.2.2). The truth of sentences containing these adjectives depends on someone’s state of information or on their experience, and thus on a judge. But these are not evaluative expressions, as they are not action-guiding. Thus, the notion of a judge is too broad to capture evaluativity, the way we have characterized it (we come back to this issue in Chapter 5).

We can now return to our question: for any evaluative sentence E , what does the truth of E depend on, beyond what the world is like? Moral or taste standards are too specific (because some evaluatives are insensitive to them) and judges are too general (because some non-evaluatives

²Coppock (2018), taking inspiration from Kölbel (2002), proposes a monist view where, rather than enrich semantics with an extra parameter that allows for, say, taste-sensitivity, possible world semantics is simplified to *outlook*-semantics, where OUTLOOKS are finer-grained objects than worlds that are sensitive to worldly and perspectival matters.

are sensitive to them). But then, how can we answer that question? I want to propose here that the correct answer is what we may loosely call *practical intentions, preferences or plans*. And moreover, I want to propose that we model this parameter using Gibbard’s notion of hyperplans. The basic thought is that the truth of any sentence containing evaluatives (such as (3.1)-(3.4) below) depends, not just on how the world is—that is, on the properties of the act of volunteering, sushi or a picture—but in addition to that, it depends on the practical intentions, preferences or plans of the people that use those sentences.

- (3.1) Volunteering is good.
- (3.2) Sushi is tasty.
- (3.3) This picture is beautiful.
- (3.4) My bike is better than yours.

That is the approach that is defended in this work. It is a monist view, according to which the single “extra” ingredient that we need to introduce in our semantics for evaluatives are practical intentions, preferences or plans (we’ll talk of *plans* for simplicity’s sake).

3.2.2 Semantic contextualism vs. relativism about evaluatives

Let’s turn now to the second axis along which these proposals differ, namely the various ways of implementing the type of semantics that I am interested in. Adopting plans as the new ingredient in our “enriched” semantics, the question arises as to how to implement a plan-sensitive semantics. This takes us into the contextualism-relativism debate (see Stojanovic 2017a,b).

We can organize different theoretical options around two questions. The first one is: how do different plans affect the meaning of evaluative sentences? There are two principal answers to that question. The most straightforward is: by adopting some kind of CONTEXTUALISM, which is the view that evaluatives encode a contextual sensitivity to plans as part of their lexical meaning.³ The standard way to cash this out is to claim that those expressions behave like indexicals (1989) (some authors call this view ‘indexical’ contextualism). Just as the meaning of a word like *today* can be factored out in a constant, functional description of a day relative to the context of use (its Kaplanian CHARACTER, i.e., *the day of the context*) and a variable day depending on the context (its Kaplanian CONTENT, i.e., *February 26th*), the meaning of a word like *good* could be factored out into a Kaplanian character (i.e. *supported by the plan of the context*), and the actual varying plan of whoever are taking part in the conversation.

Contextualists would hold that an utterance of a sentence like (3.1) states a property of a contextually-determined plan, namely that it is a plan to support the act of volunteering. Thus, (3.1) would be truth-conditionally equivalent to a sentence like (3.5):

- (3.5) Volunteering is supported by the plan of the context.

A further question then arises: what is *the plan of the context*? Is it the speaker’s? The hearer’s? Both of them? Each of these options faces different challenges. Briefly: if the speaker describes

³Contextualism about different types of expressions has been vindicated by authors like DeRose (1992; 1991), Kratzer (2012), Stojanovic (2007), Glanzberg (2007), Sæbø (2009), Moltmann (2010) von Fintel & Gillies (2011) or Silk (2016, 2017).

either their own or their hearer’s plan, then disagreement vanishes, as different speakers engaging in disagreement would be expressing semantically compatible propositions, and therefore talking past each other.⁴ If, on the other hand, speakers are describing their mutual plan so to speak, there is a sense in which their utterances are unwarranted—most of the time, we cannot know what the evaluative stance of our interlocutor is and yet that does not seem to prevent us from making value judgments (a thorough review of these problems for contextualism can be found in Lasersohn 2005, esp. §§3–4).

Furthermore, there is a general difficulty for any version of contextualism that was raised early on by proponents of expressivism like Ayer (1946): it purports to understand evaluative sentences as though they were descriptions of plans, rather than avowals thereof. This contrast was presented briefly in §2.3.2, where we discussed the difference between expressing a non-doxastic attitude and expressing a belief that one has a non-doxastic attitude. Since contextualism incorporates a reference to plans as part of the regular, descriptive proposition expressed by an utterance of an evaluative sentence at a context, evaluative sentences turn out to do the latter in a contextualist semantics: they express beliefs that one has a non-doxastic attitude, in this case a plan, rather than the endorsement or avowal of such plan. To see this more clearly, consider the contrast between the following sentences:

(3.5) Volunteering is supported by the plan of the context.

(3.6) Let’s support volunteering.

The basic criticism directed by non-factualists at contextualists is that, by assigning sentences like (3.1) a semantics along the lines of (3.5), they lose track of the fact that the communicative function of evaluative sentences seems much closer to sentences like (3.6)—where a plan is proposed or endorsed by the speaker for the hearer’s uptake—than to sentences like (3.5)—where a plan is merely described. Any view that incorporates a reference to plans as part of a regular, descriptive propositional content is prey to this criticism (see Gibbard 2003, 85 and ff for discussion of this point).

The least straightforward answer to the question ‘how do different plans affect the meaning of evaluative sentences?’ (and one that avoids the previous difficulty) is to say that different contexts do not change the Kaplanian content of (3.1), which is stable across contexts. Nevertheless, again following Kaplan, one and the same Kaplanian content receives different truth values at different CIRCUMSTANCES OF EVALUATION (Kaplan 1989). So a different way of incorporating plans into our semantics is to consider them part of the circumstances of evaluation of sentences. This is semantic RELATIVISM (Coppock 2018; Egan 2012; Kölbel 2003; Kölbel 2002; Lasersohn 2005; MacFarlane 2014; Stephenson 2007b).

The relevant analogy for this family of view is the relativity of propositions to possible worlds. Consider a regular, descriptive sentence like (3.7):

⁴There are ways out of this objection; a possibility that has recently gained traction is that speakers expressing compatible contextualist propositions could nevertheless be having a *metalinguistic* disagreement about how to use evaluative terms. The basic idea behind this position is that expressing compatible propositions need not entail that speakers cannot have a dispute, and therefore, that contextualism can, after all, accommodate disagreement (C. Barker 2013; Plunkett and Sundell 2013; Umbach 2016, see). A different possibility is that, in conversations about evaluative matters, it is presupposed that speakers share the relevant evaluative parameter. This is called a PRESUPPOSITION OF COMMONALITY (Blome-Tillmann 2009; López de Sa 2015). If speakers presuppose that they share a plan, then their disagreement could turn on the content of their shared plan. These positions face further problems. For discussion, see Stojanovic 2017a.

(3.7) Volunteering is widespread.

(3.7) expresses one and the same proposition at every context of use. However, it has different truth-values at different possible worlds. In the actual world, where volunteering is not that widespread, that proposition is (arguably) false. But things might have been different, and under different circumstances, the proposition expressed by (3.7) might be true. Importantly, note that the fact that we can evaluate one and the same sentence at different circumstances of evaluation does not imply that (3.7) has a different Kaplanian content at each of those circumstances of evaluation, but only that it has a different truth-value. Authors who adopt this option hold that the truth of an evaluative sentence like (3.1) is sensitive to plans in just the same way that the truth of (3.7) is sensitive to possible worlds. (3.1) always has as its content the proposition that volunteering is good, *simpliciter*. But that proposition can be true or false depending on the plans of the context in which we evaluate it.

A further axis of variability concerning semantic relativism is: what is the appropriate context from which to evaluate the truth of a given evaluative sentence? The most obvious answer is: the context of *use*. If relativism is defined in this way, then the context-sensitivity of evaluatives is not substantially different from other forms of context-dependency attested in natural language (such as e.g., sensitivity to times). MacFarlane (2014) advocates for another possibility: rather than holding that the appropriate context from which to evaluate the truth of an evaluative sentence like (3.1) is the context in which that sentence is *uttered*, we can consider the possibility that the right context is the context in which an utterance of such sentence is *assessed*. This proposal has been dubbed ASSESSMENT RELATIVITY. This form of context-dependency is radically different from the kind of context-sensitivity attested in natural language, and therefore more difficult and problematic to defend.

Nonetheless, that is the type of position that we want to defend in this work. That being said, it is not my objective to launch a full-fledged defense of MacFarlane's views here, nor to further explore its virtues. Moreover, an important point of contrast between his view and mine is that MacFarlane's view is pluralist, whereas the view defended here is a monist view, according to which the single semantic parameter that a semantics for evaluatives requires is a plan parameter. But before moving on, let me stress briefly the aspect that I consider most attractive of this type of semantics. This aspect arises when we consider a disagreement scenario such as the following:

- (3.8) a. Annag: Volunteering is good.
b. Liam: No, it's not.

Suppose that we eavesdrop on that conversation. Consider Annag's utterance, and assume that the correct semantics for *good* is relativist. This means that the truth of (3.8a) depends on someone's plans, even though plans are not part of the proposition expressed by her utterance. Now, the question arises as to whose plans are appropriate for assessing the truth of (3.8a). In a standard relativist framework (that of Lasnik 2005; Stephenson 2007b or Kolbel 2002), the relevant plans would be Annag's. But suppose that we, the eavesdroppers, disagree—we think that volunteering is not a good thing to do. Our intuition about Annag's utterance, then, is that she said something false (and that Liam said something true). Assessment relativity respects this intuition, and predicts that, in the situation just described, Annag's utterance expresses a falsity. This is because, from the eavesdropper's point of view, the correct plans relative to which to assess Annag's utterance are the eavesdropper's, not Annag's.

By making the truth-value of evaluative sentences always depend on the assessor’s plans, this type of semantics guarantees two things: first, the truth-value of any evaluative sentence will not depend on who utters that sentence, but on the plans of whoever is evaluating it. In practice, this means that plan-sensitivity works like possible world-sensitivity: just as a descriptive sentence (such as *volunteering is widespread*) is true or false at a possible world regardless of when, where or by whom it is uttered, an evaluative sentence is true or false relative to a plan regardless of when, where or by whom it is uttered. And secondly, by adopting a semantics whereby plans and possible worlds can vary independently of each other, the truth of evaluative sentences does not depend (merely) on how the world is: holding the possible world fixed, different choices of plans will make evaluative sentences true or false. This respects the basic assumption behind an “enriched” semantics, namely that knowing what the world is like is not enough for determining truth conditions.

Let us note, however, that assessment sensitivity approaches are not free from controversy. Stojanovic (2007, 2017a) makes a compelling case that this type of semantics has trouble accommodating the intuitions behind disagreement dialogues like (3.8), which is one of the main purported advantages of this view advertised by MacFarlane. Disagreement is not the focus of this work, and intuitions about truth are markedly shaky, so we will leave it at this.

3.2.3 Action-guidance in plan-semantics

In the semantics to be presented in the following sections, the truth of evaluative sentences in contexts of use will depend on plans. Plans, as we will see shortly, will be formalized as sets of practical instructions that tell us what to do in certain situations. For a certain state of affairs to be good relative to a plan, will be for that plan to pursue that state of affairs over various alternative states of affairs. Importantly, to say that a state of affairs is good is not to describe a plan; it is to endorse it.

An urgent question, then, is the following: does this view predict the action-guiding properties of evaluatives, and if so, how? Recall our diagnosis from §2.3.1: evaluatives invite a *practical* inference about the practical intentions or dispositions of speakers:

- | | | |
|--------|-----------------------------|-------------------------------------|
| (2.35) | a. Today’s exam was good. | ↪ the speaker supports today’s exam |
| | b. The floor is still ugly. | ↪ the speaker rejects the floor |
| | c. The soup is awful. | ↪ the speaker rejects the soup |

Non-factualists then claim that these sentences express a non-representational component that stands for an action-guiding or practical attitude. But how does this come about, and does can a plan-sensitive semantics really accomplish this? It’s best to couch the answer to that question in terms of a middle ground between two possible positions that one could adopt *vis-à-vis* action-guidance. Recall, first, that the action-guiding inferences associated with evaluatives are *defeasible*: the fact that Natalia thinks that it is better to volunteer than to donate does not guarantee that she will volunteer rather than donate; the fact that Matheus thinks that today’s exam was good does not guarantee that he will in fact recommend or support the exam in any way. Nonetheless, the inference is there, and there is a strong reason to think that it depends exclusively on the presence of evaluative language. Now, suppose that, wanting to make good on the presence of that inference, we adopted a semantics for evaluatives according to which those sentences would be false if the relevant actions were not performed. That is, according

to this type of semantics, if Natalia does not volunteer rather than donate, then it would be false that she thinks that volunteering is better than donating; and if Matheus does not in fact recommend or support the exam in any way, then he does not think that the exam was good. In light of the defeasibility of those practical inferences, such a semantics would be too strong.

Consider, on the other hand, a (contextualist) semantics according to which to utter an evaluative sentence is to describe a plan of action. Since describing a plan of action is independent from carrying it out, this would accommodate the defeasibility of the action-guiding inference. However, and setting aside the usual problems of contextualist semantics to accommodate disagreement (or, more generally, cross-contextual) phenomena, this approach now appears too weak *vis-à-vis* action-guidance. That is precisely *because* describing a plan is independent from carrying it out. At most, one could say that for a speaker to describe a plan of action *as her own* could be understood as having the performative effect of communicating her adoption of that plan. If that were the case, then a contextualist plan-semantics for evaluative sentences could have a better chance at accommodating action-guidance: by saying something like *volunteering is good*, Natalia would be describing her plan of action, but might be taken to performatively communicate her endorsement of that plan. And then, it is sensible to infer that she will carry it out. Perhaps this view can be developed, and it does not require a relativist/non-factualist semantics requiring anything beyond standard possible-world semantics. Regardless, the option that I propose here is capable of similar predictions *vis-à-vis* action-guidance without the usual problems of contextualism with regards to disagreement phenomena.

Let us think that to assert an evaluative sentence is not to describe a plan of action, but to endorse it. As we have said before, according to the type of semantics that we favor saying that volunteering is good resembles something like an invitation, or a proposal, to adopt a plan of action. If that is the case, then the inference that the speaker will carry out the relevant plan seems like a sensible, yet defeasible inference to make. If I endorse a plan, it is natural to infer that I will carry it out. A final question then arises: is the action-guiding inference, in the type of plan-sensitive semantics that I am advocating, a semantic or a pragmatic inference? In contemporary linguistic research, this question might be more accurately posed as: does this inference fall in the categories traditionally distinguished for natural language inferences (namely entailments, presuppositions or various kinds of implicatures), or is it a different phenomenon altogether? Answering this question would require much more than what this dissertation aims to cover, but we can say the following: the fact that this inference appears only associated with evaluative or normative language seems to suggest that it is a very *sui generis* type of phenomenon. Moreover, a brief overview of basic properties of natural language inferences suffices to show that it does not fall clearly under any well-established category: it can be defeated, which speaks against treating it like a standard entailment. It does not project (under negation, for instance), which suggests that it is not a presupposition (nor a conventional implicature). And it is hard to devise the conversational principles behind its presence, which suggests that it is not a conversational implicature. Barring a more thorough investigation, we will consider it a *sui generis* inference prompted by the peculiar act of asserting an evaluative sentence which, according to the semantics proposed here, amounts to the endorsement of a plan of action. The inference that the speaker will carry out that plan falls out naturally from such an act of endorsement; but that inference can be defeated if outweighed or undercut by countervailing motivations. We now turn to a fuller exposition of the specific semantic parameter that is avowed in our semantics, namely “Gibbardian” plans.

3.3 Hyperplans

As discussed briefly in §2.4, hyperplans are a conceptual tool proposed by Gibbard (2003) in order to model the connection between normative language and action. My purpose is to develop this conceptual tool into a full-fledged scalar semantics for evaluative adjectives. To do this, I proceed in three steps. I start in §3.3.1 by presenting the notion of a hyperplan, proposed by Gibbard in order to model the content of normative *judgment*. A key feature of this account is that hyperplans are considered model-theoretical primitives, just like the possible worlds of intensional semantics. Next, in §3.3.2 we will see how, following Yalcin 2012, Gibbard’s view about normative judgment can be cashed out in standard intensional semantics as a view about normative *sentences*. This requires taking possible worlds as the sole semantic primitive and defining hyperplans as functions over sets of possible worlds. However, there are two problems with Yalcin’s view: the first is that, the way the semantics is set up, it predicts that anything that is taken to be known is normatively required, which is undesirable (this is a version of a problem raised by Frank 1997; Zvolenszky 2002). The second problem is that Yalcin’s semantics is not scalar, so if we want to turn this into a view about evaluatives, we need to bring scalarity into the picture. Overcoming this problem will motivate us to change Yalcin’s definition of hyperplans, as we will see in §3.3.3.

3.3.1 Gibbard 2003’s proposal

In his 2003, Gibbard sets out to characterize a basic normative concept, that of something *being the thing to do*. For Gibbard, to judge that an action is the thing to do is to plan to perform that action under certain circumstances. Various predicates denote that concept in natural language; and Gibbard takes deontic and normative sentences such as *x ought to φ* , *φ -ing makes sense* or *φ -ing is the thing to do* to all express the speaker’s plan to φ under certain circumstances.

For example, to judge that Camila ought to start packing expresses the speaker’s plan to start packing if she were in Camila’s circumstances. This view is *expressivist*, or more generally *non-factualist*, because it relies crucially on the idea that declarative sentences that contain normative vocabulary do not, or do not only, express doxastic mental states, such as states of belief or knowledge; rather, those sentences express non-doxastic states of *planning*.

Gibbard, however, does not rely too much on the intuitive notion of a plan—we can think of plans as practical commitments or strategies for behavior.⁵ The crucial aspect of his account is that the content of a normative judgment is (partially) represented using what he calls HYPERPLANS, a modelling tool that is meant to be the non-doxastic counterpart of the possible worlds of intensional semantics: just as the content of doxastic mental states is represented *via* sets of possible worlds, Gibbard proposes that planning states be represented as sets of hyperplans. And in order to understand this notion, it is useful to keep in mind the parallelism with possible worlds. A possible world is a maximally specific way the world could be. Alternatively, one might think of a possible world as a maximally opinionated doxastic state. Analogously, a hyperplan is a maximally specific plan of action, or alternatively, a maximally decided state of planning, or as Gibbard calls it, a HYPERSTATE (Gibbard 2003, 48 and ff.)

To stress the analogy with possible worlds, Gibbard takes hyperplans as primitive. But in order

⁵As suggested by Charlow 2014. The concept of a plan has an obvious temporal component that we will abstract away from.

to understand the notion better, it is perspicuous to think of a hyperplan as a set of conditional imperatives (Field 2009, p. 258), with a centered proposition—centered on the agent of the hyperplan—in their antecedent and an imperative, addressed at that very agent, as its consequent.⁶ A hyperplan lists absolutely all the situations that the agent could find herself in, and dictates what to do in each of them. It looks like a set of instructions:

$h =$ {if you see your friend Matt, say hi;
if you miss the bus, take the metro;
if you get hungry before lunch, eat an apple;
if you finish this chapter, plan a vacation;...}

Table 3.1: A hyperplan h whose instructions are spelled out as a set of conditional imperatives.

Importantly, Gibbard’s is a view about the content of judgment and not directly about any particular fragment of language. Just as the descriptive content of a judgment is to be modelled as a set of possible worlds, i.e. a classical proposition, its normative content is to be modelled as a set of hyperplans, that is, as a PLAN. A plan can be thought of as a set of hyperplans that agree on what to do in certain situations but not others; or alternatively, we can think of hyperplans and plans as *total* and *partial* sets of instructions, respectively.

However, hyperplans and possible worlds do not live separate lives; judgments can be descriptive, normative or a mixture of both. Hence, rather than modelling normative content using sets of hyperplans and modelling factual or descriptive content using possible worlds, Gibbard offers a uniform account of the content of judgment in terms of sets of world-hyperplan pairs, or as he calls them, FACT-PRAC worlds (2003, p. 47). A fact-prac world is a pair consisting of a possible world and a hyperplan. The content of a judgment is given in terms of the fact-prac worlds compatible with it. To see how this proposal works, we can rehearse truth conditions for various kinds of judgments:

1. (*Descriptive*): A subject a judges that p just in case every fact-prac world $\langle w, h \rangle$ compatible with a ’s judgment is such that p is true at w , according to h relative to the actual circumstances.

This is true if and only if a believes that p , regardless of what a ’s planning state looks like. That can be read off the truth conditions just stated: since the adjunct *according to h relative to the actual circumstances* is vacuous, judging that p is trivially compatible with any hyperplan. So the fact-prac worlds compatible with this judgment are those that result from pairing each of a certain subset of possible worlds (those in which p is true) with every hyperplan.

2. (*Normative*): A subject a judges that it ought to be the case that p just in case every fact-prac world $\langle w, h \rangle$ compatible with a ’s judgment is such that it ought to be the case that p is true at w , according to h relative to the actual circumstances.

This is true if and only if a plans to make p true in the actual circumstances, regardless of what her beliefs are. This time, the adjunct *according to h relative to the actual circumstances* is not vacuous, and thus a condition is imposed on the subject’s planning state. By contrast, there is no condition on the subject’s belief state, since the *ought* “trumps” the statement that p is true

⁶Gibbard is explicit (2003, p. 51) that hyperplans are centered on an agent, and for this reason he recruits Lewis’ notion of a centered world (D. Lewis 1979a) A centered world is a triple formed by a world, an individual (its center) and a time.

at w : the condition is not that the subject's fact-prac worlds $\langle w, h \rangle$ be such that p is true at w , but rather that p ought to be true at w according to h relative to the actual circumstances. The fact-prac worlds compatible with this judgment result from pairing every possible world with each of a certain subset of hyperplans (those according to which p ought to be the case in the actual circumstances).

Two points about this proposal: first, note that normative judgment is spelled out in terms of quantification over fact-prac worlds. This results in the familiar inter-definability of obligation and permission (p is not obligatory iff not- p is permitted):

- a. (p is not obligatory): a judges that it is not the case that it ought to be the case that p just in case not every fact-prac world $\langle w, h \rangle$ compatible with a 's judgment is such that it ought to be the case that p is true at w , according to h relative to the actual circumstances.
- b. (not- p is permitted): a judges that it is not the case that it ought to be the case that p just in case some fact-prac world $\langle w, h \rangle$ compatible with a 's judgment is such that it ought to be the case that p is false at w , according to h relative to the actual circumstances.

Secondly, note that even though making a normative judgment is not to believe that anything is the case, it requires that the subject has certain beliefs about the world. To see this, consider an example: Camila is deliberating whether to start packing. She looks at her phone, sees that it is late, and then judges that she ought to start packing. Her plan to start packing is contingent on her belief that it is late. If it were not late yet, she would not start packing. This illustrates that normative judgments *rely on*, but do not *offer*, factual information. This sensitivity to available information is captured by the qualification *given the actual circumstances* in the truth conditions for attributions of normative judgments. This captures the fact that deontic modals are information-sensitive, in the sense of Kolodny and Macfarlane 2010; Yalcin 2012 and many others.

Finally, this is how Gibbard's proposal accommodates mixed judgments:

3. (*Mixed*): A subject a judges that, if p , it ought to be the case that q just in case every fact-prac world $\langle w, h \rangle$ compatible with a 's judgment is such that p is false at w , according to h or every fact-prac world $\langle w', h' \rangle$ compatible with a 's judgment is such that it ought to be the case that q is true at w' , according to h' .

This is true just in case a believes that not- p or plans to make q true.⁷

Before moving on, let us flag a problem that this type of view about normative judgment faces. Note that we often judge what other people ought to do. But we cannot understand this in terms of a plan that *we* adopt. The reason for this is, simply, that we can never those other people. Gibbard's view, as we have presented it, seems to apply most clearly to self-ascribed normative judgments, such as thinking that one ought to pay taxes, that it's rational to hide if a stranger enters one's house or that it makes sense to recycle. But understanding third personal ought-judgments, such as my judging that Camila ought to pack, in terms of plans is not so straightforward. According to the truth conditions given, if I judge that Camila ought to pack, I endorse a plan to pack if I am in Camila's circumstances. But obviously, I cannot literally be in Camila's circumstances—because I am not Camila! Gibbard considers

⁷This gloss makes any mixed judgment equivalent to the relevant Boolean combination of the judgments of its proper parts. That is, judging *that, if p , then it ought to be the case that q* is equivalent to either judging *that not- p* or judging *that it ought to be the case that q* . This can rightly be called into question, although this is not a particular commitment of Gibbard's proposal (see Charlow 2015, §3.3 and ff).

this problem (2003, 48 and ff), and his reply is to understand talk of circumstances in qualitative terms. That is, even though I cannot be in Camila's circumstances, I can very well find myself in qualitatively similar circumstances (in fact, my circumstances can be similar to the point that the only difference between Camila's circumstances and mine is the lack of identity between Camila and me). So when I judge that Camila ought to pack, I am avowing a plan to pack if I find myself in circumstances *similar* to Camila's.

This helps with an additional problem, which is the fact that we can also judge the actions of, say, fictional characters. Since neither these nor their circumstances are actual, there is a similar problem in avowing a plan to behave in a particular way when in a fictional character's circumstances: after all, those circumstances will never arise. That is, I might judge that Neo ought to have taken the blue pill, but that cannot mean that I plan to take the blue pill if I find myself in Neo's circumstances, because I know that I will not be in those circumstances. And this is not just because I am not Neo, but moreover because those circumstances are fictional (whereas Camila's circumstances are part of our world). Nonetheless, the same strategy that we considered before can be applied here: when I judge that Neo ought to have taken the blue pill, I am avowing a plan to take the blue pill in similar circumstances, where those circumstances do not include the fact that I am not Neo nor the fact that those circumstances are part of a fictional story. What I mean is that I plan to take the blue pill if I face a similar choice, simply put. The correct way of interpreting the phrase 'the actual circumstances', therefore, is in terms of a property, or more formally, a centered world. This is why Gibbard invokes D. Lewis 1979a as the correct way of modelling the content of hyperplans. A centered world is neutral with respect to its center's identity as well as to whether the property that it attributes to the center is instantiated or not.

There is, however, a further difficulty that Gibbard's simple recourse to a Lewisian semantics cannot overcome. We often judge that things *beyond anyone's control* ought to be. For example, we might judge that an earthquake ought not have happened. But this surely cannot mean that we avow a plan to prevent an earthquake from happening! How are we to understand these ought-judgments? There are two things we can say. First, one might question the naturalness of those judgments. After all, it is a little strange to judge that things beyond anyone's control should not happen. Does not *ought* imply *can*? However, if one does not have any problem with this type of speech, there is still the possibility that we are evaluating things from a "God's eye" perspective. Perhaps what we mean when we say that the earthquake ought not have happened is that, if there were a god, he acted wrongly in causing the earthquake. Thus, we avow a plan to not cause an earthquake if we find ourselves in god's circumstances (and more importantly than considerations about the god's frame of mind, here 'god's circumstances' probably include considerations about those in the receiving end of the earthquake, that is, considerations about their innocency and so on).

Note that some evaluatives involve in their meaning the very idea that the object of evaluation is beyond anyone's control. Take the case of *lucky*, for instance. Intuitively, *lucky* carries a positive evaluation. But it is odd to understand that evaluation in terms of plans, as *lucky* also conveys that whatever is positively evaluated also happened at chance. Consider the following sentence (Chris Barker, *p. c.*):

(3.9) I was lucky to be born in the XXth century.

Will a plan-semantics be able to accommodate *lucky*? I think so, and that the chancy and evaluative components of *lucky* can co-exist. Suppose that we can distinguish the following

three components in the meaning of *lucky*:

(3.9) I was lucky to be born in the XXth century.

- | | |
|--|----------------------|
| a. \leadsto I was born in the XX th century | veridical component |
| b. \leadsto It is good that I was born in the XX th century | evaluative component |
| c. \leadsto It is random that I was born in the XX th century | chancy component |

Now, the proposal would be to understand the evaluative component of *lucky* much in the same way as one can understand the evaluativity of a sentence like *the earthquake ought not have happened*: we evaluate things from a god's perspective, if you will, and we express a plan to choose the relevant outcome over certain alternatives (for instance, being born in the XIXth century). This is compatible with the observation that my birth was a chancy event, as (3.9) also says.⁸

For yet another example, I think that the extreme right ought not have won any seats in the Spanish parliament in the last general election. But this need not mean that I avow a plan to prevent them from taking part on the election, or to take away their seats (although I might). Rather, what I mean is that nobody ought to have voted for them. So, I avow a plan for no one to vote for them. This case can be accommodated by Gibbard's proposal, if we understand those ought-statements as covertly quantified ought-statements about what a set of individuals ought to do. I will use a similar example later in this chapter, so I might as well clarify a possible misunderstanding here. We will use as an example (containing evaluative adjectives rather than *ought*, but it doesn't matter) the judgment that Kanye ought to get two out of the two nominations that he aims to get in the following Grammys. Again, this need not be thought of as a proposal to adopt a plan such that, when faced between a choice between giving to Kanye zero, one or two nominations, one chooses to give him two—because that choice seems to imply that the agent is in the position to freely choose that outcome. To avoid this, all we need to do is to consider such a judgment as covertly quantified over: what we mean when we judge that Kanye ought to get 2/2 nominations is that we endorse a plan for each individual who takes part in the Grammy vote to choose appropriately.

Nonetheless, one might think, all of these troubles would be avoided if we switched from talk of plans to talk of *preferences*. In contrast to plans, preferences are not constrained by actual possibilities. There is nothing strange about preferring that Camila packs, that Neo takes the blue pill, that the earthquake did not occur or that Kanye gets as many Grammy nominations as he wishes. Talk of preferences is suggestive, but it will not be adopted it here, for the following reason. Consider what it is for something to be preferred by someone. Arguably, this is a relation that holds between something (the object of preference) and an intensional individual—preference is an intensional relation, much like belief or desire. We can go on to think what the conditions of application of this relation are. And here we face two clear choices. On one construal, preference bottoms out in certain private happenings occurring in the individual's head. I prefer ice-cream to waffles just in case I am in a certain state of mind involving ice-cream and waffles. But on a much more natural construal of preference, preferring something is a relation that holds between an object and an individual in virtue of the individual's *dispositions* to act in certain ways. That is, for me to prefer ice-cream to waffles is a matter of my disposition to choose ice-cream over waffles when the occasion arises. In this

⁸This being said, it bears mentioning that *lucky* and similar adjectives that admit constructions like (3.9)—many of which are evaluative—raise a host of interesting issues. See e.g., C. Barker 2002, §4 and ff, and Karttunen 2014 for a thorough discussion.

second sense, preference collapses into something very much like a plan. For this reason, we will stick to talk of plans—even though some might find it somewhat unnatural.

A further issue concerns to what extent one might fail to carry through a plan or a preference. This is a thorny question—as we saw in §2.3.1, the practical inference associated with evaluatives (or normative expressions more generally) is a *defeasible* inference, in the sense that it can be outweighed or undercut by countervailing motivations. But without going into this issue, we might think of the endorsement of a plan in terms of the adoption of a public commitment to behave in certain ways. This idea has been defended in various realms, and Gibbard himself proposes it when he describes the hyperplan framework (2003, p. 45). More recently, Krifka (2014), taking inspiration from Recanati 1987; Szabolcsi 1982, defends a view according to which speech acts in general can be thought in terms of public adoption of commitments of various kinds (see also García-Carpintero 2015). And inferentialists of various persuasions have insisted that semantic content at large ought to be thought in terms of the adoption of commitments to reason in various ways (this arguably includes practical reasoning, c.f. Brandom 1983; Chrisman 2008; Frápolli and Villanueva Fernandez 2013). This type of view seems amenable to the way in which plans are thought of here. To avow or endorse a plan can be thought as the public adoption of a commitment to carry out that plan, even though this commitment can be undercut or overridden by other considerations.

3.3.2 Yalcin 2012’s intensional implementation

In Gibbard’s approach, hyperplans are primitive. That serves his purpose of pushing an illustrative parallelism with intensional semantics, as we have seen: just as one uses possible worlds to model the logical relations between descriptive judgments, one can use hyperplans to similarly model logical relations between normative judgments. However, Yalcin 2012 argues that Gibbard’s approach obscures two important points about hyperplans: (i) hyperplans can be defined in terms of possible worlds and (ii) hyperplans can be integrated in compositional semantics.

Importantly, those two observations are independent: one can give a compositional semantics for normative language taking hyperplans as primitives, and one can cast the previous Gibbardian story about normative judgment using a non-primitive notion of hyperplan. Nonetheless, by both defining hyperplans in terms of possible worlds and giving a compositional semantics, Yalcin aims to show that the Gibbardian view on content is compatible with standard intensional semantics (see also Charlow 2015, 6 and ff).

Let us review how this works.⁹ Yalcin’s proposal is to define hyperplans as functions from sets of worlds to sets of worlds—making good on point (i) above—and have declarative sentences denote sets of fact-prac worlds rather than sets of worlds—making good on (ii). A hyperplan, in this view, is no longer a primitive. Rather, it is a function from a set of possible worlds to a non-empty subset thereof where the impermissible options at the former set are false. For example, a hyperplan is a function that takes as argument the proposition that it is late, and returns the proposition that Camila starts packing.

⁹Another signpost: I will ultimately settle for a view in which hyperplans are not primitive, although they are not defined in the same way that Yalcin defines them—for reasons that will soon become clear. In particular, whereas Yalcin defines hyperplans as functions from sets of worlds to sets of worlds, I will define them as functions from sets of *sets* of worlds to sets of *sets* of worlds.

Implementing this in a standard intensional and compositional semantics is relatively straightforward. All that one needs to do is assign as the semantic value of declarative sentences fact-prac worlds rather than worlds, and the semantics essentially falls out from the Gibbardian proposal about descriptive/normative judgment just reviewed.¹⁰

To see this, assume a simple propositional language containing a set of sentential constants p, q, r, \dots representing a set of simple sentences of English (*it's raining*, *Paris is in France*, *Camila packs*, ...). These are the descriptive sentences. Let us add to this language a unary sentential operator, *ought*. Where φ is a descriptive sentence, sentences of the form *ought* φ are the normative sentences (there are Boolean connectives too, but we ignore them).

On the semantic side, we add to our intensional semantic structure a set of hyperplans and an information state. A semantic structure S consists, then, of the usual set of possible worlds W , a set of possible hyperplans H and a set of possible information states E .

Definition 5 (Yalcinian semantic structure) $S = \langle W, H, E \rangle$

As we said, a hyperplan is a function h from subsets to subsets of W such that, for any set $X \subseteq W$, $h(X)$ is the subset of X where the impermissible outcomes at X are false.¹¹ An information state, on the other hand, is a sphere of epistemic accessibility, that is, a set of worlds that for all that is known could be the actual world:¹²

Definition 6 (Yalcinian Hyperplans) A “Yalcinian” hyperplan is a total function h from non-empty sets of worlds to non-empty sets of worlds such that, for every set of worlds u, u' such that $u' \subseteq u$, $h(u) = u'$ just in case all the impermissible outcomes at u are false at u' .

Definition 7 (Epistemic state) An epistemic state is a function e from possible worlds to sets of possible worlds such that, for every world w and set of worlds u , $e(w) = u$ just in case u is the set of worlds that are epistemically accessible from w .

A semantic model M is a tuple containing a structure S and a valuation function V mapping the descriptive sentences onto S , more specifically, onto sets of triplets $\langle w, h, e \rangle$ of world, hyperplan and information state.

Definition 8 (Model) $M = \langle S, V \rangle$

Subsequently, we define an index of evaluation i as a triplet of elements from the structure S , that is, a world-hyperplan-information state triplet:

Definition 9 (Yalcinian Index) $\langle w_i, h_i, e_i \rangle$

Importantly, e_i should be some set $e (\in E)$ of epistemically accessible worlds, defined relative to the world of the index. That is, $e_i = e(w_i)$.

¹⁰ Making this fully explicit would require incorporating time-sensitivity into the picture as well, but we leave it aside to avoid clutter.

¹¹ Yalcin writes: ‘A hyperplan is a function that takes a set of possible worlds (a set which reflects a possible informational situation *vis-à-vis* what the world is like; a set reflecting ‘an occasion of choice’) to some non-empty subset of that set (a set reflecting outcomes which it is permissible to realize according to the hyperplan, given the informational situation)’ (Yalcin 2012, p. 147). The idea is that a hyperplan takes you from a set of options to a smaller set where only the permissible options remain, or alternatively, the impermissible options have been discarded. Yalcin’s definition is not reproduced *verbatim* so as to keep definitions as uniform as possible.

¹² It is standard to posit this type of parameter in the literature on epistemic and deontic modality. The motivation in the case of deontic modality comes from cases in which the body of information available to a subject can change the truth value of a deontic statement. A well-known example is the Miners’ puzzle (Kolodny and Macfarlane 2010; see also Cariani 2013; Kratzer 2012; Yalcin 2012, *a. m. o.*).

All sentences are assigned uniform semantic values—sets of $\langle w, h, e \rangle$ triplets. But sentences can be descriptive or normative, and the difference consists in whether they impose a condition on the world or on the hyperplan and information parameters of the index: descriptive sentences are true at the world of the index but impose no condition on the hyperplan nor the information parameter; normative sentences do the opposite: they have to be true relative to the hyperplan of the index given a certain state of information, but leave the world-parameter untouched.

First, consider descriptive sentences such as (3.10).

(3.10) Camila packs.

The semantic value of a sentence like (3.10), assigned by V , is a set of world-hyperplan-information triplets. (3.10) is true at an index of evaluation $\langle w_i, h_i, e_i \rangle$ just in case $\langle w_i, h_i, e_i \rangle$ belongs in its denotation (we omit reference to the model M in what follows):

$$(3.11) \quad \llbracket \text{Camila packs} \rrbracket_{\langle w_i, h_i, e_i \rangle}^M = 1 \text{ iff } \langle w_i, h_i, e_i \rangle \in V(\text{Camila packs})$$

We said that descriptive sentences like (3.10) impose a condition on the world but not on the hyperplan nor on the information state parameter of the index. To capture that, all we need is for V to uniformly assign semantic values to atomic, descriptive sentences that result from pairwise combining each of a subset of worlds (those in which Camila packs) with every possible hyperplan and every possible information state. In other words, (3.10) is world-sensitive, but it is hyperplan- and information state-*insensitive*. This is the right result, if we want to represent the fact that sentences like (3.10) only offer factual information, but nothing normative nor epistemic.

Now let us turn to normative sentences. In this language, normative sentences are formed by applying the operator *ought* to descriptive sentences. Where φ is a descriptive sentence, *ought* has the following meaning (Yalcin 2012, p. 148) (remember that Boolean connectives are set aside):

$$(3.12) \quad \llbracket \text{ought } \varphi \rrbracket_{\langle w_i, h_i, e_i \rangle} = 1 \text{ iff } \forall w' \in h_i(e_i) \text{ is such that } \llbracket \varphi \rrbracket_{\langle w', h_i, e_i \rangle} = 1$$

Ought φ is true, relative to a world-hyperplan-information state triplet, just in case every world in the set of worlds that results from applying the information state to the hyperplan of the index is a world that makes φ true.

When applied to (3.10), the result is *Camila ought to pack*, which has the following truth-conditions:

$$(3.13) \quad \llbracket \text{Camila ought to pack} \rrbracket_{\langle w_i, h_i, e_i \rangle} = 1 \text{ iff } \forall w' \in h_i(e_i) \text{ is such that } \llbracket \text{Camila packs} \rrbracket_{\langle w', h_i, e_i \rangle} = 1$$

Just like descriptive sentences, normative sentences denote a set of world, hyperplan and information state triplets. But in contrast to descriptive sentences, normative sentences are hyperplan- and information-sensitive, although world-insensitive. Note that the sentence imposes a condition both on the hyperplan and information state parameters of the index, but not on its world parameter. In order for *Camila ought to pack* to be true at an index i , the hyperplan of i has to be such that, relative to the information state at i , it is required that Camila packs.¹³

¹³ Note that, in this view, deontic force would be captured *via* quantification over the worlds in the domain of a

Two problems for Yalcin's proposal

Could this proposal be applied to evaluative sentences (unembedded sentences containing evaluative adjectives)? Since evaluative language is a species of normative language, we could try to understand evaluatives as imposing certain conditions on the plans of interlocutors, relative to their epistemic state. In a slogan, the idea would be to cash out the linguistic expression of value in terms of plans (and model the content of plans using hyperplans).

As it stands however, Yalcin's proposal cannot be straightforwardly applied to evaluatives. First, because it predicts that everything that is taken to be true across an epistemic state is normatively required by any hyperplan given that epistemic state. And secondly, Yalcin's semantics is not scalar, which means that it needs to be refined in order to represent the meaning of evaluative adjectives, which are gradable. Let us review these issues in turn.

The first problem is one version of a general problem for Kratzer semantics for modals, which Yalcin's proposal inherits, and was raised by Frank (1997) and Zvolenszky (2002). Recall that a hyperplan is a function that, relative to an epistemic state e , will return a subset of e where all the permissible options at e are true. A proposition p ought to be the case relative to a hyperplan h and epistemic state e just in case $h(e)$ entails p ; that is, just in case all worlds in $h(e)$ are worlds where p is true. Conversely, if a proposition is entailed by a $h(e)$, then it ought to be the case. The problem is, then, that any proposition that is true across e will also be true at $h(e)$. In other words, everything that is known is *eo ipso* normatively required.

To see how this is a problem, consider the following situation: Camila needs to pack a bag and go catch her train. All her evidence suggests that it is late. This means that her epistemic state, call it e , is such that the proposition that it is late is true at every world w in e . But if the proposition that it is late is true at every world in e , then it is true at all the worlds in the set that results from applying any hyperplan h to her epistemic state. Therefore, it ought to be the case that it is late.

Importantly, this is not a fatal blow to this proposal. For example, one could say that the result of applying a hyperplan to a set of worlds relative to which a proposition p is true is not a subset of those worlds, but rather a different set of worlds altogether, maybe one where $\text{not-}p$ is an open possibility (Frank 1997). Or one could say that Yalcin's truth conditions for *ought* φ actually state only a necessary condition for the truth of *ought* φ . However, rather than look for ways of fixing Yalcin's view, I think we had better replace certain key elements altogether. We will see why in the following subsection.

The second issue with Yalcin's (and Gibbard's) view is that it is not scalar: propositions are either required or not required (they are or are not *the thing to do*, they do or do not *make sense*, etc.). In terms of the hyperplan semantics just given, this means that the speaker either plans or does not plan to make the proposition true (under certain circumstances). However, as we saw in Chapter 2, in order to account for the scalar properties of evaluative adjectives, we need a more sophisticated story than a simple non-factualist story in which evaluative sentences express outright practical attitudes (in this case, an attitude of planning to make a proposition true). For an illustration, consider the following sentences:

(3.14) Volunteering is better than donating.

hyperplan, relative to an epistemic state. That is: p ought to be the case relative to a hyperplan h and information state e just in case every world in $h(e)$ makes p true; p may be the case relative to h, e just in case some world in $h(e)$ makes p true; etc. See Yalcin 2012, p. 148.

- (3.15) It is good that you sent that e-mail.
 (3.16) It is extremely good that you sent that e-mail.

First, in order to account for (3.14) we need to be able to compare the value that a speaker assigns to each *relata*. In a Gibbard-Yalcin framework, we would need to say something about the plan that the speaker of (3.14) is advocating for. But the story given so far only lets us say something to the effect that the speaker plans to volunteer; or plans to donate; or plans to volunteer and donate, etc. None of these options apply to (3.14). Similarly, how could we account for the difference between (3.15) and (3.16)? Suppose that we assign to (3.15) a roughly equivalent meaning as *ought p*. What meaning could we give (3.16) then? There is no straightforward way of strengthening *ought p* into something that captures the addition of the modifier in (3.16).

Recall however, that Yalcin is taking resources from Kratzer semantics for modals. And even though Yalcin’s proposal is not scalar, Kratzer’s is. So one could think that, in order to make Yalcin’s proposal scalar, one simply needs to incorporate more Kratzerian (or Lewisian) elements (Kratzer 2012; D. Lewis 1973). An obvious strategy would be to bring ORDERING SOURCES into the picture. Here is how it might work: rather than have a hyperplan filter out the non-permissible worlds of an epistemic state, we could let an ordering source order—rather than rule out—those same worlds in terms of their closeness to an ideal. Then, we could cash out the contrast between (3.15) and (3.16) by saying that a proposition is *good* just in case it satisfies the relevant ideal to some contextually specified extent; and we could say that a proposition is *extremely good* just in case it satisfies the same ideal to some “extreme” extent.¹⁴

This is a possible strategy, but not the one that I will pursue, for two reasons. First, ordering sources do not immediately help with our first problem. And secondly, it is not easy to capture the scalar properties of relative-standard gradable adjectives in Kratzer’s semantics, a problem pressed by Lassiter (2017). Nonetheless, remember that Yalcin was only trying to show that Gibbard’s hyperplans were compatible with a standard story about deontics, and the standard story about deontics is basically Kratzer’s. Nothing commits *us* to the details of Kratzer semantics for deontic modals. So rather than explore further ways of building a scalar system for evaluatives based on Kratzer semantics, we will take a slightly different route.

3.3.3 Hyperplans as functions over alternatives

The shortcomings of Gibbard-Yalcin semantics are clear, but they can be overcome. Put simply: if one wants to understand value in terms of properties of hyperplans, we need to be able to represent a scale of value using hyperplans. This is what we are going to do in this section. To that effect, we propose to change the Gibbard-Yalcin semantics in two ways. First, let us drop the assumption that the arguments of hyperplans are *information* states. This assumption was necessary to account for deontic modals, but it is not clear that evaluatives rely on epistemic information in the same way that deontics do.

Secondly, the way that we are going to characterize hyperplans, the property of an epistemic state (or of any set of possible worlds) that the hyperplan will “see”, so to speak, are not simply

¹⁴A different Kratzerian strategy would be to turn to the distinction between strong and weak necessity modals (Fintel and Iatridou 2005; Von Stechow and Iatridou 2008), which has been applied to the contrast between extreme and non-extreme deontic adjectives (*crucial* vs. *important*, see Portner and Rubinstein 2016). However, that distinction is still too coarse to capture the scalarity of evaluatives.

the possible worlds contained in it, but rather the possible alternatives that are salient at that state. In our characterization, hyperplans are going to be functions from sets of sets of worlds to subsets thereof containing the chosen, or preferable, alternatives at the original set. This second tweak will allow us to model degrees of positive and negative value as degrees of priority: a proposition has a high degree of positive value just in case one’s plan is to give priority to that proposition over certain alternatives.

Besides the technical problem that was pointed out in the previous subsection—that everything that is known is normatively required—my main reason for turning to a different definition of hyperplans is that the Gibbard-Yalcin proposal is too coarse-grained. Intuitively, hyperplans select desirable, permissible or ideal states of affairs among different options. But the worlds that make those states of affairs true might also make a lot of other things true, things that we do not want a hyperplan to “care for”, so to speak.

Recall our previous example: we want Camila’s planning state to be represented by hyperplans that require her to start packing. But given that it is late, the worlds in which she starts packing are also worlds in which it is late. However, we do not want to say that hyperplans that require her to start packing also require it to be late. Being late is not part of what is required; rather, it is part of the circumstances relative to which the hyperplan requires the agent to start packing. More generally, whatever is required by a hyperplan should not include circumstances such as it being late. To capture this, rather than have hyperplans select sets of possible worlds, we can make them select sets of propositions. In the case of Camila, a hyperplan would be a function that, given her situation, selects the singleton set containing the proposition that Camila starts packing, rather than the set of worlds that make that proposition true. By making a hyperplan select sets of propositions, we are not wedded anymore to the view that whatever is already true across the selected worlds comes out required by a hyperplan. In order to be required by a hyperplan, it is not enough for a proposition to be true across worlds that are selected as preferable. Rather, the relevant proposition has to be selected by the hyperplan among a set of possible alternatives.

For further illustration, let us consider the example that we previewed at the end of §3.3.1 (and which will be a running example in this and the following chapter). Suppose that Kanye West has released a new album, and we are about to learn what Grammy nominations he got. Since he is an ambitious, global artist, he only cares for two categories: Song of the Year and Album of the Year; but since he is African American and the Grammys are famously racist, it is not guaranteed that he will be nominated, let alone awarded—however good the new album is.¹⁵

We can think of Kanye’s predicament in terms of his epistemic state, defined by the proposition that he releases a new album. In Yalcin’s view, a hyperplan takes you from an information state, understood as something like what you see on the left column of Table 3.2, to a subset of that state where the impermissible options are ruled out. Alternatively however, we can also think of that very same situation in terms of the set of possible alternatives or developments ahead of him: namely, that his new album gets 0, 1 or 2 nominations (out of the categories that he cares about). This is represented on the right column of Table 3.2:

Consider again information state e : that state is characterized by a set of worlds, namely worlds at which Kanye releases a new album. We will say that a hyperplan, when fed that set of

¹⁵ ‘ “[My Beautiful] Dark [Twisted] Fantasy” and “Watch the Throne”: neither was nominated for Album of the Year, and I made both of those in one year. I don’t know if this is statistically right, but I’m assuming I have the most Grammys of anyone my age, but I haven’t won one against a white person.’ (K. West, NY Times, 2013).

Epistemic state e	Set U of alternatives at e
Kanye releases a new album	Kanye's new album receives 0 nomination
	Kanye's new album receives 1 nomination
	Kanye's new album receives 2 nominations

Table 3.2: Kanye's epistemic state e as a set of worlds and as a set of alternatives.

possible worlds, distinguishes the different alternatives at that information state and selects one or more. That is, hyperplans see the right, rather than the left column of Table 3.2. In order to model such fine-grainedness, we will say that hyperplans do not take sets of possible worlds as arguments and return subsets of those sets as their value. Rather, we will say that their arguments are sets of propositions that represent the different alternatives that are open in a given situation. Their values will be subsets of those sets of alternatives, which represent the preferable or chosen alternatives relative to the former set of alternatives.

Hyperplans are defined as follows (an illustration is given in Table 3.3):¹⁶

Definition 10 (Hyperplans) *A hyperplan is a total function h from non-empty sets of propositions to non-empty sets of propositions such that, for every non-empty set of propositions U, V such that $V \subseteq U$, $h(U) = V$ just in case V are the chosen alternatives at U .*

Set U of alternatives		Set U of alternatives
Kanye's new album receives 0 nomination	\mapsto	Kanye's new album receives 0 nomination
Kanye's new album receives 1 nomination		Kanye's new album receives 1 nomination
Kanye's new album receives 2 nominations		Kanye's new album receives 2 nominations

Table 3.3: A hyperplan as a function from sets of alternatives to sets of alternatives.

Note that we are not too far from Gibbard's characterization. In our view, a hyperplan is a function from sets of alternatives to subsets thereof. But we can formulate this as a set of conditional imperatives, with an alternative set in the antecedent and the preferred options in the consequent. For instance, the conditional imperative version of the hyperplan characterized by Table 3.3 would be:

(3.17) If given a choice between 0, 1 and 2 nominations, let it be 2 nominations!

¹⁶ The notion of alternatives comes mainly from question and focus semantics, but the details of that literature are not important for our present concerns (see e.g. Aloni 2007, 2018; Alonso-Ovalle 2006; Hamblin 1973; Karttunen 2008; Rooth 1992 for classic and recent references). All we need to do is keep in mind that a set of alternatives is a set of propositions, which in turn are sets of possible worlds. This notion should also not be confused with the notion of (relevant) alternatives often used in (meta-)epistemology (D. Lewis 1996; Stine 1976). Similarly, the expression *alternative* should not make the reader think that hyperplans can select only one. We might as well speak of sets of *options* (where it seems more intuitive to say that one can choose more than one) or *outcomes*—although the latter seems to suggest that these are outcomes of something, which need not be the case (thanks to Carla Umbach for suggesting this).

Closer to our present interests, a similar notion of alternatives to the one employed here is used by Starr to model imperatives (Starr 2018). It bears mentioning that hyperplans, in this characterization, are formally very similar to choice functions from economics, which select the most preferable among a set of given options (see Hansson and Grüne-Yanoff 2018 for an introduction, esp. §5 and ff.). Indeed, Silk (2015) has defended a preference semantics for expressivists which recruits and deploys such concepts from economics. His proposal is scalar and has much in common with the present proposal, although it is not aimed specifically at capturing the scalarity of particular natural language items.

Note, however, that we are describing a model, but we haven't said yet what sentences or formulae this model could represent. In the next section of this chapter, we will use this model to describe a formal language, and in Chapter 4 we will use a slightly more sophisticated version of this model to represent the meaning of evaluative sentences of natural language. At this stage, we are simply considering some formal properties of a possible plan of action, one that has as its outcome that Kanye obtains two out of two Grammy nominations. Eventually, this plan will figure in the semantics of sentences that evaluate positively that situation, that is, sentences like *it'd be great if Kanye got 2/2 nominations* or *it's good that he got 2/2 nominations*. Those sentences will be true relative to a point of evaluation that contains a plan according to which, when faced with a choice between 0, 1 and 2 nominations, 2 nominations are preferred. But recall that, even if it is Kanye who is evaluating those sentences, the plan according to which he gets 2/2 nominations will not be such that Kanye *himself* chooses for him to obtain 2/2 nominations. This would be obviously bad for reasons given in §3.3.1: first, Kanye is not the one making the choice. And secondly, no single person decides the outcome of the Grammy nominations. As we said, we need to understand these sentences as endorsing a plan for an agent to choose whatever action leads to that outcome (in the case of a vote, this should most reasonably be understood as a plan for each of the people taking part in the vote to vote appropriately).

The general picture will be that evaluative sentences represent properties of hyperplans; and to make an evaluative judgment is, as we will see in later chapters, to propose to adopt hyperplans that share a certain property. More specifically, *positive* evaluative adjectives will be characterized by hyperplans that prefer, or choose, the things under evaluation over certain alternatives; while *negative* evaluative adjectives will be characterized by hyperplans that dis-prefer, or avoid, the object of evaluation. That is, in a nutshell, our approach.

It is important to stress another parallelism with possible worlds and propositions that was already mentioned: just as possible worlds might be thought of as maximally opinionated states of belief, hyperplans can be thought of as maximally specific states of planning. In our terms, this means that hyperplans are defined for absolutely any set of possible alternatives (as said in the definition, hyperplans are total functions). That is, for any set of propositions U , a hyperplan returns a subset of U consisting of the chosen alternatives at U . But just as actual states of belief are never maximally opinionated, real-life plans (not hyperplans) are never defined for absolutely any contingency. Therefore, we can define PLANS as *partial* functions from sets of propositions to preferred alternatives; or correspondingly as sets of hyperplans that agree on some of their instructions but not on others.

Hyperplans can avoid or dis-prefer alternatives as well, which will be crucial to the semantics of evaluative adjectives of negative polarity, such as *bad*, *cruel* or *ugly*. The easiest (if not the most elegant) way to model negative preference is to assign a negative function to each hyperplan h , call it \mathcal{U} , that selects the alternative(s) that h avoids or dis-prefers out of every set of alternatives:

Definition 11 (Negative hyperplans) *A negative hyperplan is a total function \mathcal{U} from hyperplans and non-empty sets of propositions to non-empty sets of propositions such that, for every hyperplan h and sets of propositions U, V such that $V \subseteq U \setminus h(U)$, $\mathcal{U}\langle h, U \rangle = V$ just in case V are the alternatives at U that h dis-prefers.*

The function \mathcal{U} behaves like a hyperplan in that it is a total function on sets of alternatives, and it can select only some or all alternatives as undesirable. In addition, note that, since the range

of a negative hyperplan relative to a hyperplan h and a set of alternatives X is a subset of X minus the result of applying h to X , the positive and negative choices of a hyperplan never overlap.

To sum up: in this section, we introduced Gibbard’s notion of hyperplan as a view about the content of normative judgments. In Gibbard’s view, hyperplans are semantic primitives, just like possible worlds. Then, we saw *via* Yalcin’s 2012 proposal how to implement Gibbard’s ideas in a compositional, intensional semantics for deontic *ought* that, in addition, no longer treats hyperplans as primitives, but rather as functions over sets of possible worlds. However, we argued that, in order to apply a hyperplan semantics to evaluative sentences, we needed to move beyond Yalcin’s characterization of hyperplans. By conceiving of hyperplans as functions that select preferred and dis-preferred sets of alternatives, we have seen how to overcome the problems that Yalcin’s implementation had, while preserving much of the spirit of that proposal. In the next section, we make all of the preceding observations more formal by defining a simple propositional language that introduces scalar evaluative operators whose semantics make use of hyperplans.

3.4 Evaluating alternatives: the language of \mathcal{L}_V

In this section, we introduce a simple propositional language that contains a pair of positive and negative unary evaluative operators \Uparrow and \Downarrow , as well as three binary evaluative operators: \geq , $>$ and \approx ; the semantics of which crucially operate on a hyperplan parameter in the index of evaluation. As we will see in Chapter 4, this language is meant to capture a basic semantic property of sentences in which positive and negative evaluative adjectives are applied to sentences. To be more specific, the unary connectives \Uparrow and \Downarrow are meant to represent the positive form of evaluative adjectives of positive and negative polarity respectively; \geq represents equatives (*at least as good/bad as*), $>$ represents comparatives and \approx is an relation of equal value (*exactly as good/bad as*), which we can use to define a degree system. For example, our claim will be that an English sentence like *it’s good that Bill arrived* has, to a large extent, the same semantics as the \mathcal{L}_V formula $\Uparrow \varphi$ (where φ stands for the English clause *that Bill arrived*).

3.4.1 Syntax

In addition to propositional constants and Boolean connectives, \mathcal{L}_V introduces five *sui generis* sentential operators: two unary operators \Uparrow and \Downarrow and three binary operators \geq , $>$ and \approx .

The vocabulary of \mathcal{L}_V contains:

- Sentential constants: p, q, r, \dots
- Boolean connectives: $\wedge, \vee, \rightarrow, \sim$
- Operators: $\Uparrow, \Downarrow, \geq, >, \approx$

The syntax of \mathcal{L}_V works as follows:

- Sentential constants are *sentences*.
- If φ and ψ are sentences, then

$$\sim \varphi \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \varphi \rightarrow \psi \mid \uparrow \varphi \mid \downarrow \varphi \mid \varphi \geq \psi \mid \varphi > \psi \mid \varphi \approx \psi$$

are sentences.

- Nothing else is a sentence.

3.4.2 Semantics

To give a semantics for \mathcal{L}_V , first we define a semantic structure, a semantic model and a notion of truth-at-an-index of evaluation. Then, we give compositional semantic rules for the sentences of \mathcal{L}_V .

First, a structure S is a tuple formed by a set W of possible worlds and a set of hyperplans H .

Definition 12 (Semantic structure for \mathcal{L}_V) $S = \langle W, H \rangle$

These parameters play a role in evaluating the truth of sentences of \mathcal{L}_V at an index of evaluation. All sentences of \mathcal{L}_V have uniform semantic values, but whereas sentential constants will impose a certain condition on the set of possible worlds W , formulae containing the operators of \mathcal{L}_V will impose conditions on the set of hyperplans H . H is a set of hyperplans, which are functions from sets of propositions to sets of propositions. Propositions are sets of possible worlds, and since the domain and range of any hyperplan in H are sets of propositions, the domain and range of any hyperplan are subsets of the power set of W , $\mathcal{P}(W)$, that is, the set of all propositions.

In §3.3.3 we defined hyperplans as follows:

Definition 10 (Hyperplans) A hyperplan is a total function h from non-empty sets of propositions to non-empty sets of propositions such that, for every non-empty set of propositions U , V such that $V \subseteq U$, $h(U) = V$ just in case V are the chosen alternatives at U .

And as an auxiliary notion, we have a negative preference function on hyperplans, \mathcal{U} . For any set of propositions U and any hyperplan h , \mathcal{U} selects the propositions among U that are dis-preferred. Recall our previous definition:

Definition 11 (Negative hyperplans) A negative hyperplan is a total function \mathcal{U} from hyperplans and non-empty sets of propositions to non-empty sets of propositions such that, for every hyperplan h and sets of propositions U, V such that $V \subseteq U \setminus h(U)$, $\mathcal{U}\langle h, U \rangle = V$ just in case V are the alternatives at U that h dis-prefers.

A model is a tuple M consisting of a structure S and a valuation function V mapping every sentential constant of \mathcal{L}_V to a subset of $W \times H$, that is, to a set of world-hyperplan pairs.

Definition 13 (Model for \mathcal{L}_V) $M = \langle S, V \rangle$

Sentences of \mathcal{L}_V are evaluated for truth at indices of evaluation, which are triplets formed by a possible world, a hyperplan and the set of all propositions $\mathcal{P}(W)$. $\mathcal{P}(W)$ is held constant across all sentences of \mathcal{L}_V , but it will be useful to have it in the index in order to show how the binary operators of \mathcal{L}_V work.

Definition 14 (Index for \mathcal{L}_V) $\langle w_i, h_i, \mathcal{P}(W) \rangle$

That is, the index provides a world a hyperplan and the set of all propositions (we will often use i as variable over indices).

Sentences are evaluated for truth at indices of evaluation. As we saw in §3.3.2, we want the sentential constants of \mathcal{L}_V to be descriptive. Ignoring $\mathcal{P}(W)$, which is held constant, this means that sentential constants should impose a condition on the world but not on the hyperplan parameter of the index. And to do that, we need the valuation function V to assign to the sentential constants sets of $w, h, \mathcal{P}(W)$ triplets formed by pairwise combining each of a subset of worlds with every possible hyperplan and $\mathcal{P}(W)$. That way, sentential constants will come out world-, but not hyperplan-sensitive.¹⁷

$$(3.18) \quad \llbracket p \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle}^M = 1 \text{ iff } \langle w_i, h_i, \mathcal{P}(W) \rangle \in V(p) \text{ (we skip superscript } M \text{ in what follows)}$$

The semantic rules of composition for the Boolean connectives are the following:

$$(3.19) \quad \llbracket \sim \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 0$$

$$(3.20) \quad \llbracket \varphi \wedge \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ and } \llbracket \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1$$

$$(3.21) \quad \llbracket \varphi \vee \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ or } \llbracket \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1$$

$$(3.22) \quad \llbracket \varphi \rightarrow \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 0 \text{ or } \llbracket \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1$$

Let us turn now to the evaluative operators introduced by \mathcal{L}_V .

The positive form: \Uparrow, \Downarrow

The unary operators \Uparrow, \Downarrow have the following semantic entries:

$$(3.23) \quad \llbracket \Uparrow \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(\mathcal{P}(W))$$

In words: ‘ $\Uparrow \varphi$ ’ is true at an index i just in case the set of possible worlds in the worldly component of φ ’s denotation is among the alternatives chosen by the hyperplan of the index, relative to the set of all propositions.

$$(3.24) \quad \llbracket \Downarrow \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in \mathcal{U}(h_i, \mathcal{P}(W))$$

In words: ‘ $\Downarrow \varphi$ ’ is true at an index i just in case the hyperplan of i is such that the set of worlds in the worldly component of φ is among the dis-preferred outcomes.

Let us highlight three aspects that can be potentially confusing at this stage: first, note that the only relationship between \Uparrow and \Downarrow that falls out from the semantics so far is that their extensions are non-intersective (this is a consequence of the definition of negative hyperplans given in §3.3.3). This establishes that if $\Uparrow \varphi$ is true, then $\Downarrow \varphi$ is false; and if $\Downarrow \varphi$ is true, then $\Uparrow \varphi$ is false. But that does not hold in the opposite direction: if $\Uparrow \varphi$ is false, it does not follow that $\Downarrow \varphi$ is true, and *viceversa*. This is how it should be, since \Uparrow and \Downarrow are not duals. This will be clearer in Chapter 4 when we map \Uparrow onto *good* and \Downarrow onto *bad*, as *good* and *bad* are not duals: all that is not good isn’t bad, and all that is not bad isn’t good. As we will see shortly though, \Uparrow and \Downarrow are further related *via* the binary operators of \mathcal{L}_V .

¹⁷Remember that letting V assign sets of possible worlds (rather than sets of $w, h, \mathcal{P}(W)$ triplets) to sentential constants while reserving the more “complex” denotations for evaluative sentences gets us into trouble: if we did that, we would not be assigning uniform semantic values to all formulae. In turn, this would make it problematic to account for Boolean combinations of descriptive and non-descriptive sentences. See §3.3.2.

Secondly, recall that we uniformly assign to any sentence of \mathcal{L}_V a set of world-hyperplan- $\mathcal{P}(W)$ triplets. This means that we cannot simply say that an index i makes ‘ $\uparrow \varphi$ ’ true if $\llbracket \varphi \rrbracket$ is one of the chosen alternatives at h_i ... because $\llbracket \varphi \rrbracket$ is not (only) a set of worlds—it is a set of world-hyperplan- $\mathcal{P}(W)$ triplets. In order to make a sentence like φ the argument of an operator like \uparrow , we need to abstract away from its hyperplan and alternative components and retain its “wordly” component, the set of worlds in the denotation of φ . To do that, we take $\{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\}$, that is, the “worldly” component of φ , and we say that that set of worlds is among the chosen alternatives. So for instance, if φ is the sentence *Mora runs*, then $\llbracket \text{Mora runs} \rrbracket$ is the set of world w , hyperplan h and $\mathcal{P}(W)$ triplets such that Mora runs at w . Now if *Mora runs* is embedded under \uparrow , we get the formula ‘ $\uparrow \text{Mora runs}$ ’. But, according to our semantics, the semantic argument of \uparrow is not a set of world-hyperplan- $\mathcal{P}(W)$ triplets, but only the worldly component of that set—the worlds at which Mora runs. And then, ‘ $\uparrow \text{Mora runs}$ ’ is true at an index of evaluation i just in case the set formed by those worlds belongs in the set of preferred alternatives of the hyperplan of i relative to the set of alternatives $\mathcal{P}(W)$.

Finally, note that, in the truth conditions for \uparrow and \downarrow , the hyperplan of the index is invariably fed the set of all sets of worlds, $\mathcal{P}(W)$, which is also a constant parameter in the index of evaluation. In practice, this means that the hyperplan of the index does not behave like a function—because it receives a constant argument. This is a formal oddity, but it is an idealization that for now will make our lives simpler. When we move from \mathcal{L}_V to natural language evaluatives, we will see how hyperplans can and do take as arguments varying sets of alternatives that are also supplied by the index of evaluation (analogously to the way in which, in Yalcin’s proposal, hyperplans were fed epistemic states).

Equatives & comparatives: $\geq, >$

The binary operators are defined with the aid of the semantics of either \uparrow or \downarrow . Let’s start with \geq :

$$(3.25) \quad \llbracket \varphi \geq \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) \text{ such that } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X \ \& \ \{w' : \llbracket \psi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X, \text{ if } \llbracket \uparrow \psi \rrbracket_{\langle w_i, h_i, X \rangle} = 1, \text{ then } \llbracket \uparrow \varphi \rrbracket_{\langle w_i, h_i, X \rangle} = 1.$$

In words, a proposition φ is at least as valuable as a proposition ψ just in case, relative to an index i , every subset of alternatives containing φ and ψ among which h_i chooses ψ is one among which φ is also chosen. In other words, there is no set of alternatives containing φ and ψ relative to which h_i chooses ψ but not φ .

Alternatively, we can define \geq with \downarrow :

$$(3.26) \quad \llbracket \varphi \geq \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) \text{ such that } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X \ \& \ \{w' : \llbracket \psi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X, \text{ if } \llbracket \downarrow \varphi \rrbracket_{\langle w_i, h_i, X \rangle} = 1, \text{ then } \llbracket \downarrow \psi \rrbracket_{\langle w_i, h_i, X \rangle} = 1.$$

In words, a proposition φ is at least as valuable as a proposition ψ just in case, relative to an index i , every subset of alternatives containing φ and ψ among which h_i disprefers φ is one among which ψ is also dis-preferred. In other words, there is no set of alternatives containing φ and ψ relative to which h_i dis-prefers φ but not ψ .

From this we can easily define the comparative connective $>$:

$$(3.27) \quad \llbracket \varphi > \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \geq \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ and } \llbracket \psi \geq \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 0$$

In words: relative to an index i , a proposition φ is more valuable than a proposition ψ just in case φ is at least as valuable as ψ but not the other way around. Again, this can be defined with \Uparrow or \Downarrow . That is, one cannot find a set of alternatives (a subset of $\mathcal{P}(W)$ containing φ and ψ) at which ψ \Uparrow φ but not φ \Uparrow ψ is chosen [dis-preferred], although one can find a set of alternatives at which φ \Uparrow ψ but not ψ \Uparrow φ is chosen [dis-preferred]. In other words, the hyperplan of the index is such that there is at least one subset of $\mathcal{P}(W)$ containing φ and ψ at which φ \Uparrow ψ is chosen [dis-preferred] over ψ \Uparrow φ :

$$(3.28) \quad \llbracket \varphi > \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \exists X \subseteq \mathcal{P}(W) \text{ such that } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X \ \& \ \{w' : \llbracket \psi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X, \text{ and } \llbracket \Uparrow \varphi \rrbracket_{\langle w_i, h_i, X \rangle} = 1 \text{ and } \llbracket \Uparrow \psi \rrbracket_{\langle w_i, h_i, X \rangle} = 0.$$

$$(3.29) \quad \llbracket \varphi > \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \exists X \subseteq \mathcal{P}(W) \text{ such that } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X \ \& \ \{w' : \llbracket \psi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in X, \text{ and } \llbracket \Downarrow \psi \rrbracket_{\langle w_i, h_i, X \rangle} = 1 \text{ and } \llbracket \Downarrow \varphi \rrbracket_{\langle w_i, h_i, X \rangle} = 0.$$

Defining the binary operators (\geq , $>$) of \mathcal{L}_V using either \Uparrow or \Downarrow has the consequence that a relation is established between \Uparrow and \Downarrow . Simply put, that relation is that the orderings induced by recursive applications of either \Uparrow or \Downarrow mirror each other. If we use \Uparrow to order two propositions, say φ and ψ , it follows that an ordering has been established involving \Downarrow (and viceversa).¹⁸ This can be seen more clearly if we suppose that, where \Uparrow is mapped onto *good* and \Downarrow is mapped onto *bad*, $>$ is mapped onto *better* if read left to right and *worse* if read right to left. We will discuss this further in the next chapter.

By purposefully holding fixed the hyperplan's argument in the semantic entry for the unary operators, we can see the difference between the unary and the binary operators of \mathcal{L}_V : at an index of evaluation i , the unary operators \Uparrow and \Downarrow impose certain conditions on the hyperplan of i relative to a fixed set of alternatives, namely $\mathcal{P}(W)$. By contrast, relative to the same index of evaluation, when evaluating formulae containing the binary operators, the hyperplan of i is not fed $\mathcal{P}(W)$. Rather, the hyperplan is fed subsets of $\mathcal{P}(W)$; and the various binary relations are defined by quantifying over those subsets.

The equal-value connective \approx & defining a degree system

The equative connective \geq also lets us define an equal-value relation \approx as one would expect:

$$(3.30) \quad \llbracket \varphi \approx \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \llbracket \varphi \geq \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ and } \llbracket \psi \geq \varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1$$

Additionally, using \approx we can define a degree system by identifying a proposition φ 's degree of value, written $\mu_V(\varphi)$, with the set of propositions that stand in the \approx relation to φ relative to an index i .

$$(3.31) \quad \llbracket \mu_V(\varphi) \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = \{\psi : \llbracket \varphi \approx \psi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1\}$$

Note, however, that μ_V is not part of the vocabulary of \mathcal{L}_V . The sole purpose of defining μ_V is to show that we can conceive of degrees as equivalence classes of propositions (Cresswell 1976; Klein 1980, *a.o.*), and use the formal apparatus of degree semantics (that is, measure

¹⁸(Informal) proof: if $\Uparrow \varphi$ and $\sim \Uparrow \psi$, then $\varphi > \psi$. By the inter-definability of $>$ via \Uparrow or \Downarrow , it follows that there is a subset X of alternatives of $\mathcal{P}(W)$ including φ and ψ such that $\Downarrow \psi$ and $\sim \Downarrow \varphi$ relative to X .

functions like μ) to represent the meaning of certain constructions. This will be convenient in §4.6.

To proceed in this way is to use a Delineation approach to the semantics of the operators of \mathcal{L}_V : we start by giving a non-scalar semantics for the unary operators \uparrow and \downarrow , and we then give a scalar semantics for the binary operators \geq , $>$ and \approx by quantifying over subsets of the set of alternatives relative to which the unary operators are evaluated, namely $\mathcal{P}(W)$. The procedure is similar to the one we presented at the end of Chapter 2 for *tall*: the positive form *tall* is assigned a non-scalar extension relative to a comparison class. Any individual in that comparison class will fall in the extension, anti-extension or extension gap of *tall*. And in order to define the comparative *taller*, we introduce quantification over comparison classes: an individual x is taller than an individual y just in case there exists a comparison class relative to which x is tall but y is not (introducing quantification over comparison classes, in turn, makes comparatives appropriately *insensitive* to the comparison class supplied by context).

Iterated \mathcal{L}_V operators?

Before moving on, let us observe that it is not straightforward to say what, if anything, the meaning of formulae containing iterated instances of the operators of \mathcal{L}_V should be. What is the meaning of ‘ $\uparrow\uparrow\varphi$ ’, ‘ $\downarrow\sim\uparrow\varphi$ ’ or ‘ $\uparrow\varphi > \psi$ ’?

At the outset, let me point out that it is not clear how to interpret the corresponding English sentences. Assume, as we will, that \uparrow represents the positive form of an evaluative adjective, e.g. *good*, and that $>$ represents the comparative, e.g. *better*. Then, a formula like ‘ $\uparrow\uparrow\varphi$ ’ would be equivalent to the English sentence *it is good that it is good that φ* . And a formula like ‘ $\uparrow\varphi > \psi$ ’ would be equivalent to *It is good that it is better that φ than that ψ* . These are sentences whose meaning is intuitively difficult to spell out.

Regardless, it is important to say what our semantics predicts the corresponding formulae to denote. In general, the meaning of any evaluative formula that is under the scope of an evaluative operator is equivalent to the “total” proposition, i.e. W , for the following reason: evaluative operators impose a condition on, and only on, the worldly component of the formulae that they take as arguments. But the worldly component of any evaluative formula is the set of all worlds. Therefore, any evaluative operator that takes an evaluative formula as its argument will only “see” the worldly component of that formula, namely the set of all possible worlds W . Consider this example:

$$(3.32) \quad \llbracket \uparrow\uparrow\varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \{w' : \llbracket \uparrow\varphi \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(\mathcal{P}(W))$$

Since ‘ $\uparrow\varphi$ ’ imposes no condition on the world parameter, the set of worlds at which $\uparrow\varphi$ is true is the set of all worlds. Thus,

$$(3.33) \quad \llbracket \uparrow\uparrow\varphi \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } W \in h_i(\mathcal{P}(W))$$

Admittedly, this is a little strange: ‘ $\uparrow\uparrow\varphi$ ’ is true just in case W is among the chosen set of alternatives relative to the hyperplan of the index. However, the set of all worlds makes *no* proposition false, and therefore, it does not seem like a valid choice for a hyperplan.

Furthermore, consider Boolean combinations including evaluative formulae under the scope of other evaluative operators, such as ‘ $\uparrow\sim\downarrow\varphi$ ’. If our semantics predicts that evaluative formulae in the scope of evaluative operators contribute the total proposition, and since the embedded formula is negated, the meaning of ‘ $\sim\downarrow\varphi$ ’ under \uparrow turns out to be the empty set.

$$(3.34) \quad \begin{aligned} & [[\uparrow\sim\downarrow\varphi]]_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff} \\ & \{w' : [[\sim\downarrow\varphi]]_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(\mathcal{P}(W)) = 1 \text{ iff} \\ & \{w' : [[\downarrow\varphi]]_{\langle w', h_i, \mathcal{P}(W) \rangle} = 0\} \in h_i(\mathcal{P}(W)) = 1 \text{ iff} \\ & \emptyset \in h_i(\mathcal{P}(W)) \end{aligned}$$

And given that the empty set is included in every set whatsoever, it is included in the set of preferred alternatives, thereby making ‘ $\uparrow\sim\downarrow\varphi$ ’ trivially true at any index of evaluation.

There are solutions that one might try to apply to make these results less awkward. Syntactically, one could ban evaluative operators from embedding evaluative formulae and Boolean combinations thereof. Semantically, it might make sense to preclude the total proposition W from being a possible value of $h(\mathcal{P}(W))$, for any h . That would make evaluative formulae embedded under evaluative operators uninterpretable. Either move could be welcome, in view of the observation that the corresponding natural language sentence are hard to interpret. But given that they are hard to interpret, it is not clear either that the predictions that we obtain from our semantics are actually bad in the first place. We leave this issue for future consideration.

3.5 A toy fragment & structure

To see the operators of \mathcal{L}_V at work, let us map a fragment of \mathcal{L}_V to a small toy structure, using Kanye’s example again. Our fragment has the following sentences:

- Atomic sentences:
 - p = ‘Kanye gets 0 nominations’
 - q = ‘Kanye gets 1 nomination’
 - r = ‘Kanye gets 2 nominations’
- Non-atomic sentences: $\uparrow\varphi \mid \downarrow\varphi \mid \varphi \geq \psi \mid \varphi > \psi \mid \varphi \approx \psi$

And we map sentences onto the following structure:

- S is a structure composed of $\langle W, H \rangle$, where,
 - W is a finite set $\{w_0, \dots, w_n\}$ of possible worlds
 - U is a finite set of propositions $\begin{cases} u_0 = \{w : \text{Kanye gets 0 nominations in } w\} \\ u_1 = \{w : \text{Kanye gets 1 nominations in } w\} \\ u_2 = \{w : \text{Kanye gets 2 nominations in } w\} \end{cases}$
 - H is a finite set $\{h_0, h_1, h_2, h_3, h_4, h_5\}$ of hyperplans

Remember that hyperplans select their preferences and dis-preferences among sets of propositions. Hyperplans are total functions, which means that every hyperplan in H is such that, for each subset of U , it tells you which ones are (dis-)preferable. Table 3.4 shows the “preferential”

profile of each hyperplan in H , relative to each set of propositions whose cardinality is > 1 (we leave out the empty set because hyperplans are not defined for it and we also leave out singleton sets because for any hyperplan h and singleton set of propositions U , $h(U) = U$):

$\{X \subseteq U : X > 1\}$	h_0	h_1	h_2	h_3	h_4	h_5
$U = \{u_0, u_1, u_2\}$	$\{u_1, u_2\}$	$\{u_1\}$	$\{u_0, u_1\}$	$\{u_0\}$	$\{u_0, u_1, u_2\}$	$\{u_2\}$
$\{u_1, u_2\}$	$\{u_2\}$	$\{u_1\}$	$\{u_1\}$	$\{u_2\}$	$\{u_2\}$	$\{u_2\}$
$\{u_0, u_1\}$	$\{u_1\}$	$\{u_1\}$	$\{u_0\}$	$\{u_1\}$	$\{u_0\}$	$\{u_0, u_1\}$
$\{u_0, u_2\}$	$\{u_2\}$	$\{u_0, u_2\}$	$\{u_0\}$	$\{u_0\}$	$\{u_0\}$	$\{u_0, u_2\}$

Table 3.4: Preferences of hyperplans in H relative to subsets of U with cardinality > 1

Let M be a model $\langle S, V \rangle$ for our fragment of \mathcal{L}_V such that:

- $V(p) = u_0 \times H \times \{U\}$
- $V(q) = u_1 \times H \times \{U\}$
- $V(r) = u_2 \times H \times \{U\}$

That is, the denotation of each sentential constant results from combining each element of a certain subset of worlds—a worldly component—with each and every hyperplan in H and with the set of U of propositions (note that what figures in the denotation of these sentences is U itself, not the elements in U). For example, the denotation of p is the set of w, h, U triplets such that every world w is a world in which Kanye gets 0 nominations and h is every possible hyperplan. This results in a world-sensitive but hyperplan-insensitive denotation for p .

Relative to model M , the following hold:

- (3.35) $\llbracket \uparrow q \rrbracket_{\langle w_0, h_0, U \rangle} = 1$
...because the set of worlds w' such that $\llbracket q \rrbracket_{\langle w', h_0, U \rangle} = 1$, namely u_1 , is among the preferred propositions for h_0 at U .
- (3.36) $\llbracket q \geq r \rrbracket_{\langle w_0, h_0, U \rangle} = 0$
...because not every subset of U containing the worldly components of q and r (u_1 and u_2 respectively) is such that, if h_0 chooses u_2 , u_1 is also chosen: relative to the set $\{u_1, u_2\}$, h_0 chooses u_2 but not u_1 . Conversely,
- (3.37) $\llbracket r > q \rrbracket_{\langle w_0, h_0, U \rangle} = 1$
...because there is a subset of U containing u_1 and u_2 , namely $\{u_1, u_2\}$, at which h_0 prefers u_2 but not u_1 .
- (3.38) $\llbracket p \approx r \rrbracket_{\langle w_0, h_2, U \rangle} = 0$
...because not every subset of U containing u_0 and u_2 at which h_2 chooses u_2 is one where u_0 is also chosen. To wit: $\{u_0, u_1, u_2\}$. By contrast,
- (3.39) $\llbracket p \approx r \rrbracket_{\langle w_0, h_1, U \rangle} = 1$
...because every subset of U at which h_1 prefers u_0 is one at which h_1 prefers u_2 , and vice versa (there is only one such subset, $\{u_0, u_2\}$).

We will not bother the reader by giving the “dis-preferential” profile of all these hyperplans, but to illustrate how \Downarrow works, let the negative preference function \mathcal{U} for h_0 be as follows:

$$\begin{aligned}\mathcal{U}\langle h_0, \{u_0, u_1, u_2\} \rangle &= \{u_0\} \\ \mathcal{U}\langle h_0, \{u_1, u_2\} \rangle &= \{u_1\}\end{aligned}$$

$$\begin{aligned}\mathcal{U}\langle h_0, \{u_0, u_1\} \rangle &= \{u_0\} \\ \mathcal{U}\langle h_0, \{u_0, u_2\} \rangle &= \{u_0\}\end{aligned}$$

In virtue of this, the following holds:

$$(3.40) \quad \llbracket \Downarrow p \rrbracket_{\langle w_0, h_0, U \rangle} = 1$$

...because u_0 is among the dis-preferred propositions for h_0 at U .

3.6 Order and consistency

In this section, we take up the issue of how a scalar system is built using the kind of semantics that we have proposed for \mathcal{L}_V . The main problem that this type of Delineation-inspired framework has is that it does not guarantee consistency across sets of alternatives—a hyperplan could, in principle, “switch” its choice, resulting in an inconsistent preference pattern. To prevent this, we need to say something about how an ordering system is induced on a set of alternatives by recursively applying a hyperplan to it.

3.6.1 Inducing order

Consider the set of alternatives U , where Kanye releases a new album and there are three possible developments: he receives 0, 1 or 2 nominations (assuming there are no other possibilities).

$$(3.41) \quad U = \left\{ \begin{array}{l} 0 \text{ nominations} \\ 1 \text{ nomination} \\ 2 \text{ nominations} \end{array} \right\}$$

Let h be a hyperplan that, faced with this description, prefers the development where Kanye receives two nominations. We can write this as follows:

$$(3.42) \quad h(U) = \{2 \text{ nominations}\}$$

Given that the development in which Kanye gets 2 nominations is preferred over each of the other two, h induces a very simple order onto U . We can represent this as a set of ordered pairs:

$$(3.43) \quad \left\{ \begin{array}{l} \langle 2 \text{ nominations}, 1 \text{ nomination} \rangle \\ \langle 2 \text{ nominations}, 0 \text{ nomination} \rangle \end{array} \right\}$$

This represents the fact that hyperplan h prefers that Kanye receive 2 nominations to him receiving 1, and that h prefers him receiving 2 nominations to him receiving none (the first element in each pair is the preferred one). But note that the outcomes in which Kanye receives 0 and 1 nominations are not ordered with respect to each other. This order is therefore partial.¹⁹

¹⁹More specifically, it is a STRICT PARTIAL ORDER. A relation R on a set X is a strict partial order iff R is:

1. Partial: there are elements x, y in X such that neither xRy nor yRx is the case;
2. Irreflexive: for every element x in X , it is not the case that xRx ;
3. Antisymmetric: for every elements x, y in X , if xRy and yRx then $x = y$;

As we said, hyperplans may select more than one alternative. Let us consider what the situation would be like in that case. Suppose that another hyperplan h' is such that both 1 and 2 nominations are selected at U . We can write that like this:

$$(3.44) \quad h'(U) = \begin{cases} 2 \text{ nominations} \\ 1 \text{ nomination} \end{cases}$$

This would then induce the following ordering on U ,

$$(3.45) \quad \begin{cases} \langle 2 \text{ nominations}, 0 \text{ nomination} \rangle \\ \langle 1 \text{ nomination}, 0 \text{ nomination} \rangle \end{cases}$$

where 2 and 1 nominations are preferred to 0 nominations, and 2 and 1 nominations are not ordered with respect to each other.

Back to hyperplan h again. Let us now consider what h would prefer, if faced with a subset of the original set of alternatives U . Consider the subset U' of U containing only the two alternatives that were not originally ranked by h , namely 1 and 0 nominations,

$$(3.46) \quad U' = \begin{cases} 0 \text{ nominations} \\ 1 \text{ nomination} \end{cases}$$

and let $h(U')$ be the result of applying hyperplan h to U' :

$$(3.47) \quad h(U') = \{1 \text{ nomination}\}$$

This induces the following ordering onto U' :

$$(3.48) \quad \{\langle 1 \text{ nomination}, 0 \text{ nominations} \rangle\}$$

We may wonder now what is the relation between the preferences of h at U and U' , or more generally, between sets of alternatives related by set inclusion. As far as we have established, a hyperplan could reverse its preferences when applied to “bigger” or “smaller” sets of alternatives. For instance, we said that, relative to set of alternatives U , h selects the alternative in which Kanye receives 2 nominations to both 1 and 0. Now consider a different set $U'' (\subseteq U)$,

$$(3.49) \quad U'' = \begin{cases} 2 \text{ nominations} \\ 1 \text{ nomination} \end{cases}$$

and suppose that h returns the following subset of U'' :

$$(3.50) \quad h(U'') = \{1 \text{ nomination}\}$$

4. Transitive: for every elements x, y, z in X , if xRy and yRz then xRz .

Elements x, y in X of which it holds that neither xRy nor yRx are INCOMPARABLE. This is different from saying that they are symmetric with respect to R , which would imply (by the anti-symmetric property) that they are identical.

Nothing we have said so far prevents hyperplan h from selecting the alternative in which Kanye receives 1 nomination relative to U'' , while at the same time selecting the outcome in which Kanye receives 2 nominations relative to U . That, however, would be undesirable. Why? There is a short and a long answer to that question. We finish this section with the short answer, and we tackle the longer answer in §3.6.2.

The short answer is that we want hyperplans to represent the meaning of evaluative sentences, and if we do not ensure that our hyperplan semantics rules out orderings such as these, down the line we would be allowing models that can represent the meaning of inconsistent sets of natural language sentences (that is, sets of sentences all of whose members cannot be true). For example, we would end up allowing possible indices of evaluation relative to which all three of the following sentences are true: *p is good*, *q is bad*, *q is better than p*.

To rule out such models, we need to recur to the same principles of consistency that guide the successive application of the predicate *tall* in the Delineation approach framework, for the same reason: just as recursively applying the predicate *tall* to comparison classes and their subsets does not by itself guarantee that the resulting extension of *tall* represents a consistent ordering of individuals, recursively applying a hyperplan to sets and subsets of alternatives does not guarantee that a consistent ordering of those alternatives will be obtained. That is, essentially, the answer we gave in §2.6 for why we need an axiom of **No Reversal**.

From a semantic standpoint, the short answer is enough. The long answer is nonetheless interesting and philosophically relevant, given that we are avowing a form of non-factualism about value. Simply put, a non-factualist cannot (simply) say that the reason why a certain set of evaluative sentences is inconsistent is because it misrepresents reality, or because it represents an impossible state of affairs, or something similar. But the reason why the short answer is sufficient from a semantic standpoint is that it is not at all clear that it is the semanticist's job to say *why* any set of sentences is inconsistent, nor that the non-factualist is in a worse position to answer that question than the descriptivist is. Semanticists can simply enforce certain properties on the models that they use so that those models accurately represent the meaning of the expressions that they wish to apply them to.

3.6.2 Preserving consistency

In order to tackle the longer answer to the question of why we want to rule out inconsistent hyperplans, let us recap our argument so far. The goal of this chapter has been (and part of the next will be) to argue that the operators introduced by \mathcal{L}_V represent basic semantic properties of positive and negative evaluative adjectives when they appear in positive, comparative and equative form with full clauses as their arguments.

For example: \Uparrow represents the positive form of an evaluative adjective of positive valence when applied to a sentence (*it is good that Kanye gets 0 nominations*); \Downarrow represents the positive form of an evaluative adjective of negative valence when applied to a sentence (*it is bad that Kanye gets 1 nomination*); $>$ represents the comparative form of an evaluative adjective of positive or negative valence when applied to a pair of propositions (*it is better that Kanye gets 1 nomination than that he gets 0*), and so on.

The problem is that nothing in our semantics for \mathcal{L}_V blocks a sentence like (3.51),

- (3.51) it is good that Kanye gets 0 nominations, it is bad that Kanye gets 1 nomination and it is better that Kanye gets 1 nomination than that he gets 0.

which in \mathcal{L}_V would correspond to the conjunction of the following three formulae:

- (3.52) a. $\uparrow p$
 b. $\downarrow q$
 c. $q > p$

In the semantics we just gave, those formulae are consistent. This means that we can find an index of evaluation that makes all of them true.

To see this, let us turn again to our previous toy fragment and structure (§3.5). This time however, consider the preferences of hyperplan h_3 in Table 3.4. In addition, let h_3 dis-prefer $\{u_1, u_2\}$ relative to U . That is, let $\mathcal{U}(h_3, U) = \{u_1, u_2\}$.

h_3 makes each of the formulae in (3.52) true: h_3 prefers outcome u_0 , making $\uparrow p$ true; it dis-prefers outcome u_1 , making $\downarrow q$ true too; and finally there is a subset of U , namely $\{u_0, u_1\}$, relative to which h_3 prefers u_1 and not u_0 , thus making $q > p$ true.

But if this is how \mathcal{L}_V works, then it cannot possibly represent a basic semantic property of English evaluative adjectives, because as we have seen, those adjectives do not work like that: if a proposition p has positive value, and another proposition q has negative value, then it cannot be the case that q has more value than p .

To prevent that, we need to block either of those three formulae from being true at an index that makes the other two true as well. Here is where the axiom of **No Reversal** comes in.

Axiom 4 (No Reversal) *For any structure $S = \langle W, H \rangle$, any hyperplan h , any two propositions u, u' and any set of propositions U such that $u \in h(U)$ and $u' \notin h(U)$, there is no $U' \subseteq U$ such that $u' \in h(U')$ and $u \notin h(U')$.*

The first axiom blocks any hyperplan from behaving like h_3 and *a fortiori* prevents any index from making (3.52a)-(3.52c) true: for any propositions u, u' , if any hyperplan prefers u but fails to prefer u' , then by **No Reversal** we won't find a smaller set of propositions such that that hyperplan's preferences are reversed.²⁰

²⁰ As we saw in §2.6, two other axioms are often invoked in Delineation semantics approaches (Benthem 1982; Burnett 2017; Klein 1980), although strictly speaking we do not need them to block (3.52a)-(3.52c). For any structure $S = \langle W, H \rangle$, any hyperplan h , any two propositions u, u' and any set of propositions U such that $u \in h(U)$ and $u' \notin h(U)$,

Axiom 5 (Upward Difference) *For all U' such that $U \subseteq U'$, there are propositions $v, v' : v \in h(U')$ and $v' \notin h(U')$;*

Axiom 6 (Downward Difference) *For all U' such that $U' \subseteq U$ and $u, u' \in U'$, there are some propositions $v, v' : v \in h(U')$ and $v' \notin h(U')$.*

Upward Difference guarantees that if a hyperplan selects a proposition and does not select another, the hyperplan will select and not select some propositions at any bigger sets of alternatives. In other words, a choice relative to a set of alternatives implies a choice at any bigger set of alternatives (not necessarily involving the same propositions). We illustrated this in Table 3.4 with h_4 : relative to $\{u_0, u_1\}$, for example, h_4 prefers u_0 and not u_1 . But relative to the bigger U , all propositions are preferred, and therefore, h_4 does not respect **Upward Difference**.

Downward Difference guarantees that any choice pattern will imply a similar choice pattern at any smaller set of alternatives containing the two propositions that were respectively selected and not selected. Note however, that the axiom requires that the choice pattern is preserved, but not that it involves the *same* propositions. h_5 in Table 3.4 illustrates a failure to respect **Downward Difference**: relative to U , u_2 is preferred and u_0, u_1 are not;

We can now come back to the longer answer to the question of why we need to prevent our hyperplan semantics from allowing certain orderings. In other words, why we need NO REVERSAL. The question is essentially the same as with the Delineation approach to dimensional adjectives like *tall*: in both cases, recursively applying a non-scalar predicate across sets and subsets of objects does not by itself guarantee that the resulting extension is a consistent ordering. But the longer answer to that question is prompted by a contrast between non-evaluative adjectives like *tall* and evaluative adjectives like *good*.

For the non-factualist about *good*, a straightforward answer for why we need **No Reversal** becomes unavailable. Simply put, failure to apply the adjective *tall* consistently to a set of individuals results in a misrepresentation of reality—for example, two individuals might end up being represented as taller than each other, and that cannot possibly be true. By contrast, since we are non-factualists, we cannot say that failure to apply the adjective *good* consistently to a set of alternatives—for example, if two propositions are represented as better than each other—would result in a misrepresentation of reality... because we are not representing reality when we use *good*! Given our hyperplan semantics, at most we can say that a failure for a hyperplan to apply consistently across a set of alternatives would be a misrepresentation of the hyperplan itself, which would represent an inconsistent plan of action.

This is a version of the problem of inconsistency for expressivists. In particular, it comes closest to the version of the problem that Gibbard’s *plan*-expressivism faces (Baker and Woods 2015; Charlow 2014, 2015; Schroeder 2008c; Silk 2015; Willer 2017). The question is: why must hyperplans be consistent, in other words, what is wrong about a hyperplan with contradictory preferences? If there were inconsistent hyperplans, they could figure in indices of evaluation relative to which certain inconsistent set of evaluative formula would be true.

However, note that this is not a question that arises in our semantics. This is because we are giving a thoroughly *disquotational* treatment of Boolean connectives. By doing this, we account for object-language (that is, \mathcal{L}_V) inconsistencies in virtue of inconsistencies in the metalanguage. Consider what our semantics predicts the meaning of any formula and its negation to be, that is, $\varphi \wedge \sim \varphi$ (abstracting away from the set of alternatives $\mathcal{P}(W)$):

$$(3.53) \quad \begin{aligned} & \llbracket \varphi \wedge \sim \varphi \rrbracket_{\langle w, h \rangle} = 1 \text{ iff} \\ & \llbracket \varphi \rrbracket_{\langle w, h \rangle} = 1 \text{ and } \llbracket \sim \varphi \rrbracket_{\langle w, h \rangle} = 1 \text{ iff} \\ & \llbracket \varphi \rrbracket_{\langle w, h \rangle} = 1 \text{ and } \llbracket \varphi \rrbracket_{\langle w, h \rangle} = 0 \end{aligned}$$

Our semantics predicts that there is no index of evaluation that can make this sentence true, because no index of evaluation can both have and not have a property, namely the property of being such that φ is true at it. This is the case regardless of whether the relevant formula is descriptive or evaluative, that is, regardless of whether it imposes a condition on the worldly or hyperplan component of the index. And this is so regardless, too, of whether we admit inconsistent hyperplans in our semantic structure—after all, in virtue of our composition rules, a contradictory set of formulae results in a contradiction in the metalanguage, and that is all we need.²¹

but at the smaller set containing both u_0 and u_2 , namely $\{u_0, u_2\}$, both are preferred, thereby blurring the choice pattern at U .

²¹See Charlow 2015 for a thorough exploration of the perils of this type of disquotational account of semantic properties *vis-à-vis* certain species of expressivism.

3.7 Conclusion

This completes our exposition of the propositional language \mathcal{L}_V . In this chapter, we have presented Gibbard's notion of hyperplans as a way of accounting for the normative content of judgments. We explored Yalcin's compositional and intensional implementation of Gibbard's approach, and we found some problems with it. To overcome those problems, we proposed to think of hyperplans as functions from sets of alternatives to subsets thereof representing the preferred, or chosen alternatives within those sets. We then saw how this hyperplan framework can be offered as a semantics for a simple propositional language enriched with evaluative operators representing evaluative properties and relations over propositions. In the next chapter, we will use this semantics as a template from which to start looking at the behavior of evaluative adjectives in natural language.

Chapter 4

From \mathcal{L}_V to natural language

Summary

This chapter aims to apply the scalar hyperplan framework presented in Chapter 3 to evaluative adjectives across the board. The chapter is divided in three parts. The first explores how to map the propositional language explored in chapter 3 onto all-purpose evaluatives like *good*, *bad*, and *better*, when these take propositional complements. The second part is devoted to both applying this semantics to uses of those adjectives with other types of arguments. The last part of the chapter is dedicated to exploring further semantic properties of evaluative adjectives, such as the relative-absolute distinction and the internal properties of their scales.

4.1 Introduction

The formulae that result from applying the operators introduced by \mathcal{L}_V , namely \Uparrow , \Downarrow , \geq , $>$ and \approx , are intended to represent evaluative sentences of natural language. Two features of \mathcal{L}_V introduce an initial constraint in this project: first, the evaluative operators of \mathcal{L}_V behave like relative-standard adjectives. This is shown, for example, by the fact that $\varphi > \psi$ does not entail $\Uparrow \varphi$ nor $\Downarrow \psi$. Indeed, we have been assuming that evaluative adjectives are relative-standard, and we have used that observation to argue for our hyperplan semantics (later in this chapter we will find reasons to challenge this assumption, but for now let us set it aside). Secondly, the arguments of the evaluative operators of \mathcal{L}_V are propositions. This means that the evaluative formulae of \mathcal{L}_V most accurately represent evaluative adjectives' lives as sentential operators, such as in the following examples:

(4.1) It is bad that you sent that e-mail.

(4.2) It would be good if you also took the trash out one of these days.

(4.3) What he did to her is so much crueler than what she did to him!

In all these cases, the arguments of those adjectives are propositions (or sets thereof). Our claim in the first part of this chapter is that the evaluative formulae of \mathcal{L}_V represent an essential semantic component of those sentences, so we will start by giving entries for thin evaluatives (*good*, *bad*) taking propositions as arguments, based on the operators of \mathcal{L}_V .

However, evaluation in natural language takes many forms, and in the second part of this chapter we will see that considering other kinds of evaluative sentences quickly complicates things.

We will enrich our approach in three (interrelated) directions: the first is to consider how to apply the hyperplan semantics to objects of evaluation other than propositions, chiefly action types and individuals. The following are examples of evaluatives applied to action types:¹

(4.4) Volunteering is good.

(4.5) Stealing is bad.

This extension of our proposal is unproblematic. It simply requires treating the grammatical subject of those sentences as containing a silent *PRO* pronoun, which we can fill out in various ways. A sentence like (4.4) will turn out to be equivalent to a sentence like *it is good that PRO volunteers*, where *PRO* might be interpreted in various ways.

Applying the hyperplan semantics to individuals, on the other hand, requires further tweaks. In a nutshell, we will say that a sentence where an evaluative adjective is applied to an individual does two things: (i) it predicates certain non-evaluative and potentially very underspecified properties of that individual, and (ii) it avows a preference for those properties—importantly, *not* for the individual *nor* for the proposition that the individual has those properties—over certain alternatives.

(4.6) This car is good.

(4.7) The mushrooms are even worse than the spinach.

In other words, the evaluation of an individual will be factored out into a regular, descriptive statement about the properties of that individual, together with an evaluation of those properties, which essentially follows the model for the evaluation of action types. An utterance of (4.6) attributes certain underspecified descriptive properties to a car and expresses, or avows, a preference for those properties over certain alternative properties.

The second direction in which we will enrich our proposal is applying it to evaluatives of other flavors, in particular aesthetic adjectives and predicates of personal taste (PPTs), such as the following:

(4.8) This car is beautiful.

(4.9) This soup is disgusting.

Aesthetic adjectives and PPTs are for the most part predicated of individuals or action types (rather than full propositions), so our proposal will essentially follow the model of evaluatives applied to individuals or action types. The fact that these predicates introduce particular flavors of evaluation will be cashed out by introducing restrictions on the type of non-evaluative properties that these adjectives attribute to their objects: whereas a sentence like (4.6) predicates certain non-evaluative and underspecified properties of a car, (4.8) does basically the same thing, but in addition it introduces certain restrictions on the type of non-evaluative properties that are predicated of the car—in the case of aesthetic adjectives and PPTs for example, those properties have to be perceivable. In other words, evaluatives of particular flavors reduce or constraint the underspecificity in the descriptive component of these adjectives.

¹By action types, we mean things like *stealing*, *loitering* or *volunteering*. Evaluation of action-tokens, by contrast, fall under the propositional model (i.e., (4.3)).

Finally, the third way in which we are going to expand our account are thick evaluative adjectives. Thick adjectives, as we saw in the introductory chapter, are those that carry in their meaning a descriptive and a evaluative component that are often taken to be separable from each other. These are some examples:

- (4.10) Marielle Franco was very brave.
- (4.11) Nathalie is extremely generous with her time.
- (4.12) Hiding your office mate's keys is cruel.

Being *generous*, for example, has the descriptive meaning of *giving beyond what is due*. But in addition to this, calling a person *generous* involves a positive evaluation, which for *generous* can be paraphrased as *being good in virtue of giving beyond what is due*.

To extend our account to thick terms, all we need to do is plug our story about thin evaluation into the evaluative component of these terms. So when we call someone *generous*, we are both saying that she gives beyond what is due, and in addition to this, we would be expressing a preference for the type of action that makes her generous (namely, giving beyond what is due). In our view, the difference between thin and thick evaluatives (including thin adjectives of specific flavors of evaluation) is not that the former have one less component of meaning than the latter. At least when evaluative adjectives are applied to individuals (as they are most of the time) thin and thick adjectives both involve a descriptive and an evaluative component. The difference between them, then, lies in the fact that thin evaluatives admit a much higher underspecificity in their descriptive component. Whereas a *good* person can have any combination of properties whatsoever, a *generous* person is invariably someone who gives beyond what is due. And in addition, in both cases the speaker is also expressing a preference for those sets of properties over certain alternatives.

Crucially, what we say here makes only one assumption about the debate on thick adjectives, namely that the descriptive and evaluative component of these terms can be distinguished. But it remains neutral with respect to the status of the evaluative component—in our case the hyperplan component, that is, whether it is an entailment, a presupposition or an implicature (or yet another thing altogether; see Cepollaro and Stojanovic 2016; Väyrynen 2013 for discussion).

Finally, in the third part of this chapter we will explore further scalar properties of evaluatives, some of which appear to be shared among them and some of which not so clearly. In particular, we will consider their behavior with respect to different types of adjectival modifiers, such as measure phrases (# \$200 *generous*), ratio (*twice as beautiful*) or so-called endpoint modifiers (*slightly / perfectly cruel*), which offer hints about the nature of their scales and the relation to their thresholds for the positive form (whether they are absolute- or relative-standard adjectives). Furthermore, we will consider the behavior of these adjectives with respect to PPs denoting comparison classes, and we will tackle the question of their scale structure. The most common types of scale used in measurement theory are ORDINAL, RATIO and INTERVAL scales, and we will consider arguments and data points in favor of assimilating evaluative adjectives to each scale type. We will conclude that there are reasons both in favor of adopting an interval and a ratio scale for evaluative adjectives, and we will remain undecided.

This chapter is structured as follows: In §4.2, we will consider how to extend \mathcal{L}_V to natural language. We will first discuss how to move from \mathcal{L}_V , where the alternatives relative to which the unary evaluative operators are defined were held fixed, to natural language, where the relevant alternatives change with context (§4.2.1, §4.2.2). Then, we will give truth conditions to

sentences containing thin evaluative adjectives taking clausal arguments (§4.2.3). In section §4.3 we will look into various possible extensions of our proposal. In §4.3.1 we will see how to account for thin evaluative adjectives when they take objects of evaluation other than propositions, in particular action-types and individuals. We will then see how to extend our proposal to thin evaluative adjectives of specific flavors of evaluation (§4.3.2); and then to thick evaluatives (§4.3.3). We will then discuss our proposed solution to the Frege-Geach problem in both its “supra-” and “sub-sentential” version, which we introduced in Chapter 2 (§4.4). After that, in the last two sections of the chapter we will look at further scalar properties of evaluative adjectives, in particular the relative-absolute distinction §4.5 and the question of their scale structure (§4.6). §4.7 recaps and concludes.

4.2 Thin evaluatives as sentential operators

Recall that the purpose of introducing \mathcal{L}_V was to capture a basic semantic component of evaluative adjectives when they take full clauses as their arguments. We can put this idea to test by considering actual evaluative adjectives of English taking sentential clauses as arguments. After that, we will be in a position to offer a solution to the “sub-sentential” Frege-Geach problem, namely the problem of why is it the case that only the positive form of evaluative adjectives express outright practical attitudes, even though the comparative form of those adjectives is compositionally derived from the positive form. Solving that problem closes a long-winded argument that started in Chapter 2.

The sentences of English that the evaluative formulae of \mathcal{L}_V most closely resemble are adjectives like *good*, *bad*, *better* & *worse*, when applied to full sentences (another obvious candidate would be an adjective like *valuable*, or the predicate *to have value*). These adjectives are often called THIN evaluative terms in the metaethics literature (Väyrynen 2013), in virtue of the fact that they seem to do nothing except attribute value. If to call something *good* or *bad* is nothing but to attribute value to it, and if one is a non-factualist about value, then one may aim to give a non-factualist semantics to those adjectives. This is what we will do here: we assign lexical entries to *good*, *bad*, *better* & *worse* such that they behave just like some of the operators of \mathcal{L}_V .²

The chief difference between the semantics of natural language evaluatives and the semantics of \mathcal{L}_V is that the sets of propositions that feed a hyperplan’s argument can vary across indices of evaluation. When we gave the semantics for \mathcal{L}_V , we let hyperplans range over the set of all propositions (sets of worlds). In natural language, this is not the case—hyperplans take as argument contextually dependent sets of alternatives. To see this, remember Kanye’s predicament.

Following Kanye’s thinking, we are only considering as alternatives getting nominated for Song or Album of the Year. But we might as well ponder what would be more preferable, that Kanye, Kendrick or Beyoncé got the nomination for one of those two categories, say, Song of the Year. Thus, the particular set of alternatives among which preferences are selected and

²A signpost/reminder: this view is non-factualist because the type of semantic value that we assign to evaluative sentences imposes a condition on a hyperplan parameter of a semantic structure (which represents plans or commitments for action and/or behavior) and not the world parameter of that structure (which represents how the world is).

thereby ordered is also contextually sensitive, and it can change the truth value of a sentence like (4.13):

(4.13) It is good if Kanye gets 1 nomination.

We might clearly prefer Kanye to get 1 nomination, relative to a choice between 0, 1 or 2 nominations (for the main categories). That would make (4.13) true relative to that set of alternatives. But suppose that the alternatives were not {Kanye gets 0 nominations, Kanye gets 1 nomination, Kanye gets 2 nominations}, but rather, {Kanye gets 1 nomination, Beyoncé gets 1 nomination, Kendrick gets 1 nomination}. Relative to this other set of alternatives, we might prefer Kendrick or Beyoncé getting 1 nomination and not Kanye. Then, (4.13) would come out false.

There is, in addition, a second contextual aspect on which the truth-value of a sentence like (4.13) depends: holding fixed a set of alternatives (say, Kanye getting 0, 1 or 2 nominations), we could also change our mind about (4.13) if certain background assumptions were different. Remember that we are in a “racist Grammys” scenario: the Grammys systematically award African American artists only in “racialized” categories such as rap or R’n’B, thereby preventing them from winning in the main categories (Song and Album of the Year). But suppose that things were different. Suppose that we were in a “color-blind Grammys” scenario. Then, given the megalomaniac aspirations of Kanye (which we are identifying with, for the moment), getting only 1 out of 2 nominations in the main categories would be awful. We might then judge (4.13) to be false.

Importantly, neither of these two axes of contextual variability (set of alternatives and general background conditions) need to be considered as variations of our *values*, which we represent as a set of hyperplans. Rather, such variability goes into determining the set of alternatives that are fed as the argument of hyperplans.

The semantic structure for \mathcal{L}_V ignored all that contextual variability, since the set of alternatives was uniformly the set $\mathcal{P}(W)$ of all propositions. In natural language however, sets of alternatives are contextually modulated in the way that we have seen. As Kanye’s example suggests, such contextual modulation works in two ways: on the one hand, the set of alternatives to some proposition can be different at different contexts (e.g., Kanye getting 0, 1 or 2 nominations vs. Kanye, Beyoncé or Kendrick getting 1 nomination). On the other hand, background information about how the world is can also affect a hyperplan’s choice pattern relative to a fixed set of alternatives (the “racist Grammys” vs. “color-blind Grammys” backdrop). More generally, we can think of sets of alternatives as sets of propositions that represent possible choices, relative to certain backdrop assumptions: there is a choice between Kanye getting 0, 1 or 2 nominations (against a “racist Grammys” backdrop); and there is a choice between Kanye, Kendrick or Beyoncé getting 1 nomination (against a similar backdrop). Similarly, there is a choice between those different set of alternatives against a “color-blind Grammys” backdrop.

There is an important point of contrast between these two axes of contextual variability, the alternative-sensitivity and the “backdrop”-sensitivity: while it seems really difficult to tinker with the latter, the set of alternatives can clearly be affected by prosody and focus. In the following subsection, we are going to consider how this works. But in what remains we will abstract away from the second kind of contextual sensitivity, the “backdrop” sensitivity. This is because, even though backdrop assumptions do seem to play an important role in determining the truth-value of evaluative sentences, it is not completely clear how they do that. In

addition, ignoring that second axis of contextual sensitivity will also keep our formalism relatively manageable. In sum, in the next two subsections we are going to see how to incorporate alternative-sensitivity into the picture.

4.2.1 Alternatives, focus and background information

As we said in Chapter 3, the notion of alternatives comes from the formal semantics literature on a number of topics (focus, questions, free-choice, disjunction) and should not be confused with the notion of relevant alternatives sometimes used in (meta)epistemology. The alternatives of focus semantics are generated as alternatives to particular propositions, and are modulated *via* contrastive focus, among other mechanisms (see note 16 on Chapter 3). Let us see here how focus interacts with the alternatives mobilized by sentential *good*. This has been observed by Lassiter (2017, 204 and ff), who relies on a well-known point from Dretske 1972 about the interaction of focus and statements about reasons.

To start with, consider the following sentence:

(4.14) It's good that Ceasar crossed the Rubicon.

As we discussed, it is natural to think that the proposition *that Ceasar crosses the Rubicon* can be evaluated relative to two independent axes: a set of alternatives and a set of background assumptions. Hence, this proposition can figure in different sets of alternatives all of which can build the same available background information; and conversely, it can figure in a single set of alternatives defined against different background information.

Suppose, for simplicity, that the only relevant piece of information about the world is that according to Roman law, for a Roman general to cross the Rubicon is tantamount to declaring war on Rome. Against this backdrop, we can still consider the proposition that Ceasar crosses the Rubicon as belonging to different sets of alternatives. For example, other sets of alternatives could include people *other than Ceasar* crossing the Rubicon as well as Ceasar crossing rivers *other than the Rubicon*. Crucially, those two sets of alternatives can both incorporate the information that whoever crosses the Rubicon is declaring war. And obviously, that piece of information affects how we wage a certain proposition against its alternatives: when we consider that Ceasar crosses the Rubicon as opposed to other people crossing the Rubicon, we might be waging the consequences of people of different statuses entering Italy (in particular, if the person is not a general, doing that would not mean declaring war on Rome); and when we consider that Ceasar crosses the Rubicon as opposed to Ceasar crossing other rivers, we are waging the consequences of Ceasar doing something with legal consequences as opposed something inane.

Conversely, we can think of how different background information can affect the status of different members of a single set of alternatives. Suppose that we are considering the possibility of Ceasar crossing the Rubicon *vs.* him crossing the Metauro and the Po. But this time, suppose that, according to Roman law, crossing the Metauro (not the Rubicon!) is tantamount to declaring war. If that were the case, then the alternative with the gravest consequences in that set would be Ceasar crossing the Metauro; while Ceasar crossing the Rubicon would be inconsequential.

Now, note that the relevant alternative set—but not the background information—is affected by

focus. The basic phenomenon is exemplified by the following contrast (CAPS indicate prosodic stress):

- (4.15) a. It's good that CEASAR crossed the Rubicon.
 b. It's good that Ceasar crossed the RUBICON.

It seems clear that (4.15a) invokes an alternative set formed by substituting Ceasar for other relevant individuals in the scope of *it's good that*. What (4.15a) tells us is that Ceasar crossing the Rubicon is among the preferable options in a set that contains, say, Sallust crossing the Rubicon and Pompei crossing the Rubicon. By contrast, (4.15b) invokes an alternative set formed by substituting the Rubicon for other relevant rivers (or landmarks, or what have you). What (4.15b) says is that Ceasar crossing the Rubicon is preferable relative to an alternative set containing, say, Ceasar crossing the Po and Ceasar crossing the Metauro. To see more clearly, this consider the following: if you wanted to express a preference for Ceasar crossing the Rubicon *vs.* other Roman generals, you would not use (4.15b); and if you wanted to express a preference for Ceasar crossing the Rubicon *vs.* other rivers, (4.15a) would be out.

Importantly though, both (4.15a) and (4.15b) would be equally acceptable in contexts with different background information: whether crossing the Rubicon or the Metauro means war can affect the truth value of (4.14), but it doesn't affect the acceptability of (4.15a)/(4.15b). The point of all of this is simply to insist that we need to keep separated the contribution of the relevant alternative-set and the background information. Alternatives can be affected by prosody, background information cannot. Perhaps more would need to be said about how alternatives are defined, but in what follows we are going to assume that the set of alternatives for computing an evaluative adjective coincides with the set of alternatives relevant for computing a focus alternative set, and can be affected in a similar way.

4.2.2 Semantics: introducing the alternative-generating function

The simplest way to incorporate alternative-sensitivity into our semantics is to make it part of our semantic structure and indices of evaluation, so this is what we will do. Just as Yalcin (2012) introduced an information or epistemic state parameter, whose value was fed to the hyperplan of the index, and in the same way that Delineation semanticists introduce a comparison class parameter in the index of evaluation relative to which dimensional adjectives are evaluated (see §2.6; and e.g. Burnett 2017; Klein 1980, 1991, *a.m.o.*), we introduce into the semantic structure a parameter representing possible alternatives that can figure as arguments of hyperplans.

Definition 15 (Alternatives) *An alternative set is a function a from possible worlds to sets of sets of possible worlds such that, for every world w and for every set of sets of worlds U , $a(w) = U$ just in case every set of worlds in U is a potential outcome at w .*

At each world in W , an alternative-generating function generates a set of possible alternatives, and this set can serve as the argument of a hyperplan. To see how this function works, suppose that, at some world w , it generates the following set:

$$\left\{ \begin{array}{l} \text{Kanye gets 0 nominations} \\ \text{Kanye gets 1 nomination} \\ \text{Kanye gets 2 nominations} \end{array} \right\}.$$

Alternatively, a different function might generate, at w , the following set:

$$\left\{ \begin{array}{l} \text{Kanye gets 1 nomination} \\ \text{Beyoncé gets 1 nomination} \\ \text{Kendrick gets 1 nomination} \end{array} \right\}.$$

As we will see, when we evaluate sentences in which an evaluative adjective is applied to a proposition, the index of evaluation will supply a particular alternative-generating function, defined relative to the world of evaluation. This function will generate a set of alternatives. So for a sentence such as *it's good that Kanye gets 1 nomination*, the index of evaluation will supply a set of alternatives. Then, the semantics of the evaluative operator *it is good that* will ensure that the proposition in its scope (*that Kanye gets 1 nomination*) is among the possible alternatives, and the sentence will be true at that index of evaluation just in case the hyperplan of the index is such that it selects the proposition that Kanye gets 1 nomination as one of the preferred alternatives in that set.

Thus, our semantic structure will be composed of a set W of possible worlds, a set of hyperplans H and a set of alternative-generating functions A :

Definition 16 (Alternative-enriched semantic structure) $S = \langle W, H, A \rangle$

Expressions are evaluated relative to a model formed by structure S and the interpretation function V :

Definition 17 (Model) $M = \langle S, V \rangle$

And indices of evaluation define a value for the alternative-generating function parameter, in addition to the world and hyperplan parameters:

Definition 18 (Alternative-enriched index) $i = \langle w_i, h_i, a_i \rangle$

Declarative sentences are evaluated for truth at indices of evaluation. The valuation function V assigns subsets of $W \times H \times A$ as the semantic value of sentential constants.

Just as we saw in Yalcin's proposal (§3.3.2), we want atomic declarative sentences of English containing no evaluative vocabulary to be descriptive. This means that they should impose a condition on the world but not on the hyperplan nor on the alternative-generating function parameter of the index.

To do that, we let function V assign to the descriptive sentences sets of w, h, a triplets formed by pairwise combining each of a subset of worlds with every possible hyperplan and every possible alternative generating function. That way, sentential constants will come out world-but not hyperplan- nor alternative-sensitive. Where φ is any descriptive and atomic (i.e. non-Boolean) sentence of English,

$$(4.16) \quad \llbracket \varphi \rrbracket_{\langle w_i, h_i, a_i \rangle}^M = 1 \text{ iff } \langle w_i, h_i, a_i \rangle \in V(\varphi) \text{ (again, we skip superscript } M \text{ from now on.)}$$

Since φ is descriptive, the semantic value that V assigns to this sentence contains all hyperplans and all alternative-generating functions, but not all worlds. Therefore, only those worlds can fail to make φ true. In other words, the truth of this sentence is insensitive to the hyperplan and alternative-generating function of the index; it is only sensitive to the world parameter.

Let us take stock. What we have done so far—at least in terms of semantic machinery—is not too different from Yalcin’s proposal for deontic *ought*, which is the first semantics that we presented. Both Yalcin and I propose to introduce two things in the semantic structure relative to which natural language sentences are interpreted: first, a hyperplan parameter; and second, a parameter supplying the arguments of hyperplans. For Yalcin, hyperplans were functions from epistemic states to subsets thereof. The hyperplan parameter was defined accordingly, and the parameter supplying the hyperplans’ arguments was the usual sphere of epistemic accessibility.

On the other hand, I have defined hyperplans as functions from/to sets of sets of worlds (i.e., sets of classic propositions), and therefore, a sphere of epistemic accessibility is not the right kind of object to supply a hyperplan’s argument. Thus, we have swapped that parameter for an alternative-generating function, which gives us the right kind of object to serve as the argument of hyperplans (as we define them) and it is not *epistemic* (which is not necessary for the kind of normative expressions we are dealing with, i.e., evaluatives).

4.2.3 *Good, bad, better & worse*

We turn now to the semantics of evaluative sentences of natural language with thin evaluative adjectives taking sentential clauses as their arguments.

It is good that

It is good that will have roughly the same semantics as the operator \uparrow of \mathcal{L}_V , except for the fact that the hyperplan’s argument will be supplied by the alternative-generating function of the index.

In addition, we need to ensure that the proposition under evaluation—or more exactly, its “worldly” component, the set of worlds that make it true—is among the alternatives determined by the index. We can think of this as a kind of semantically hardwired “ought implies can” principle; in order to evaluate anything, it has to be an option. We write it ‘ $a_i[\varphi](w_i)$ ’; which reads: *the set of alternatives a_i including (the worldly component of) φ that are determined at world w_i .*

Thus, *it is good that* is a sentential operator: its arguments are propositions, and when it receives an argument, it returns a set of world-hyperplan-alternative set triplets. The truth-conditions of a sentence like *it is good that φ* are as follows:

$$(4.17) \quad \llbracket \text{it is good that } \varphi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, a_i \rangle} = 1\} \in h_i(a_i[\varphi](w_i))$$

A proposition φ is good, relative to an index of evaluation i determining a world w_i , a hyperplan h_i and a set of alternatives a_i just in case the set of worlds that constitute the *worldly* component of φ (or simply φ , for short) is among the set of alternatives a_i (including φ) that are chosen by h_i . In other words: the index supplies a set of alternatives, which includes φ ; and then the hyperplan selects a subset of those alternatives. If and only if φ is among the selected propositions, it is good that φ .

Turning to our previous example (*IN* stands for the sentence *Kanye gets 1 nomination*):

$$(4.18) \quad \llbracket \text{it is good that Kanye gets 1 nomination} \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ \{w' : \llbracket 1N \rrbracket_{\langle w', h_i, a_i \rangle} = 1\} \in h_i(a_i[1N](w_i))$$

In words: relative to an index $\langle w_i, h_i, a_i \rangle$, the sentence *it is good that Kanye gets 1 nomination* will be true just in case the worldly component of the sentence *Kanye gets 1 nomination* (the set of worlds at which Kanye gets 1 nomination) is among the alternatives that are chosen by the hyperplan h_i relative to the set of alternatives a_i (which includes Kanye getting 1 nomination).

Suppose that the alternatives at some index of evaluation are Kanye getting 0, 1 or 2 nominations. Then, this sentence will be true just in case the hyperplan of the index is such that, when it takes as argument the set of propositions $\{\text{Kanye gets 0 nominations}, \text{Kanye gets 1 nomination}, \text{Kanye gets 2 nominations}\}$, it returns a subset of that set that at least contains the proposition *that Kanye gets 1 nomination*.³

At least as good as, better than

Similarly to the operator $>$ of \mathcal{L}_V , the meaning of the comparative operator *it is better that* results from applying *it is good that* across subsets of the alternatives determined by the index.

But let us first look at the more basic *it is at least as good that*. This corresponds to the \geq operator of \mathcal{L}_V , and so will have the following denotation:

$$(4.19) \quad \llbracket \text{that } \varphi \text{ is at least as good as that } \psi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \forall a' \subseteq a_i, \\ \text{if } \llbracket \text{it is good that } \psi \rrbracket_{\langle w_i, h_i, a' \rangle} = 1, \text{ then } \llbracket \text{it is good that } \varphi \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$$

(4.19) says that, relative to an index of evaluation i , a proposition φ is at least as good as another proposition ψ just in case every subset a' of the set a_i of alternatives determined by the index is such that, if ψ is *good* relative to $h_i(a')$, then φ is *good* too relative to $h_i(a')$. In other words, there is no subset of a_i relative to which ψ comes out *good* but φ does not.

In turn, this means the following: for every subset a' of a_i , the following must hold: if ψ is *good* is true relative to $\langle w_i, h_i, a' \rangle$, then φ is *good* is true relative to $\langle w_i, h_i, a' \rangle$. This requires evaluating those sentences along the model of *it is good that*, which we have just seen. Each of those sentences has the following truth-conditions:

³We are purposefully abstracting away from a number of features of evaluative adjectives in their use as sentential operators. For example, sentence mood and whether the complement clause is an *if*- or *that*-clause alter what we may call the “epistemic status” of the complement clause in a *it is good that*-sentence. Consider the following generalization from Sode 2018, p. 407:

- | | | |
|-------|--------------------------------------|---------------------------------|
| (i) | It is good that the cat is fat | \leadsto the cat is fat |
| (ii) | It is good if the cat is fat | \leadsto the cat might be fat |
| (iii) | It would be good if the cat were fat | \leadsto the cat is not fat |

See Sode 2018 for discussion of this feature of *good* and its obvious connection to the semantics of conditionals and mood.

Another general feature of thin evaluatives like *good* and *bad* that we are bracketing for now is the fact that they have beneficiary arguments:

- (iv) It is good for you that the cat is fat.

The prepositional phrase *for you* in that sentence does not denote neither an experiencer (as similar PPs do when they appear with predicates of personal taste Glanzberg 2007; Stephenson 2007b; Stojanovic 2007) nor an opinion holder. We come back to beneficiary arguments in Chapter 5.

- (4.20) a. $\llbracket \text{it is good that } \psi \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$ iff $\{w' : \llbracket \psi \rrbracket_{\langle w', h_i, a' \rangle} = 1\} \in h_i(a'(w_i))$
 b. $\llbracket \text{it is good that } \varphi \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$ iff $\{w' : \llbracket \varphi \rrbracket_{\langle w', h_i, a' \rangle} = 1\} \in h_i(a'(w_i))$

And each of them is true just in case the proposition in its scope is among the alternatives selected by the hyperplan of the index, relative to the alternatives of the index. For (4.19) to be false, one must be able to find a subset of alternatives that makes (4.20a) true but (4.20b) false. That subset of alternatives must include both ψ and φ and it must be such that the hyperplan of the index selects ψ but not φ .

Let us look at an example:

- (4.21) That Kanye gets 2 nominations is at least as good as him getting 1 nomination.

This sentence has the following truth conditions (2N and 1N stand for *Kanye gets 2 nominations* and *Kanye gets 1 nomination*, respectively):

- (4.22) $\llbracket \text{That 2N is at least as good as that 1N} \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1$ iff $\forall a' \subseteq a_i$,
 if $\llbracket \text{it is good that 1N} \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$, then $\llbracket \text{it is good that 2N} \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$

(4.21) is true at an index i just in case every subset a' of the alternatives of the index is such that, if the proposition that Kanye gets 1 nomination is selected by the hyperplan of i at a' (where a' contains the proposition that Kanye gets 1 nomination), then the proposition that Kanye gets 2 nominations is equally selected by the hyperplan of i at a' (where a' contains the proposition that Kanye gets 2 nominations).

Suppose that the alternatives of the index a_i are Kanye getting 0, 1 or 2 nominations ($a_i = \{2N, 1N, 0N\}$). And suppose that the hyperplan of the index is as follows:

$$\begin{aligned} h_i, (\{2N, 1N, 0N\}) &= \{2N, 1N\} \\ h_i, (\{2N, 1N\}) &= \{2N, 1N\} \\ h_i, (\{2N, 0N\}) &= \{2N\} \\ h_i, (\{1N, 0N\}) &= \{1N\} \end{aligned}$$

Relative to this index, (4.21) comes out true, because there is no subset of a_i that contains both 2N and 1N as alternatives and such that h_i chooses 1N but not 2N.

Now suppose that we evaluate (4.21) relative to the same set of alternatives, but relative to a different hyperplan h' :

$$\begin{aligned} h', (\{2N, 1N, 0N\}) &= \{2N, 1N\} \\ h', (\{2N, 1N\}) &= \{1N\} \\ h', (\{2N, 0N\}) &= \{2N\} \\ h', (\{1N, 0N\}) &= \{1N\} \end{aligned}$$

Relative to this hyperplan, (4.21) comes out false, because there is a subset of a_i at which h' chooses 1N but not 2N, namely the set $\{2N, 1N\}$.

As we saw in the semantics of \mathcal{L}_V , it is straightforward to define the comparative *it is better than* from the equative *it is at least as good as*:

- (4.23) $\llbracket \text{it is better that } \varphi \text{ than that } \psi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1$ iff
 $\llbracket \text{that } \varphi \text{ is at least as good as that } \psi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1$ and
 $\llbracket \text{that } \psi \text{ is at least as good as that } \varphi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 0$.

For (4.23) to be true, we need to be able to find a subset of the alternatives of the index including both *relata* relative to which the hyperplan of the index selects φ but not ψ .

$$(4.23) \quad \begin{aligned} &[[\text{it is better that } \varphi \text{ than that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \exists a' \subseteq a_i \text{ such that} \\ &[[\text{it is good that } \varphi]]_{\langle w_i, h_i, a' \rangle} = 1 \text{ and } [[\text{it is good that } \psi]]_{\langle w_i, h_i, a' \rangle} = 0 \end{aligned}$$

Now consider the following sentence:

$$(4.24) \quad \text{It is better that Kanye gets 1 nominations than that he gets 2.}$$

Relative to our previous set of alternatives a_i and hyperplan h_i , (4.24) comes out false because every subset of a_i containing $2N$ and $1N$ relative to which h_i selects $1N$ is also a subset relative to which h_i selects $2N$. In order to make (4.24) true, we would need to find one such subset selecting $1N$ but not $2N$. Hyperplan h' , on the other hand, is like that: relative to the set $\{2N, 1N\}$, h' selects $1N$ but not $2N$. Therefore, (4.24) is true relative to h' (and a_i).

Summing up, in this semantics, *being good* just means being chosen among the relevant set of alternatives. And for something to be better than another thing is just for there to exist a set of alternatives including both options relative to which the former option is chosen but the second is not. Comparative “betterness” is derived from absolute “goodness”, following a Delineation approach.

It is bad that, it is worse that

As the reader probably expects, we assign to *it is bad that* a denotation based on that of the operator \Downarrow of \mathcal{L}_V :

$$(4.25) \quad [[\text{it is bad that } \varphi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \{w' : [[\varphi]]_{\langle w', h_i, a_i \rangle} = 1\} \in \mathcal{U}\langle h_i, a_i[\varphi](w_i) \rangle$$

A proposition φ is bad relative to an index i that determines a hyperplan h_i and a set of alternatives a_i just in case φ is a proposition that h_i dis-prefers at a_i (where a_i includes φ).

The meaning of the negative polarity comparative *worse* derives from the equative *at least as bad as*, which in turn is derived from *it is bad that* (just as *better* was derived from *at least as good as*, which was derived from *it is good that*). The equative would mean the following:

$$(4.26) \quad \begin{aligned} &[[\text{that } \varphi \text{ is at least as bad as that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \forall a' \subseteq a_i, \\ &\text{if } [[\text{it is bad that } \psi]]_{\langle w_i, h_i, a' \rangle} = 1, \text{ then } [[\text{it is bad that } \varphi]]_{\langle w_i, h_i, a' \rangle} = 1 \end{aligned}$$

(4.26) says that, relative to an index of evaluation i , a proposition φ is at least as bad as another proposition ψ just in case every subset a' of the set a_i of alternatives determined by the index is such that, if ψ is *bad* relative to $h_i(a')$, then φ is *bad* too relative to $h_i(a')$. In other words, there is no subset of a_i relative to which ψ comes out *bad* but φ does not. The comparative *worse* would be defined as follows:

$$(4.27) \quad \begin{aligned} &[[\text{it is worse that } \varphi \text{ than that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ &[[\text{that } \varphi \text{ is at least as bad as that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ and} \\ &[[\text{that } \psi \text{ is at least as bad as that } \varphi]]_{\langle w_i, h_i, a_i \rangle} = 0 \end{aligned}$$

For (4.27) to be true, we need to be able to find a subset of the alternatives of the index including both *relata* and relative to which the hyperplan of the index dis-prefers φ but not ψ .

$$(4.27) \quad \begin{aligned} &[[\text{it is worse that } \varphi \text{ than that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \exists a' \subseteq a_i \text{ s.t.} \\ &[[\text{it is bad that } \varphi]]_{\langle w_i, h_i, a' \rangle} = 1 \text{ and } [[\text{it is bad that } \psi]]_{\langle w_i, h_i, a' \rangle} = 0 \end{aligned}$$

Before moving on, a complication for this proposal should be highlighted (Carla Umbach, *p. c.*). Note that, as was pointed out in §2.5.1, positive form equatives invite an inference to the positive form, while negative form equatives do not:

(2.59) The first movie was as good as the second. \rightsquigarrow Both movies were good

(2.60) The first movie was as bad as the second. \rightsquigarrow Both movies were bad

Our semantics does not predict this, as the equative *as bad as* does not involve a reference to the positive form *bad* relative to the alternatives defined at the index. A simple solution would be to include that in the meaning of the equative, perhaps as a presupposition.

$$(4.28) \quad \begin{aligned} &[[\text{that } \varphi \text{ is at least as bad as that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ &[[\text{it is bad that } \varphi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ and } [[\text{it is bad that } \psi]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ and} \\ &\forall a' \subseteq a_i, \text{ if } [[\text{it is bad that } \psi]]_{\langle w_i, h_i, a' \rangle} = 1, \text{ then } [[\text{it is bad that } \varphi]]_{\langle w_i, h_i, a' \rangle} = 1 \end{aligned}$$

But since the comparative *worse* is derived from the equative, this simple solution would have the bad prediction that the comparative *worse* would also involve a reference to the positive form.

Rett (2007, p. 217), albeit with different motivations, offers the following alternative path: assume, first, that the equative means something stronger than *at least as*, namely *exactly as*. Then, equatives formed with antonym adjectives have the same truth-conditions: φ is as good as ψ is equivalent to φ is as bad as ψ . Secondly, assume that equatives are ambiguous between a “positive-entailing” (c.f., (4.28)) and a “positive-neutral” interpretation (c.f., (4.26)). That is, *as good/bad as* would each have the following interpretations:

(4.29) φ is as good as ψ

a. $\varphi \approx \psi$

b. $\varphi \approx \psi$ and both are good

(4.30) φ is as bad as ψ

a. $\varphi \approx \psi$

b. $\varphi \approx \psi$ and both are bad

Third, negative antonyms are in general marked with respect to their positive counterparts. For adjectives like *tall* and *short*, this is shown by e.g., the fact that only the positive antonym takes measure phrases (*she is 1.63m tall* {# *short*}). Thus, assume that *bad* is equally marked with respect to *good*. Finally, assume a principled dis-preference to marked forms when these are truth-conditionally equivalent. The result of all this is that the “positive-neutral” interpretation of *bad* is ruled out, and only the “positive-entailing” interpretation remains:

(4.29) φ is as good as ψ

a. $\varphi \approx \psi$

b. $\varphi \approx \psi$ and both are good

(4.30) φ is as bad as ψ

a. ~~$\varphi \approx \psi$~~

b. $\varphi \approx \psi$ and both are bad

On the other hand, this view predicts that positive form equatives are ambiguous between a “positive-neutral” and “positive-entailing” interpretations. We remain unsure as to whether this is the right prediction, but we will leave this issue for future work.

Relation between *good*, *better*, *worse* and *bad*

In Chapter 3, we assumed that the binary relations \geq and $>$ were inter-definable using \Uparrow and \Downarrow . Now, by contrast, we have two linguistic items, *better* and *worse*, each of which respectively defined using the simple predicates *good* and *bad*. What is, then, the relation between them? As we said, *good* and *bad* should not be duals: all that is not good isn't bad, and all that is not bad isn't good. We want to allow for certain things to be neither good nor bad. On the other hand, what we do want is that *better* and *worse* mirror each other. The semantics does not establish this, so we need to establish by brute force:

Axiom 7 (*Better* \approx *worse*) for any descriptive and atomic (i.e. non-Boolean) sentences of English φ and ψ and any index of evaluation i , φ is better than ψ is true i if and only if ψ is worse than φ is true at i .

Defining pairs of antonyms in this way has been observed to incur into problems (Kennedy 1999, 2013a), in particular when dealing with cross-polar anomalies, but we will ignore that issue for the time being.

The non-intersective character of the extensions of positive and negative hyperplans (let's call it *non-intersectionality*; see the definition of negative hyperplans in §3.3.3), the axiom of **No Reversal** defined in Chapter 3 and the axiom **Better** \approx **worse** defined here result in a number of welcome predictions in this semantics, which we present informally here.

First, a basic antonymity pattern holds between *good* and *bad*. In virtue of *non-intersectionality*, if *it is good that* φ is true, then *it is bad that* φ is false. And conversely, if *it is bad that* φ is true, then *it is good that* φ is false.

Secondly, note however that the previous entailment does not hold in the opposite direction: if *it is good that* φ is false, it does not follow that *it is bad that* φ is true. And conversely, if *it is bad that* φ is false, it does not follow that *it is good that* φ is true. This means two things: first, *good* and *bad* are not duals: all that is not good isn't bad, and all that is not bad isn't good. And second, some things are neither good nor bad. Both of these are welcome.⁴

Thirdly, φ is better than ψ is compatible with any application of *good* to any of φ and ψ , except *it is good that* ψ and *it is not good that* φ . This holds in virtue of **No Reversal**. Similarly, ψ is worse than φ is compatible with any application of *bad* to any of φ and ψ except *it is bad that* φ and *it is not bad that* ψ , also in virtue of **No Reversal**. But now we can go further. In virtue of **Better** \approx **worse**, these consequences of *better* / *worse* are interchangeable: φ is better than ψ rules out the conjunction of *it is bad that* φ and *it is not bad that* ψ ; and ψ is worse than φ rules out the conjunction of *it is good that* ψ and *it is not good that* φ . And finally, in virtue of *non-intersectionality*, the conjunction of *it is bad that* φ and *it is good that* ψ is ruled out too. In sum, if φ is better than ψ , then it cannot be the case that ψ is good and φ is bad.

This concludes our initial application of the language of \mathcal{L}_V to evaluative adjectives of natural

⁴Perhaps we would like as well to predict intra-adjectival gaps, that is, some things being neither good nor not good, and neither bad nor not bad. Since hyperplans are maximally decided, they will be defined for each and every possible alternative. To circumvent this, a procedure of supervaluation over hyperplans could be introduced such that φ is good relative to a set of hyperplans iff it is good relative to all hyperplans in the set, and so on. This might be independently desirable: recall that hyperplans are maximally decided planning states, but real world planners are hardly (if ever) maximally decided. Quantifying over hyperplans would be the obvious way of formalizing undecidedness.

language. Remember that we have focused on a relatively small fragment of English, namely the two simplest evaluative adjectives, *good* and *bad*, in positive and comparative form and taking sentences as arguments. It is time to expand this story, and we'll do that in three directions: first, we'll look at arguments *other* than sentences, namely action types and individuals. Then, we will see how to incorporate increasing degrees of “thickness” into the picture, that is, of descriptive elements.

4.3 Beyond evaluative operators

The sole bearers of value in \mathcal{L}_V are propositions, but very often we evaluate things other than propositions, such as action-types or properties (*murdering is morally bad*; *having a visa is more important than having an address*), people (*John is cruel*) or objects (*this knife is crap*). Moreover, *good* and *bad* are thin adjectives, but a large portion of our evaluative judgments are expressed using thin adjectives of particular flavors of evaluation (such as moral, aesthetic or personal taste adjectives) as well as so-called thick adjectives (*cruel*, *generous*, *nasty*).

Table 4.1 shows the different varieties of evaluation that we will explore in sections §4.3.1 to §4.3.3 of this chapter, resulting from the interaction of thin to thick forms of evaluation (left to right) and different objects of evaluation (top to bottom). It requires some guidance to navigate this table; this is found in the following paragraphs.

	Thin evaluation	Specific flavors of evaluation	Thick evaluation
Propositions	φ is good – $\Uparrow \varphi$	φ is beautiful – φ is [...] * – \Uparrow [...] *	φ is generous – φ is P – $\Uparrow P$
Action types / properties	F -ing is good – $\Uparrow F$ -ing	F -ing is beautiful – F -ing is [...] * – \Uparrow [...] *	F -ing is generous – F -ing is P – $\Uparrow P$
Individuals	x is good – x is [...] – \Uparrow [...]	x is beautiful – x is [...] * – \Uparrow [...] *	x is generous – x is P – $\Uparrow P$

Table 4.1: Varieties of evaluation

The main complication introduced by the new varieties of evaluation is the occasional addition of a descriptive component. Recall that our view about sentential *good* (§4.2) is that these sentences express a practical commitment or intention to prefer a proposition over certain alternatives. This is represented in the uppermost left cell of Table 4.1 with the operator \Uparrow of \mathcal{L}_V . Crucially, those sentences do not predicate any further properties of the proposition under evaluation. All that there is to sentential *good* is (roughly) $\Uparrow \varphi$. By contrast, thick evaluatives have a descriptive component in addition to their evaluative component. As we will see, while the descriptive component simply consists in the predication of a set of descriptive properties, we will plug our hyperplan semantics to account for the evaluative component. As represented in Table 4.1 (right column), thick adjectives do two things: (i) they predicate descriptive properties (... is P), and (ii) they evaluate those properties ($\Uparrow P$).

The rest of varieties of evaluation fall somewhere between the simpler propositional model and the more complicated thick model. Moving downwards on the leftmost column to the evaluation of action-types or properties, we see the operator \Uparrow being applied to an action type or prop-

erty (*F-ing*), rather than a proposition. As we will see, thin evaluation of action-type/properties is very similar to the thin evaluation of propositions. Things get trickier when we consider thin evaluation of *individuals* (bottom left cell). There, even though we are still in the thin evaluation column, we already have a descriptive component. But importantly—and in contrast thick evaluatives—the descriptive component is a radically underspecified set of properties (this is written [...] in the table).

Moving rightwards, a descriptive component is always present. This is seen on the central (evaluatives of specific flavor) and right (thick evaluatives) columns. Consider the central column. Analogously to thin evaluatives when applied to individuals, evaluatives of specific flavors (moral, aesthetic, PPTs) also predicate a set of underspecified descriptive properties of their objects. But in contrast to the thin evaluation of individuals, those descriptive properties are not completely underspecified. That is, these adjectives introduce restrictions on the kind of descriptive properties that are predicated (this is written [...]*) on the table). For instance, to call an individual *beautiful* is to attribute certain, underspecified descriptive properties to it (*x* is [...]*) and, in addition, it is to express support for those properties (↑ [...]*). But since *beautiful* is an aesthetic adjective, the descriptive properties are not completely underspecified: at the very least, they have to be properties that can be experienced or perceived in some way. Thick evaluatives, as we said, are characterized by the fact that the descriptive component is not underspecified.

Finally, two cells in Table 4.1 are highlighted in red, in order to signal that we are unsure about how to treat these cases. First, it is not clear how to interpret evaluative terms of specific flavors when they are applied to full propositions. To illustrate this, consider the following two sentences (from *corpora*):

(4.31) It is beautiful that nature loves to hide.

(4.32) It is unethical that I see teaching as a job and not a “call of duty”.

It is not clear that *beautiful* works in (4.31) as anything but a term of very general praise—pretty much like *wonderful*, *great* or *good*. So we are hesitant to say that any descriptive properties are being predicated there of the proposition *that nature loves to hide*. Perhaps the right analysis of sentences like these reduces to the completely thin case. On the other hand, (4.32) sounds equivalent to a sentence like *seeing teaching as a job and not as a “call of duty” is unethical*, which pulls in the direction of assimilating sentences like (4.32) to the evaluation of action-types/properties.

Secondly, thick evaluatives taking full propositions as arguments either sound slightly odd or appear to be assimilable to the action-type/property case, just as (4.32). Here are some examples found in *corpora*:⁵

(4.33) It is generous that we get to keep him.

(4.34) It is cruel that they designed this level to be impossible to solve.

(4.35) It is rude that he didn’t show up.

Before moving on, let us mention that the view offered in the following pages is by no means intended as a total, complete theory of evaluative expressions. I am confident that the view just

⁵Examples do not abound: on the [iWeb corpus](#), *it is generous that* shows 1 hit, *cruel* shows 6 hits, *rude* shows 1 hit, and *nasty*, *lewd*, *selfish* or *noble* show no hits. Compare with *beautiful* (29 hits) and *good* (2542 hits).

defended for evaluative operators (that is, evaluative adjectives taking full clauses as arguments) is by and large correct, and I am also confident that a theory according to which evaluatives have a descriptive and an evaluative component *when they are predicated of individuals* is also necessary—for reasons to be laid out in Chapter 6. But I remain less confident about the extensions proposed in the sections to follow. The aim, however, is not to offer an exhaustive theory. Rather, it is to explore how the basic ingredients introduced for our analysis (namely, hyperplan semantics) can be applied to a variety of evaluative expressions. Each of those extensions of my proposal will surely incur individual challenges, and I will not tackle them here. In addition, nothing is said as to how this proposal could apply to evaluative terms beyond adjectives, such as nouns (including slurs) pejoratives and expressives. We turn now a more thorough analysis of (some of) the cells in Table 4.1.

4.3.1 Other bearers of value

Evaluative adjectives often take arguments other than propositions. Action types or properties and individuals are the most common. Here are some examples:

- (4.36) Stealing is bad.
- (4.37) To just stare without doing anything is just horrible.
- (4.38) That white Mercedes is good.
- (4.39) The picture that you got was uglier than the one I got.

Roughly, we are going to defend that, when evaluative adjectives take these types of arguments, they behave in pretty much the same way as when they take full clauses as arguments: they express (dis-)preference for a proposition over certain alternatives.

For instance, when an evaluative adjective takes as its argument a nominalized clause—(4.36) or (4.37), it will be argued that those clauses can be mapped onto regular propositions, and then the semantics for those sentences is no different from what we saw in previous sections. Thus, *stealing is bad* is roughly equivalent to *that [arbitrary individual] steals is bad*, which in turn expresses a dis-preference against the proposition *that [arbitrary individual] steals* over relevant alternatives.

On the other hand, the evaluation of individuals is more complicated. We propose to treat a sentence like (4.38) as the combination of a proposition that attributes certain descriptive properties to an individual, together with an evaluation of such properties that works in the same way as the evaluation of properties. Thus, (4.38) turns out to be equivalent to (i) *that white Mercedes has such-and-such properties* and (ii) *having such and such properties is good*. Let's consider the evaluation of action-types and individuals in turn.

Action types

Sentences like (4.36) are the prime examples of evaluative sentences that metaethicists tend to focus on.

- (4.36) Stealing is bad.

Our approach is to treat the subject of those sentence as denoting full propositions, which is relatively uncontroversial, and then give those sentences the same semantics we saw in §4.2. Arguably, (4.36) contains a covert *PRO*, so it means something like *PRO stealing is bad*. *PRO* stands for an arbitrary individual, and *PRO stealing* is equivalent to the proposition *that PRO steals*. Thus, (4.36) turns out to be equivalent to

(4.40) It is bad that *PRO* steals.

One may wonder about the exact reference of *PRO*. Three very salient options are an existential (= *that anyone steals*), a generic (= *that people steal*) and a *de se* reading (= *that I steal*) (see Moltmann 2006, for discussion). Each of these options have something in their favor, and something against it. The existential reading of (4.40) (*it is bad that anyone steals*) would be plausible in light of a somewhat strict set of values, but false in light of a more relaxed set of values that admits of exceptions. Things work exactly the other way around for the generic reading: under that reading (*it is bad that people steal*), (4.40) would express general dis-preference for acts of stealing, with exceptions. A sentence like (4.36) seems compatible with both interpretations. Finally, a *de se* interpretation would better capture a connection to the actual planning state of the speaker (as it would be roughly equivalent to *I choose/plan/prefer to not steal*), although it might appear too subjective—presumably, moral statements like (4.36) have some kind of universal force.

For our purposes however, it does not matter which of those readings is the right one. They key for the semantics of (4.36) is that its grammatical subject is assigned a regular, descriptive proposition. Once we have that, we can apply our hyperplan semantics in exactly the same way we saw in §4.2:

$$(4.41) \quad \begin{aligned} &[[\text{it is bad that } PRO \text{ steals}]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ &\{w' : [[PRO \text{ steals}]]_{\langle w', h_i, a_i \rangle} = 1\} \in \mathcal{U}\langle h_i, a_i[PRO \text{ steals}](w_i) \rangle \end{aligned}$$

In words: (4.40) will be true, relative to an index of evaluation, if the hyperplan of the index dis-prefers the (worldly component of the) proposition denoted by the grammatical subject of that sentence relative to the set of alternatives determined by the index (which must include that proposition as an option). Call a_i the set of alternatives at some index, and let $a_i = \{PRO \text{ steals}, PRO \text{ begs}, PRO \text{ borrows}\}$. (4.40) be true iff the proposition *that PRO steals* is in $\mathcal{U}\langle h, a_i \rangle$.

It is important to note that, even though *bad* takes a full proposition as its argument, the object of evaluation is still an action-type or property that is evaluated against other action-type or properties. In other words, what a speaker rejects when she utters (4.36) is the act of stealing, not the proposition *that PRO steals*. If it was, then in principle it should be possible to consider a set of alternatives were *PRO* is substituted for other individuals. Nonetheless, the right result is predicted by our semantics if alternatives are indeed computed as was suggested in §4.2.1, that is, if they are focus-sensitive: given that *PRO* is silent, it cannot be stressed, and thus the relevant set of alternatives can never be computed.

Incidentally, the view just laid out about sentences like (4.36) was arguably already Ayer's view, who wrote as follows (in what has come to be considered one of the foundational statements of metaethical non-cognitivism):

The presence of an ethical symbol in a proposition adds nothing to its factual content. Thus if I say to someone, 'You acted wrongly in stealing that money,' I am

not stating anything more than if I had simply said, ‘You stole that money.’ In adding that this action is wrong I am not making any further statement about it. I am simply evincing my moral disapproval of it. It is as if I had said, ‘You stole that money,’ in a peculiar tone of horror, or written with the addition of some special exclamation marks. The tone, or the exclamation marks, adds nothing to the literal meaning of the sentence. It merely serves to show that the expression of it is attended by certain feelings of the speaker. *If now I generalize my previous statement and say, ‘Stealing money is wrong,’* I produce a statement which has no factual meaning—that is, expresses no proposition which can be either true or false (Ayer 1946, p.142, my emphasis).

Ayer says that a sentence like *Stealing money is wrong* is a “generalization” from *you acted wrongly in stealing that money*. He does not have anything else to say about the relation between those two sentence types. However, his approach is similar to ours because it connects the two sentence types, and it lets him distinguish, in sentences like (4.36), a regular, factual component (i.e., a descriptive proposition), and an evaluative component. This evaluative component can easily be factored out of the factual component and treated like some kind of prosodic *addendum* (like exclamation marks).

Individuals

As we hinted at, accommodating evaluative sentences in which evaluative adjectives are predicated of individuals is a little less straightforward. Here are some examples:

(4.38) That white Mercedes is good.

(4.39) The picture that you got was uglier than the one I got.

There are at least two possible strategies: the first is to go for a view in which evaluative adjectives can take individuals as arguments, in addition to propositions. That view respects the surface grammar of sentences like (4.38) and (4.39), but it requires making evaluative adjectives semantically flexible enough to apply to both individuals and propositions. The second strategy is to assign to those sentences a structure such that the argument of the evaluative adjective is a proposition, and not an individual. This involves attributing to these adjectives a more complicated meaning, and one that does not respect the surface grammar of these sentences. By contrast to the case of nominalized clauses, this is definitely *not* standard—the orthodoxy is that nominal phrases denote individuals, or individual concepts, not propositions. However, the second approach lets us defend a single story about how evaluative adjectives work (one in which they are uniformly sentential operators), and one that, as we will see, sits nicely with available views about thick adjectives.

As we have already said (and as the reader can see from Table 4.1), we are going for the second kind of view. However, let us say something about how the first type of view could be spelled out. This view would assign a standard semantics to the grammatical subject of sentences like (4.38) and (4.39), and we would let evaluative adjectives take individuals as arguments. In turn, this requires making our hyperplan semantics liberal enough to be compatible with individuals. Informally, the idea would be the following: we have been describing hyperplans as functions that take you from a set of propositions to the (dis-)preferable proposition(s) in that set. The idea would be to generalize this, and think of hyperplans as functions that take you from a set

of objects of any type to the (dis)-preferable members of that set. The arguments of hyperplans could be propositions, but they could also be individuals, properties, etc. For instance, when we say that *a* is a good car, we would be avowing a practical commitment such that, when faced with a choice between *a* and a set of alternative cars, one chooses at least car *a*. Hyperplans, in this view, would be functions from sets of alternatives of *any semantic type* to their (dis-)preferred subsets.

In order to represent the fact that evaluative adjectives can take arguments of different semantic types, we would need to assign a variable type to their arguments. In addition to this, we would let the index supply alternative sets of the corresponding type—if one is evaluating individuals, then the alternatives should be a set of individuals; if one is evaluating propositions, the alternatives ought to be a set of propositions, etc. In this view, (4.38) would be true, relative to an index *i*, just in case the hyperplan h_i of the index selects *that White Mercedes* over whatever set of alternative cars (including the white Mercedes) is determined by *i*.

This proposal can be further spelled out, but it has drawbacks. First, adopting such a liberal account of hyperplans commits us to a rather powerful logic, whereby certain predicates (evaluative predicates) take both individuals and full propositions as arguments. From a logical point of view, that move is not for free, although that is not my concern here. Moreover, assuming that evaluatives are able to take both individuals and propositions as arguments would make evaluatives—or at least some of them—a class of natural language expressions with a rather unique distribution (besides logical connectives like *and* and *or* it is difficult to come up with other examples).

Secondly, there is an extra wrinkle to the evaluation of individuals which motivates adopting the second view we considered at the start—or so I will argue. In exploring this application of hyperplan semantics to individuals, we have been claiming that the evaluation of individuals is based on the capacity of hyperplans to select sets among sets of individuals. However, it seems more intuitive to think that the evaluation of individuals is guided, in some fundamental way, by the properties that individuals have, and not by the individuals themselves. Consider sentence (4.38) again. In virtue of what is that white Mercedes more preferable than a set of alternative cars? Intuitively, in virtue of the properties it has. That is, in virtue of its motor, speed, color, design, etc. A consequence of this is that, *ceteris paribus*, one should consistently prefer things with similar properties. And that seems to be the case. To see this, let *X* be a set $\{a, b, c\}$ of three cars, and suppose that we judge that car *a* is good relative to *X*. In addition, let *A*, *B* and *C* stand for the set of properties that each of those cars respectively and uniquely instantiates (*A*, for instance, may be a combination of different properties, such as *having low emissions*, *a large rear spoiler*, *being Japanese*...). Now, suppose that we had been presented with a different set *X'* of cars $\{a', b', c'\}$ which also instantiate properties *A*, *B* and *C* uniquely and respectively. It is very natural to assume that we would also judge that car *a'* is good relative to *X'*. Whatever set of properties one appreciates in an individual, those properties ought to make one appreciate to the same extent *any other* individual that has them—*ceteris paribus*.

If this is correct, then it seems that what we might have taken to be the evaluation of an individual “in itself” is actually the evaluation of its properties. And then, it would be more sensible to model hyperplans as functions from sets of *descriptions* of individuals, that is, properties, rather than sets of individuals. But if individuals are not, after all, the real objects of evaluation, why not just adopt our model for the evaluation of action types/properties? If saying *this car is good* is tantamount to saying *the properties of this car are good*, then we might as well model sentences like (4.38) as the evaluation of those very properties.

That will be in fact, our proposal. But note, that in addition to an evaluation of certain properties, it is also part of the meaning of (4.38) to predicate something *of* this particular individual, i.e., *that white Mercedes*.⁶ If I think that the value of a car is measured by the size of its rear spoiler, then when I utter (4.38) (somehow) I am also getting across that the white Mercedes has a large rear spoiler. And if you think that all that matters with respect to cars is their environmental friendliness, then when you utter (4.38) you are also getting across that the Mercedes is environmentally friendly. But having a large rear spoiler and being environmentally friendly are descriptive, not evaluative properties. So we propose to assign a double semantic value to (4.38): on the one hand, that sentence conveys something descriptive: that an individual has certain, non-evaluative properties. And on the other hand, it expresses an evaluation of those properties, along the model offered in §4.3.1. We can represent its meaning thus:

$$(4.42) \quad \text{That white Mercedes is good} = \begin{cases} \text{That white Mercedes is [...]} \\ \text{[...] is good} \end{cases}$$

The first component is a descriptive component, whereby a certain, contextually determined set of descriptive properties is predicated of an individual. What those properties are depends on whatever the speaker perceives the preferable properties of a car to be. It could be having low emissions, or having a large rear spoiler, or being made in Japan, or a combination of some of those things. Importantly, thin evaluatives like *good* leave massively underspecified what those descriptive properties are (we write that ‘[...]').

The evaluative component, which is an expression of support for those properties, is what is common to all and any uses of *good*. To account for this, all we need to do is plug our account for sentences like *F-ing is good*. Assume that *F* stands for the properties that we just represented as [...]. Then, the evaluative component of (4.42) would be equivalent to a sentence like *it is good that PRO is F*. We can then assign the following truth conditions to the evaluative component of (4.42):

$$(4.43) \quad \begin{aligned} &[[\text{it is good that } PRO \text{ is } F]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ &\{w' : [[PRO \text{ is } F]]_{\langle w', h_i, a_i \rangle} = 1\} \in h_i(a_i[PRO \text{ is } F](w_i)) \end{aligned}$$

As we said before about evaluation of action-types and properties, it is important to note again that sentences like (4.38) do not express preference or rejection for a full proposition, but for a set of properties. This is predicted by our semantics, insofar as the object of evaluation for a sentence like (4.38) is not the proposition *that that white Mercedes is F*, but rather the proposition *that PRO is F*. If the object of evaluation was the proposition *that that white Mercedes is F*, then alternatives could be generated were other cars are *F*, but this would misconstrue the meaning of (4.38), as (4.38) cannot be used to express a preference for this particular individual having a certain set of properties *F* over other individuals having *F* (in order to express *that*, one would have to say something like: *it is good that THIS WHITE MERCEDES is F*).

To say all of this, of course, is to follow in Hare’s footsteps (Hare 1952). Hare defended that there are two components to the meaning of *good*, one variable and the other one stable: the variable component was what he calls a varying standard for goodness, which is nothing but a certain set of non-evaluative properties. However, there is something which people with

⁶ I say part of its meaning, but not necessarily part of what is asserted when that sentence is uttered at a context.

radically different standards of goodness have in common when they use the word *good*, and that is its invariant, commendatory component: anyone who uses the word *good* is invariably commending whatever they apply the word to (see also Umbach 2016, 146 and ff, who calls the descriptive meaning of *good* a ‘quasi-denotation’).

This view about the evaluation of individuals will be crucial in Chapter 6, for the following reason: if a hyperplan selects properties of individuals, then it should be possible to learn about the properties of individuals *via* hyperplans. More specifically, if you do not know the properties of an individual, but you know (i) the standard relative to which it is evaluated and (ii) how it fares according to that standard, then you can use that information to learn about its descriptive properties.

We will see this in detail in Chapter 6, but we might as well offer the following preview: evaluative language can communicate descriptive information when speakers use evaluative sentences about individuals in contexts in which an evaluative standard is clearly shared, but descriptive information about the individuals under evaluation is not. In such contexts, a common evaluative standard can be used by speakers to share descriptive information about individuals, just as plain descriptive language can. For illustration, let us suppose that we are in such a context: they are showing three movies, *a*, *b* and *c*, and you tell me that movie *a* was fantastic. I have not seen any of them, but I know that we—you and me—would describe a movie as *fantastic* only if it had a certain property, for example if it was Korean. Then, for you to tell me that movie *a* was *fantastic* is not, or only, for you to convey to me an evaluation of *a*. Rather, it is a way for you to convey to me certain descriptive information about *a*, namely that *a* is a Korean movie. This shows that evaluative sentences have descriptive uses, which arise systematically from features of the context in which those sentences are uttered. We will spelled this out more carefully in Chapter 6.

4.3.2 Specific flavors of evaluation

In Table 4.1 we said that thin evaluative adjectives of specific flavors of evaluation occupied a sort of middle ground between the completely thin adjectives and the thick adjectives. We can now say a little more about this. By that description, we mean adjectives that do not seem to carry much descriptive content and yet they are clearly not all-purpose evaluatives like *good* and *bad*. In Chapter 1 we talked of the intuitively different flavors of evaluativity that we find in natural language: there are adjectives of (at least) moral, aesthetic, personal taste, epistemic or practical evaluation. Within each of those categories, one can distinguish thin and thick adjectives: within the moral realm, *ethical* seems to be quite thin, while *cruel* or *generous* are clearly thick. In the aesthetic domain, *beautiful* is thin while *provocative* is arguably thick. A similar contrast appears in the domain of PPTs between *tasty* and *cloying*, for instance.

Nonetheless, the thin adjectives within each specific realm of evaluation seem to carry at least *some* descriptive meaning. Otherwise, it should be possible to use them “across flavors”, just like the all-purpose *good* and *bad*. In other words, if thin adjectives of particular flavors were really thin, then just as we can call actions, sculptures or individuals *good* and *bad*, we should also be able to call any of those things *tasty*, *ethical* or *beautiful*. That is not the case, though: you cannot call an action *tasty*, for example, and you cannot call a food *ethical*.⁷ Furthermore,

⁷ Except in the derived sense in which its production is ethical. In this sense we may speak of *ethical soy*, for example. But arguably, it is the production of the food, and not the food itself, that is under (moral) evaluation.

besides there being certain restriction on the kind of things that flavor-specific thin evaluatives can be applied to, these adjectives are not interchangeable: calling an action *beautiful* and calling it *ethical* does not mean the same thing.

How can we incorporate flavor-specificity while at the same time preserving the *thinness* of these adjectives? An easy solution would be to adapt the model for *good* applied to individuals that we have just seen, and cash out flavor-specificity by introducing certain restrictions on the kind of descriptive properties that the adjective can ascribe to the individual under evaluation.

Recall our model for individual-level *good* is:

$$(4.42) \quad \text{That white Mercedes is good} = \begin{cases} \text{That white Mercedes is [...]} \\ \text{[...] is good} \end{cases}$$

The meaning of a sentence like (4.42) is factored out in a descriptive component and an evaluative component. The descriptive component is factual claim that the individual in question has certain properties. What properties are these, however, is massively underspecified by the meaning of *good*. The evaluative component, on the other hand, is an evaluation of those properties.

Our claim is that specific-flavor thin evaluatives work like *good*, except for the fact that their descriptive component is less underspecified. In order to press the contrast between the underspecificity of *good* and *beautiful*, consider the following. There is any number of things that someone could value in a car. Here is a list of properties that could stand for the descriptive component of *good*:

- That white Mercedes is environmental-friendly.
- That white Mercedes has efficient motor-performance.
- That white Mercedes is painted a flashy color.
- That white Mercedes is has a classic design.

Suppose now that instead of *good*, we used *beautiful*.

$$(4.44) \quad \text{That white Mercedes is beautiful.}$$

It does not seem like that sentence could be used to mean that the Mercedes is environmental-friendly or efficient—rather, only a set of properties like the third and the fourth item on the previous list could stand for the descriptive component of *beautiful*. Nonetheless, that still leaves a lot of room for variability: someone might use *beautiful* to mean a flashy color, or a classic design, or a classic interior; an extravagant design, etc.

In order to model this, we could still leave the descriptive component of these adjectives underspecified, while at the same time introducing certain restrictions on its specification. For instance, the descriptive component of a thin *aesthetic* adjective should intuitively be a set of properties that can be perceived by the senses, such as color, shape or sound. Or in the case of PPTs, they should be properties that can be tasted (for food-related adjectives like *tasty*) or more generally experienced (for adjectives like *fun*). In the case of moral adjectives, the relevant properties should be connected to action and decision. For instance, *ethical* could stand for *taking into consideration the interests of the relevant people* or *respecting all moral principles*,

etc. Let us write this [...]*, where the asterisk marks a flavor-specific restriction on the set of descriptive properties that someone using the adjective could be ascribing to the relevant object of evaluation:

$$(4.44) \quad \text{That white Mercedes is beautiful} = \begin{cases} \text{That white Mercedes is [...]}^* \\ \text{[...] }^* \text{ is good} \end{cases}$$

Just as in the case of individual-level *good*, this view predicts that it should be possible to use a flavor-specific thin evaluative to communicate something descriptive: if a speaker uses *beautiful* in a context in which the properties of whatever is under evaluation are not known by the audience, but by contrast the audience knows what is the accepted standard of beauty, then the audience can learn what those properties are (see e.g. dialogue (16) on Umbach 2016, p. 144).

4.3.3 Thick evaluatives

Thick terms are evaluative terms that have evaluative and descriptive (or non-evaluative) meaning; thin terms are evaluative terms that are thought to have only evaluative meaning. Examples of the first kind are words like *cruel*, *brave*, *generous* or *chaste*. Thin terms include general or all-purpose evaluative terms like *good*, *bad*, as well as the flavor-specific adjectives that we have just discussed.

The intuitive contrast lies in the fact that to call someone *cruel*, for example, is to convey something negative about that person: it is to communicate that they are bad. But not all bad things are cruel. To call someone cruel is to say something more specific, something along the lines of *wilfully causes pain or suffering to others*. That more specific content however, is not evaluative but descriptive. Therefore, thick terms like *cruel* are standardly described as having evaluative and descriptive components of their meaning, whereas thin terms like *bad* are taken to only carry an evaluative component. The evaluative component is somewhat general, while the descriptive component is more specific, hence the idea of “thickening” the evaluative component (see Foot 1958; Williams 1987, for classic references).

An important point of the debate surrounding thick terms, especially in metaethics, is whether the descriptive and evaluative component are really separable (McDowell 2002). In contemporary philosophy of language and linguistics, it is largely assumed that they are, and the debate revolves around the issue of how these two elements combine with each other. To be more precise, the debate assumes that the descriptive component provides the truth-conditions of sentences containing thick terms, and then considers the question of how the evaluative component is communicated—whether it is part of those truth conditions, a presupposition, an implicature or something else entirely (see e.g. Cepollaro and Stojanovic 2016; Kyle 2013; Väyrynen 2013). Let us note, however, that the separability thesis is relatively recent; for a long time it was thought that these adjectives had a unitary meaning. So much so that Foot (1958), for instance, argued that *rude* overcomes the ought-is gap in virtue of the fact that it expresses a concept that is descriptive and normative at once. In the aesthetic domain, one might doubt whether the separability thesis holds for adjectives like *dumpy*, *dainty* or *delicate*. After all, it seems very difficult to attain a satisfactory non-normative characterization of the properties that characterize the meaning of those adjectives. It is possible, therefore, that the

view presented here—which relies heavily on the separability of a descriptive and evaluative component—needs to be tweaked in light of these adjectives.

In order to incorporate thick evaluation into our account, all we have to do is let the meaning of these adjectives be factored out in two components, and account for their evaluative component in terms of our hyperplan semantics.

To see how this proposal could work, consider the case of *cruel*. As we saw in Chapter 1, we can factor out the meaning of a sentence predicating *cruel* of an individual in the following two components:

- (4.45) John is cruel.
- a. John wilfully causes pain or suffering to others.
 - b. John is bad in virtue of wilfully causing pain or suffering to others.

Importantly, just as we saw with individual-level *good*, we can understand the evaluative component not as the evaluation of an individual, but rather as an evaluation of the properties that make him cruel. That is, the evaluative component can be interpreted as an action-type/property evaluation:

- (4.46) Wilfully causing pain or suffering to others is bad.

If we do that, we can adopt a similar analysis to the one we gave for individual level *good* and flavor-specific thin evaluatives: the meaning of *cruel* is factored out in an individual-level ascription of descriptive properties and an evaluation of such properties.

- (4.47) John is cruel = $\begin{cases} \text{John wilfully causes pain or suffering to others.} \\ \text{Wilfully causing pain or suffering to others is bad} \end{cases}$

We can assign the following truth conditions to the evaluative component (where *PRO* is *W* stands for *PRO wilfully causes pain or suffering to others*):

- (4.48) $\llbracket \text{it is bad that } PRO \text{ is } W \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1$ iff $\{w' : \llbracket PRO \text{ is } W \rrbracket_{\langle w', h_i, a_i \rangle} = 1\} \in \mathcal{U}_{\langle h_i, a_i \rangle}[\llbracket PRO \text{ is } W \rrbracket(w_i)]$

The contrast between both individual-level *good* and flavor-specific thin evaluatives, on the one hand, and thick evaluatives, on the other, is that the descriptive properties that thick evaluatives ascribe to their objects are *not* underspecified. In other words, different people cannot by and large mean different things when they call a person *cruel*, *generous* or *nasty*. Of course, there is still a lot of wiggle room (what do you mean by *cruel*? Do you mean *generous* with time or with money? Etc.), but the descriptive meaning of those adjectives is relatively stable, or at least, comparatively much more stable than that of thin adjectives.⁸

⁸As we said in Chapter 1, there are concepts that combine descriptive and normative elements in a way that defies the model of thick evaluatives. Such is the case for some adjectives like *normal* (Bear and Knobe 2017) or so-called ‘dual-character concepts’ such as *father* or *scientist* (Del Pinal and Reuter 2017; Leslie 2015). The type of view advocated here, according to which the descriptive and evaluative components of thick adjectives are clearly separable, would have trouble accounting for these concepts (but so would most contemporary accounts of thick adjectives in the philosophy of language and linguistics literature).

Recall how we claimed that individual-level *good* and flavor-specific thin evaluatives could be used to convey an evaluation or to convey descriptive information. Since our view about thick adjectives is that they have a stable descriptive component, we predict that these adjectives always convey that descriptive component—this seems to be the case, but we will explore this matter more thoroughly in Chapter 6, where we will expand on the contrast between descriptive and evaluative uses of evaluative vocabulary.

Before moving on, let us note that the classification in Table 4.1 implies that the difference between thin and thick evaluatives is largely a matter of degree. More specifically, it is a matter of the degree of underspecificity in the descriptive component of the relevant adjective. It follows that it is not accurate to say that thin evaluatives lack description—thin evaluatives applied to individuals do carry a description, but that description is massively underspecified by the conventional meaning of such words. Moreover, this type of gradual classification permits placing thin evaluatives of specific flavor in an intermediate position between the thin and thick, which seems welcome.⁹

In this section, we have seen how to use the machinery of \mathcal{L}_V to represent the meaning of an array of evaluative adjectives, with different types of arguments and different degrees of “thickness”. As we noted, to adopt this view about the difference between thin and thick evaluation is to reject the traditional view according to which thin evaluatives simply lack a descriptive element—what we are saying here is that all evaluatives, at the very least when they take individuals as arguments, carry a descriptive component. The difference between the thin and the thick lies in the degree of underspecification of such descriptive component. As we will see in Chapter 6, our proposal makes a particular prediction about how and whether evaluative adjectives can be used to convey descriptive information. In addition to this, it is part of this proposal that the difference between the thin and thick is a matter of degree. As we noted, understanding the thin-thick distinction as a *continuum* suggests a very natural place for thin evaluatives of specific flavors, namely somewhere in the middle of that *continuum*.

4.4 Solving the Frege-Geach Problem

Once presented the basic semantic entries for *good*, *bad* and various relational constructions, and having shown how they can be applied to a range of argument types, we can come back to the Frege-Geach problem for non-factualism, and in particular, to its “sub-sentential” version, and offer a more explicit solution. Let us consider the supra- and sub-sentential incarnations of the problem in turn.

⁹This idea is present in Cepollaro and Stojanovic 2016, p. 464, who write: ‘Consider adjectives such as ‘beautiful’, ‘ugly’, or ‘evil’. When discussed within the realm of aesthetics, ‘beautiful’ and ‘ugly’ are considered to be ‘thin’. Similarly, ‘evil’, within the domain of ethics, would likely count as ‘thin’, as it means roughly the same thing as “morally bad”. At the same time, from the broader perspective of looking at evaluative expressions in general, ‘beautiful’ and ‘ugly’ are narrower than ‘good’ and ‘bad’, as they constrain the evaluation to aesthetic evaluation, and ‘evil’ is narrower than ‘bad’, as the negative evaluation that it encodes is linguistically constrained to moral evaluation.’

4.4.1 Solving the supra-sentential Frege-Geach Problem

Recall the basic outline of the traditional Frege-Geach problem for non-factualism. Once that non-factualists claim that the meaning of sentences ought to be given in terms of non-standard semantic objects, the question arises as to whether and how such non-standard semantic values can be embedded in truth-functional operators, such as logical connectives. Our answer is essentially the same as the one we sketched in §2.4 and attributed to Gibbard (2003) and others. We can be more explicit about it now.

The basic idea behind this solution is that if we adopt a moderately liberal view about logical connectives on the one hand, and we assign uniform set-theoretical semantic values to declarative sentences on the other, the problem vanishes. We need to be liberal about connectives because we need to allow them to take arguments of a semantic type *other* than classical propositions (i.e., sets of possible worlds). But as we pointed out before, this type of view is independently needed to account for connectives with non-propositional arguments, such as *and* and *or* (*She was wearing a new and expensive dress*, Partee & Rooth 1983/2002), as well as for conditionals with non-propositional consequents, such as *if I fall asleep, wake me up!* (see Charlow 2015, 5 and ff). Logical connectives are then mapped onto the usual set-theoretical operations: negation is set complementation, conjunction is set intersection, disjunction is set union and the conditional is set inclusion, but there are no restrictions as to the type of their arguments (except, of course, that they end in *t*).

On the other hand, since we are adopting sets of world-hyperplan-alternatives triplets as the uniform semantic value of declarative sentences, these can figure as the arguments of logical connectives. Let us see how this works for a case like (2.54) from §2.4:

- (2.54) a. If eating animals is bad, then eating beef is bad.
 b. Eating animals is bad.
 ∴ Eating beef is bad.

Assume, first, that the conditional operator *if... then* represents set inclusion. In a standard possible world semantics, we would represent the conditional as follows:

$$(4.49) \quad [[\text{If } p, \text{ then } q]]_{w_i} = \{w : p(w) = 1\} \subseteq \{w : q(w) = 1\}$$

Now, all we have to do is allow for the conditional operator to admit sets of world-hyperplan-alternatives triplets as its arguments:

$$(4.50) \quad [[\text{If } p, \text{ then } q]]_{\langle w_i, h_i, a_i \rangle} = \{\langle w, h, a \rangle : p(\langle w, h, a \rangle) = 1\} \subseteq \{\langle w, h, a \rangle : q(\langle w, h, a \rangle) = 1\}$$

Consider now our proposal for sentences like (2.54b). Recall that we proposed to treat these sentences as containing a hidden *PRO*. Thus, *eating animals is bad* is equivalent to:

$$(4.51) \quad \text{It is bad that } PRO \text{ eats animals.}$$

Now, the truth conditions for this sentence are the following:

$$(4.52) \quad [[\text{it is bad that } PRO \text{ eats animals}]]_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\
\{w' : [[PRO \text{ eats animals}]]_{\langle w', h_i, a_i \rangle} = 1\} \in \mathcal{U}\langle h_i, a_i[PRO \text{ eats animals}](w_i) \rangle$$

A similar story applies to *eating beef is bad*. Then, in virtue of the meaning of the conditional, what (2.54a) says is that the set of world-hyperplan-alternatives triplets in the meaning of *eating animals is bad* is a subset of the set of world-hyperplan-alternatives triplets in the meaning of *eating beef is bad*. It follows that any index $\langle w, h, a \rangle$ such that $\{w' : \llbracket PRO \text{ eats animals} \rrbracket_{\langle w', h, a \rangle} = 1\} \in \mathcal{U}\langle h, a[PRO \text{ eats animals}](w) \rangle$ is an index such that $\{w' : \llbracket PRO \text{ eats beef} \rrbracket_{\langle w', h, a \rangle} = 1\} \in \mathcal{U}\langle h, a[PRO \text{ eats beef}](w) \rangle$. In other words, if your hyperplan rejects eating animals, then it also rejects eating beef.¹⁰

4.4.2 Solving the sub-sentential Frege-Geach Problem

Let us recap the sub-sentential Frege-Problem now, and offer a solution: evaluative adjectives, non-factualists defend, are characterized by the fact that they express positive and negative practical attitudes or commitments (we can say that they are action-guiding, or they invite practical inferences, or they have motivational “oomph”, etc.). To be more precise, positive evaluatives communicate support or commendation of whatever is under evaluation, and negative communicate avoidance or rejection thereof. This was spelled out in **Definition 4**:

Definition 4 (Non-factualism about positive/negative value) *Unembedded evaluative sentences containing positive evaluative adjectives express practical attitudes of support; evaluative sentences containing negative evaluative adjectives express practical attitudes of rejection.*

As we discussed, this observation does not apply to all evaluative sentences, as many of them fail to express such outright positive and negative attitudes. Our prime example was comparatives. Consider the contrast between the following two sentences:

- (4.53) Volunteering is good. » *practical attitude for volunteering*
 (4.54) Volunteering is better than donating. » *practical attitude for/against ??*

While an utterance of (4.53) would in most contexts express a commendation or support of the action of volunteering, it does not seem like a similar thing can be said about (4.54). That sentence is compatible with despising both volunteering and donating; or commending both; or praising volunteer work while rejecting the possibility of donating money.

In fact, we claimed there to be a correlation between the capacity of an utterance of an evaluative sentence to express support and rejection and the presence of an inference to the positive form of the relevant adjective. We called this **POS-ATT**:

Observation 1 (POS-ATT) *An utterance U of an unembedded evaluative sentence S containing an evaluative adjective E expresses practical support/rejection if and only if U invites an inference to the positive form of E .*

¹⁰A problem appears here: since the alternatives to any proposition in the scope of an evaluative operator are determined partly by that proposition itself (in the sense that, e.g., the alternatives to eating animals must include eating animals), my semantics apparently fails to predict the truth of (2.54a). This is so because the alternatives to eating animals are not the same as the alternatives to eating beef. Therefore, dis-preferring eating animals relative to its alternatives is compatible with failing to dis-prefer eating beef relative to its alternatives. A possible solution would be to establish, as a principle, that the relevant set of alternatives for any complex combination of evaluative sentences is determined by the logically weakest clause in it. In (2.54a), the relevant alternative for the whole construction would be those to eating animals. This solution is somewhat *ad hoc* but it seems additionally supported by the intuitive idea that, if eating animals is bad, then eating beef is also bad *for the same reason*.

This is attested by the contrast between (4.53) and (4.54), but by many other constructions involving evaluative adjectives, as we saw in §2.5.

The problem was that accounting for the connection between the positive form of evaluative adjectives and the expression of practical support/rejection is not straightforward for standard degree semantics (e.g. Cresswell 1976; Kennedy 2007; Kennedy and McNally 2005, *a.m.o.*). Degree semantics, on its most usual construal, holds that the meaning of the positive form of any gradable adjective results from the composition of two semantic elements, the lexical meaning of the adjective and the threshold-contributing silent morpheme *POS*. The problem was that neither of those components seems to be the right source of the outright positive/negative attitudes. If we choose to assign the expression of practical support/rejection to the presence of *POS*, then one should ask oneself why are other gradable adjectives not equally action-guiding in their positive form. On the other hand, if we associate the expression of practical support/rejection with the lexical meaning of these adjectives, then the association **POS-ATT** remains a mystery (c.f., §2.5.2).

As we argued, the situation is difficult for degree semantics, but a solution follows rather straightforwardly from our semantics. Or more accurately, the problem never really arises. This is so because, in a Delineation semantics, the positive form of adjectives does not result from combining two semantic pieces; rather, it is a single piece. So the answer to the question of where the expression of support/rejection comes from is: the evaluative adjective, which in its most basic form is a simple, non-gradable predicate whose function is to express outright support/rejection.

Our proposal for (4.53), as we saw in §4.3.1, is to treat that sentence as equivalent to *it is good that PRO volunteers*, which has the following truth-conditions for (4.53):

$$(4.55) \quad \llbracket \text{it is good that } PRO \text{ volunteers} \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ \{w' : \llbracket PRO \text{ volunteers} \rrbracket_{\langle w', h_i, a_i \rangle} = 1\} \in h_i(a_i[PRO \text{ steals}](w_i))$$

In our hyperplan semantics, (4.53) is made true by indices of evaluation with the following property: that the set of alternatives determined by that index is such that, relative to that set, volunteering is preferred. That is, part of what makes an utterance of that sentence true at an index of evaluation is that the index itself determines an ‘occasion for choice’, as Gibbard calls it (2003, p. 56), and volunteering is chosen at that occasion. If (according to our semantics) by sincerely uttering (4.53), a speaker subscribes a hyperplan to choose volunteering among the alternatives determined at *this* very time and place (i.e. the alternatives of the index), then it seems reasonable for her audience to assume that the speaker will, or is disposed to, volunteer; in other words, they will assume that the speaker has the relevant practical attitude towards volunteering.

Given our Delineation framework, the pressing question is not about the positive form, but about the comparative: if the positive form of an evaluative adjective semantically expresses an outright practical attitude, why does the comparative fail express such outright practical attitudes? The answer to this question lies in the fact that the comparative introduces quantification over subsets of the alternatives determined by the index. To see this, let us consider what our analysis of (4.54) would be. In accordance with our treatment of (4.53), we treat (4.54) as equivalent to *that PRO volunteers is better than that PRO donates*:

- (4.56) $\llbracket \text{that } PRO \text{ volunteers is better than that } PRO \text{ donates} \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1$ iff $\exists a' \subseteq a_i$ s.t.
 $\llbracket \text{it is good that } PRO \text{ volunteers} \rrbracket_{\langle w_i, h_i, a' \rangle} = 1$ and
 $\llbracket \text{it is good that } PRO \text{ donates} \rrbracket_{\langle w_i, h_i, a' \rangle} = 0$

(4.54) is true at an index of evaluation i just in case one can find a smaller set of the alternatives defined at i such that, relative to that smaller set—which crucially is *not* the alternative set of the index—volunteering is chosen over donating. An utterance of (4.54) does not suggest that the speaker will volunteer (or donate) because the set of alternatives at which one would volunteer rather than donate might be indefinitely removed from the alternatives defined at the index of evaluation. For all that the speaker cares, the occasion might never arise to actually choose one over the other. This is why comparative sentences do not convey outright practical attitudes, even though those attitudes contribute semantically to the meaning of comparative sentences.

What we should say about comparatives is that they express practical attitudes that are more sophisticated than the outright practical attitudes that the positive form of evaluative adjectives convey. (4.54) does not express an outright positive or negative attitude towards any *relatum*, but it does express a practical attitude in favor of choosing volunteering over donating if the occasion arises (or conversely, an attitude against choosing donating over volunteering). We may speak of a “conditional” practical attitude.

- (4.54) Volunteering is better than donating.
 » *practical attitude for volunteering over donating*
 » *practical attitude against donating over volunteering*

Importantly, those conditional practical attitudes are compatible with having any outright positive or negative attitude towards volunteering and/or donating, *except* having a positive attitude towards donating *and* a negative attitude towards volunteering.

In turn, this suggests a revision of the previous definition of non-factualism. Instead of **Definition 4**,

Definition 4 (Non-factualism about positive/negative value) *Unembedded evaluative sentences containing positive evaluative adjectives express practical attitudes of support; evaluative sentences containing negative evaluative adjectives express practical attitudes of rejection.*

we have to withdraw the assumption that the attitudes expressed are of outright support for positive evaluatives and outright rejection for negative evaluatives, as this only applies to adjectival forms that invite an inference to the positive form of those adjectives. The resulting definition is more similar to **Definition 3**:

Definition 19 (Non-factualism about positive/negative value (2nd attempt)) *Unembedded evaluative sentences containing positive/negative evaluative adjectives express—potentially conditional—practical attitudes of support and rejection.*

Evaluatives in forms that do not invite an inference to the positive form, such as comparatives, express a conditional practical attitude; while evaluatives in forms that do invite that inference express outright attitudes of support / rejection. And as we have seen, the conditional practical attitudes are compositionally derived from the outright practical attitudes, as they should.

To sum up: by introducing quantification over subsets of alternatives in the way proposed, we achieve two things: (i) we preserve the idea that the basic meaning of an evaluative adjective in

comparative form lies in its connection to preference and action, (by compositionally deriving the comparative from the positive form) but (ii) we remove that connection to preference and action from the actual situation of speech, thereby blocking the expression of outright practical attitudes.

4.5 Relative- vs. Absolute-standard

In the two last sections of this chapter, we explore other scalar properties of evaluative adjectives. Some of these properties might not be distributed uniformly: some evaluative adjectives might be relative-standard, while some might be absolute-standard adjectives; some might have scales that intuitively have upper or lower endpoints; some might not; some might (under certain interpretations) allow for specific types of measurement, some do not. In this section, we review scalar properties in connection to the positive form (namely, whether they are relative- or absolute-standard adjectives), and in the next and last section we review the question of the type of scale that these adjectives lexicalize.

We start by considering scalar properties in connection to the positive form of these adjectives.¹¹ In line with the semantics that we have defended up to here for *good* and *bad*, we maintain that evaluative adjectives are relative-standard. However, those adjectives do not behave like paradigmatic cases of relative standard adjectives, that is, dimensional adjectives like *tall* or *wide*: for example, negative adjectives like *cruel* seem to have entailment patterns that would suggest that they are absolute-standard; and evaluative adjectives across the board seem to have a different behavior with respect to comparison classes than paradigmatic relative-standard adjectives. Nonetheless, we argue that the hypothesis that evaluative adjectives are relative-standard is better supported by available data. We will end this section by challenging the connection between the absolute/relative distinction and the presence of scale bounds.

The distinction between RELATIVE- (REL) and ABSOLUTE-STANDARD (ABS) gradable adjectives can be characterized in terms of the property that the positive form of a gradable adjective predicates of its argument: the positive form of REL-adjectives predicates *that the argument's degree of the relevant property surpasses a certain threshold*, which is an intermediate point on the appropriate scale whose precise value is determined contextually. *Tall* is a typical example of a RES adjective: to be tall amounts to possessing a degree of height that exceeds a contextually determined threshold, which is an intermediate point on the height scale.

ABS-adjectives are those whose positive form predicates *that their argument's degree of the relevant property is a maximal or minimal degree on the corresponding scale, that is, a scale endpoint*. ABS-adjectives are called *maximum* (ABS_{max}) or *minimum* (ABS_{min}) standard if the endpoint is the upper or lower scale endpoint, respectively. *Full* is an ABS_{max} adjective: to be full is to have a maximum degree of the property of fullness (namely, to have as much content as capacity); *dirty* is an ABS_{min} adjective, since to be dirty is to possess a minimal degree of dirtiness.

In addition to this, gradable adjectives are of positive or negative polarity, depending on whether modification by *-er* (or *more*) denotes a higher degree on the relevant scale. Thus, *tall* is

¹¹This section follows largely section 2 of Faroldi and Soria Ruiz 2017. The observations in that paper are restricted to moral adjectives, but they can be extended to evaluative adjectives more generally.

positive, while *short* is negative; *full* is positive while *empty* is negative; and *dirty* is positive while *clean* is negative (Kennedy and McNally 2005, a.m.o.).

Interestingly, common tests for this distinction used in the literature do not give stable results when applied to evaluative adjectives. In particular, while certain entailment patterns suggest that these adjectives are REL-standard, their behavior with respect to prepositional phrases denoting comparison classes as well as their admissibility of certain modifiers suggests that they are ABS-standard (see Liao, McNally, et al. 2016; Liao and Meskin 2015 for experimental data about aesthetic adjectives; and Stojanovic 2018 for criticism of those studies). Finally, some of these tests give different results for positive and negative adjectives, so we will consider adjectives of both polarity. We start by considering various data points in favor of each hypothesis, and we discuss them at the end of this section.

4.5.1 Entailment patterns

We'll start by looking at three tests according to which evaluative adjectives come out as REL-standard (see Faroldi and Soria Ruiz 2017 for these tests applied to moral adjectives; the tests themselves come from Kennedy and McNally 2005, *a. m. o*). First, given that ABS-adjectives in the positive form denote endpoints on a scale, the following entailment patterns hold of them:

- (4.57) a. If x is ABS_{max} , then x could not be ABS_{max} -er.
 b. If x is not ABS_{min} , then x does not possess any degree of ABS_{min} -ness.

To test these entailment patterns, we consider whether sentences that do not respect them are acceptable. That is, we consider whether the conjunction of the antecedent clause and the denial of the consequent are coherent. If not, then the pattern holds for that adjective and the adjective is ABS. Consider the following two cases with *full* (ABS_{max}) and *open* (ABS_{min}), and compare them with *tall* (which is *rel*):

- (4.58) a. # The glass is full and it could be fuller.
 b. # The door is not open and it is ajar.
 c. Natalia is tall and she could be taller.
 d. Matheus is not tall and he possesses some degree of height.

As these examples show, the relevant constructions are not acceptable when we use ABS-adjectives, thereby suggesting that these adjectives respect those entailment patterns. When we use REL-adjectives however, those constructions are acceptable, which suggests that REL-adjectives do not have those entailments.

Positive (*generous*, *virtuous*) and negative (*despicable*, *cruel*) evaluative adjectives pattern like REL-adjectives in this test:

- (4.59) a. What she did was generous and it could have been more generous.
 b. Being vegetarian is not virtuous and it possesses some degree of virtuousness.
 c. What he did was despicable and it could have been more despicable.
 d. ?? Eating animals is not a cruel thing to do and it involves some degree of cruelty.

These constructions are for the most part acceptable, which suggests that the positive form of these adjectives does not predicate that their object has a maximal or minimal degree of the relevant property.¹²

Secondly, ABS-adjectives also satisfy the following entailment pattern:

- (4.60) If x is more ABS_{min} than y , then x is ABS_{min} .
 If x is more ABS_{max} than y , then y is not ABS_{max} .

For illustration, compare *full* and *dirty* to *tall*:

- (4.61) a. # The shirt is dirtier than the jacket and the shirt is not dirty.
 b. # The glass is fuller than the vase and the vase is full.
 c. Natalia is taller than Matheus and she is not tall.
 d. Natalia is taller than Matheus and Matheus is tall.

With respect to these two patterns, positive and negative evaluative adjectives give slightly different results: positive adjectives do not show these patterns, hence they behave like REL-adjectives. Negative adjectives show the second pattern, although our judgments about them are not very clear:

- (4.62) a. Animal testing for scientific purposes is more ethical than for cosmetic purposes, and it is not ethical.
 b. Volunteering for a charity is more virtuous than donating to a charity, and donating is (also) virtuous.
 c. ?? Hiding your office mate's keys is more cruel than eating their snacks, and it is not cruel.
 d. Accepting bribes is more despicable than paying them, and paying them is (also) despicable.

Using a negative evaluative adjective like *cruel* in a comparison does seem to suggest that both terms of comparison invite an inference to the positive form. However, we are skeptical that this particular pattern of inference is an entailment. The fact that a sentence like (4.62c), even if marked, is not completely out, suggests that the inference is to some extent cancellable, so it might well arise due to implicature or some other mechanism. In general, it bears mentioning that the acceptability of all these constructions only supports the conclusion that the relevant inferential patterns are not *entailments* of those adjectives. But those inferences could arise *via* implicature or other mechanisms.

4.5.2 Comparison class sensitivity

The previous observations strongly suggest that evaluative adjectives are REL-standard. By contrast, the behavior of evaluative adjectives with respect to their thresholds' sensitivity to

¹²Admittedly, the last of those, (4.59d), is a little off. This might be due to the fact that the scale invoked, a scale of cruelty, is a negative value scale and therefore marked (compare: *Ann is not short and she possesses some degree of {#shortness} / {height}*). If height is the property that *tall/short* lexicalizes, what is the equivalent for *cruel*? Moreover, what is the antonym of *cruel*? A salient possibility would be simply negative value, or *bad*, but that is too weak—eating animals might possess some degree of badness that has nothing to do with cruelty.

comparison classes seems to suggest that they are not.

As we mentioned, the thresholds for REL-adjectives are determined in context. Additionally, the value of a threshold can be shifted by explicit reference to a comparison class: *tall for a basketball player* and *tall for a 5 year old* establish different thresholds for the positive form of the REL-adjective *tall*. Furthermore, invoking one or other comparison class does not, in general, suggest an inference to the positive form of the relevant adjective:

(4.63) Ann is tall for a 5 year old. ??~ Ann is [not] tall.

(4.64) Ann is tall for a basketball player. ??~ Ann is [not] tall.

ABS-adjectives, by contrast, do not show such sensitivity to comparison classes: modification by a comparison class forces the interpretation that the positive form of the adjective does not apply after all: the prepositional phrase *for a TV antenna* in *straight for a TV antenna* does not shift the threshold of straightness (which remains the maximum degree of straightness). Rather, it is most naturally taken to simply mean *not straight*.

(4.65) That rod is straight for a TV antenna. ~ That rod is not straight.

Liao, McNally, et al. 2016 observe that aesthetic adjectives pattern in this respect like ABS-adjectives: rather than allowing a shift in the positive form threshold, mentioning a comparison class suggests that the bare positive form does not apply (they also note that such constructions—namely [aesthetic adjective + for comparison class]—appear very rarely in *corpora*). Here are two examples that they mention:

(4.66) She is beautiful for an older woman. ~ She is not beautiful.¹³

(4.67) The hotel was elegant for a Best Western. ~ The hotel was not elegant.

Mentioning the comparison class in both cases implies that the bare positive form does not apply. The same holds of evaluative adjectives across the board: explicit reference to comparison classes does not seem to be able to shift the threshold for the positive form. Rather, it suggests that the bare positive form does not apply:

(4.68) Giving alms is a generous act for a miser. ~ Giving alms is not generous.

(4.69) What they did was not unethical for a war criminal. ~ What they did was unethical.

Interestingly, these data show only half of the picture, because these examples purposefully invoke comparison classes that are clearly on the *negative* end of the relevant scales—Best Westerns are conventionally taken to not be elegant, war criminals are taken to not be moral, etc. If we swap those comparison classes for classes of object on the positive end of the same scales, it is interesting to note that the relevant sentences become difficult to interpret and it becomes very unclear what—if any—inference remains:

(4.70) ?? The hotel was elegant for a luxurious Parisian XIXth century hotel.

¹³Liao, McNally, et al. 2016's original example is: *Anyone who calls someone 'beautiful for an older woman' does not get my love.*

- (4.71) ?? What they did was not unethical for a reputed NGO focused on human rights violations.

This piece of data suggests that evaluative adjectives behave more similarly to ABS- than to REL-standard adjectives. We come back to this observation in the discussion section.

4.5.3 Scale bounds

Lastly, the scales of gradable adjectives can be open on either, neither or both ends; and the acceptability of modifiers like *slightly* and *perfectly* (which pick out minimal/maximal scale endpoints) with positive and negative adjectives has been taken to reveal information about the type of scale lexicalized by a given pair of adjectives (Kennedy and McNally 2005). For example, the pair {*dirty*, *clean*} lexicalizes an upper-closed scale, as shown by the following pattern:

- (4.72) a. perfectly / ?? slightly clean
b. ?? perfectly / slightly dirty

The reason why these antonyms pattern differently with respect to those modifiers is that the scale of dirtiness/cleanliness is open on one end, and closed on the other. Intuitively, things can be ever dirtier, but there exists a maximum level of cleanliness. So the reason why *perfectly clean* is acceptable while *slightly clean* is not is that the positive form *clean* refers to the maximum degree on the scale of cleanliness, which is the top of that scale. By contrast, the reason why *slightly dirty* is acceptable while *perfectly dirty* is not is that the positive form *dirty* refers to some non-zero degree on the scale of cleanliness.

Interestingly, evaluative antonyms show the same pattern.¹⁴

- | | |
|---|---|
| (4.73) a. perfectly / ?? slightly good | (4.75) a. perfectly / ?? slightly beautiful |
| b. ?? perfectly / slightly bad | b. ?? perfectly / slightly ugly |
| (4.74) a. perfectly / ?? slightly ethical | (4.76) a. perfectly / ?? slightly fun ¹⁵ |
| b. ?? perfectly / slightly unethical | b. ?? perfectly / slightly boring |

Kennedy (2007) proposes the generalization that adjectives with totally open scales are REL-standard, while adjectives that lexicalize scales closed on one or two ends are ABS-standard. If

¹⁴Introspective intuitions might not be super reliable here, but a quick search on the [iWeb corpus](#) shows a strong contrast in number of hits for each of these constructions. Nonetheless, a potential confound is the choice of the particular modifiers *perfectly* and *slightly*, which might be more preferable with some adjectives than with other. A full assessment of this pattern of acceptability would require looking at other modifiers that arguably have the same meaning as *perfectly* and *slightly*, such as, e.g. *completely* or *somewhat*.

¹⁵Note that *a little / slightly fun* can be acceptable, as in *the show was a little / slightly fun*. To my ear though, that sentence requires emphasis on the modifier, which is a hint that the construction is not perfectly natural in itself. Moreover, a salient interpretation is that the show is (sufficiently, standardly) fun *but only some times*; not that its overall degree of fun is slight. That time might be playing some role can be further supported by noting the contrast between the following sentences, where *slightly beautiful* is applied to an object with and without temporal parts:

- (v) ?? The picture is slightly beautiful.
(vi) The movie is slightly beautiful.

this generalization is correct, these data count in favor of the hypothesis that evaluative adjectives are ABS-standard.

To summarize: entailment patterns suggest that evaluative adjectives are REL-standard; while lack of sensitivity to comparison classes and the fact that pairs of evaluative adjectives admit endpoint modifiers suggests that they are ABS-standard.

4.5.4 Discussion

More tests can and ought to be carried out, but we venture that the hypothesis that evaluative adjectives are REL-standard is better supported by the data (the observations about entailment patterns are particularly telling), whereas tests that suggest that evaluative adjectives are ABS-standard are based on observations that can be explained alternatively: first, the presumed insensitivity of evaluative adjectives to comparison classes *per se* does not show that evaluative adjectives are ABS-standard, but simply that their thresholds' are rigid in a way that the thresholds of other REL-standard adjectives are not. Why is this so? We may rehearse the following answer: whether an object falls under a certain evaluative concept, say whether an action-type α is despicable, depends on the pre-existing values of the person who is considering that question. Crucially however, it does not depend on the actions that we compare α to. Thus, it is to be expected that explicitly invoking a comparison class does not shift the threshold for the positive form. Immunity to comparison classes is secured in our hyperplan semantics, insofar as comparison classes play no role in the semantics of evaluative adjectives.

Secondly, the association of open scales to REL-adjectives and of partially closed scales to ABS-adjectives is a generalization that has been challenged. In particular, Lassiter 2017 challenges the claim that REL-adjectives are associated with an open scale: the epistemic modal adjective *likely*, for instance, is a REL-standard adjective with a closed scale on both ends. Evaluative adjectives could be an exception to the second generalization.

Thirdly, and regardless of whether we want to challenge the second generalization, the acceptability pattern of modifiers like *perfectly/slightly* can receive at least three alternative explanations which are compatible with evaluative adjectives being REL-standard. First, we might say that the pair *slightly, perfectly* quantifies over parts of the relevant objects, and not over the degrees of the relevant property. So we could say that *the painting is perfectly beautiful* means *every part of the painting is beautiful*; while *the painting is slightly ugly* would mean *some part of the painting is ugly*. However, that by itself does not yet explain why the acceptability of *perfectly/slightly* correlates with the choice of antonym. But it might be the case that evaluatives of positive and negative valence behave differently with respect to their objects and their parts: it could be that a beautiful object requires all its parts to be sufficiently beautiful, whereas a single sufficiently ugly part of an object makes the whole thing ugly. Crucially, this is compatible with those adjectives being REL-standard.¹⁶

A second explanation for the acceptability of *perfectly/slightly*, at least for moral adjectives, could be that they quantify over *circumstances*. *Perfectly ethical* could mean *ethical in all circumstances*; while *slightly unethical* might mean *unethical in some circumstances*. Again, this is perfectly compatible with (un)ethical being a REL-standard adjective.

¹⁶Plural arguments might be a window into this phenomenon, c.f. Maldonado et al. 2019.

A final explanation for the *perfectly/slightly* acceptability pattern could be the putative MULTIDIMENSIONALITY of evaluative adjectives. Multidimensional adjectives (Sassoon 2013a, 2016) are adjectives that integrate different respects or dimensions. The pair $\langle \textit{healthy}, \textit{sick} \rangle$ is a paradigmatic example: one can be healthy or sick with respect to various dimensions, such as blood pressure, cholesterol or blood sugar level. By contrast, there is but one dimension associated with an adjective like *tall* (i.e. height). Following Sassoon, the main available linguistic tests for multi-dimensionality are the admissibility of “dimension-accessing” operators and modifiers, such as the PPs *with respect to ...* and *in some/most/every respect(s)*. While evaluative adjectives do not pass these tests with merits, they do not fare nearly as badly as dimensional adjectives (as we will discuss in the following Chapter). Therefore, there is some independent reason to think that evaluative adjectives are multi-dimensional, or at least that their extension somehow depends on different aspects of dimensions.

With respect to our present concerns, the important thing to note is that multidimensional adjectives admit maximum/minimum modifiers like *perfectly/slightly*, and yet those adjectives are not ABS-standard. Combined with multidimensional adjectives, those modifiers do not reference endpoints on a scale, but rather, they quantify over the dimensions associated with the relevant adjectives: *perfectly healthy* does not mean having a maximum degree of health, but rather being sufficiently healthy in all respects; and *slightly sick* does not mean having a non-zero degree of sickness, but rather means being sufficiently sick in some respect (see Sassoon 2016, p. 10). Thus, that is yet another way that the adverbial phrases in (4.73)-(4.76) could be interpreted: *perfectly ethical* means ethical in every respect; and *slightly unethical* means unethical in some respect. These interpretations are consistent with the claim that such adjectives are REL-standard.

Before we move on however, one may wonder: why venture the generalization that *all* evaluative adjectives are REL- or ABS-standard? Why not entertain the possibility that, for instance, some evaluative adjectives are ABS? There is an empirical and a theoretical reason for this: on empirical grounds, we simply have yet to come across an evaluative adjective that very clearly patterns like standard ABS-adjectives. Our choice of examples is not selective; we have picked evaluative adjectives at random and observe that they reveal a—relatively stable—semantic pattern. On theoretical grounds, we consider it a plausible hypothesis that all evaluative adjectives have similar semantic properties due to their sharing certain mathematical or structural properties inherent to our notion of value (and which we aim our hyperplan semantics to broadly capture). For instance, it is plausible that orderings of objects according to their value are only partial. Thus, one might reasonably expect that all evaluative adjectives give rise to partial orderings as well (thereby allowing for incomparabilities). However, we think that if the foregoing hypothesis is true, that would constrain the mathematical properties of the scales of these adjectives, but not necessarily the properties of their thresholds. In other words, that evaluative adjectives share certain scalar properties is compatible with them being REL or ABS. So we take there to be theoretical reasons to expect some uniformity at the level of the internal scale of these adjectives, and empirical—but of course defeasible—reasons to expect uniformity at the level of thresholds for the positive form. We turn now to discussing those internal scalar properties.

4.6 Scale structure

In his 2017, Lassiter uses Representational Measurement Theory (RMT, see Krantz et al. 1971) to explore the features of a set of linguistically gradable items, mainly epistemic, probability and deontic modals, but also the evaluative adjective *good*. In this section, we present the typology of scales that he discusses as well as the relevant linguistic tests that can help diagnose the scale type of a scalar item, and we will see what best applies to evaluative adjectives and whether our hyperplan semantics can accommodate the relevant observations.

In RMT, the properties of scales are studied by considering what mathematical operations they support. The outcome of this is a typology of scales, or a set of scale types. Lassiter proceeds by attempting to subsume the scales lexicalized in various natural language expressions under scale types defined by RMT. His procedure is roughly the following: he starts from the observation that some predicates are gradable, and assumes that they denote scalar properties, or SCALES for short. Then, in order to study the properties of those scales, he does two things. The first is to observe what kind of inferences and modifiers those natural language items allow and forbid. The second is to map the various acceptable uses of those scalar items onto different mathematical relations over the real numbers, in the way that RMT tells us to. Depending on the kind of mapping from natural language onto such mathematical relations that are admissible, a scale can be subsumed under one or other scale type.

For concreteness, let us define a SCALE as a tuple $\mathcal{S} = \langle X, \geq, \dots \rangle$ containing a set of individuals X , a binary ordering relation \geq and potentially other operations. In order to determine the features of \mathcal{S} , we seek to define a structure-preserving mapping (a homomorphism) μ from \mathcal{S} onto $\langle \mathbb{R}, \geq, \dots \rangle$ (where \mathbb{R} is the set of real numbers, \geq is the usual ordering relation and other operations over \mathbb{R} might be taken into account). If a function μ is a homomorphism from \mathcal{S} onto $\langle \mathbb{R}, \geq, \dots \rangle$, then we call μ an ADMISSIBLE MEASURE FUNCTION of \mathcal{S} . And to prove that μ is an admissible measure function of \mathcal{S} is to prove a REPRESENTATION THEOREM. Different scale types are then distinguished by imposing different representation theorems that the admissible measure functions must satisfy; the more conditions they must meet; the more structure the scale has.

There is a potentially infinite number of relations that we can define over a scale \mathcal{S} . But the crucial ones for our purposes are the binary ordering relation \geq , which we have already mentioned, and the operation of CONCATENATION, (which we represent as \circ). Concatenation allows us to construct compound objects from the simple elements in a given domain. For any elements a, b in some domain, $a \circ b$ is the concatenation of a and b .

However, concatenation is not a linguistic operation. In order to represent concatenation in natural language, it has to be mapped onto some model-theoretical relation. Lassiter (2017, p. 39), following Krifka (1989), maps it to the set-theoretical operation of join, \sqcup , restricted to non-overlapping individuals: $x \circ y$ is defined if and only if

1. x and y belong to the same semantic type α , and
2. x and y do not overlap.

When defined, $x \circ y = x \sqcup y$, where \sqcup is join over domain D_α .

As we will see, since we will stick to propositional objects of evaluation, and given that the join operation over the domain of propositions is just set union, we will understand the concatenation of propositions as their union: for any non-overlapping propositions u, v , $u \circ v = u \cup v$,

which we represent linguistically as Boolean disjunction. Concatenation is crucial because mapping it to different mathematical relations helps to define different scale types (in particular ratio and interval scales).

Lassiter focuses on the three main type of scales used in social and empirical sciences, namely ORDINAL, INTERVAL and RATIO scales. These scales are defined, as we said, *via* their admissible measure functions. In order to investigate what scale a given lexical item has, we need to consider adjectival modifiers. Adjectival modifiers impose restrictions on the scales of the adjectives that they modify. The adjectival modifiers that can offer insight into scale structure include measure phrases (*two meters*, *35 years*), quantity adverbs (*much*, *a little*, *a lot*), ratio (*twice*, *6.2x*) and proportional modifiers (*half*, *3/4*, *95%*). For example, consider the contrast between the following sentences:

- (4.77) a. The car is 2 years old.
 b. # The car is 2 years/meters/horsepower/... good.

The modifiers in (4.77) provide information about the scales that correspond to each adjective. The acceptability of (4.77a) suggests that the scale of *old* accepts units of measure (in particular temporal units), while the unacceptability of (4.77b) suggests that the scale of *good* has no acceptable unit of measurement (as discussed in §2.5).

We turn now to presenting these three scale types and to consider whether the scale lexicalized by the pair *good/bad* satisfies each of them. We conclude that the scale of *good/bad* is definitely stronger than an ordinal scale, yet weaker than a ratio scale. However, we remain unsure as to whether an interval scale is suitable, as it seems to impose too strong a structure. Subsequently, we aim to show that our hyperplan semantics offers a structure that is compatible with these observations about the scale of *good/bad*.

4.6.1 Ordinal

A scale \mathcal{S} that is merely ordinal is such that, all that can be said of elements in its domain is how they are ordered with respect to each other. For this reason, all measure functions μ that preserve the ordering among the reals are admissible measure function of \mathcal{S} . No other structure is represented; we do not know anything about the *distances* between elements on the scale, for instance, or their respective distances to a zero point. The relation of set inclusion is an example of a relation with a merely ordinal structure: all the information that set inclusion represents is an ordering on its domain. More precisely:

Theorem 1 (Ordinal scale) *If a scalar property $\mathcal{S} = \langle X, \geq \rangle$ is an ORDINAL scale (disregarding concatenation and other operations), then every admissible measure function μ that maps \mathcal{S} onto $\langle \mathbb{R}, \geq \rangle$ is such that, for all $a, b \in X$ and $a \geq b$, $\mu(a) \geq \mu(b)$.*

Any measure function μ is an admissible measure function of \mathcal{S} as long as, to any two elements a, b of X that stand in the \geq relation of \mathcal{S} , μ assigns numerical values such that the value of a is a number at least as great as b (note that the theorem says nothing about the type of mathematical operation that concatenation should be mapped onto. Thus, any mathematical operation is admissible; it could be addition, subtraction or what have you).

To see how this works, consider again the set inclusion relation. Let us represent it as a structure $\langle \mathcal{P}(X), \supseteq \rangle$, where $\mathcal{P}(X)$ is the power set of some set X , and \supseteq is the superset relation. If this

structure is ordinal, then every measure function that maps it onto $\langle \mathbb{R}, \geq \rangle$ should satisfy the representation theorem above. This implies that any mapping that respects the ordering among reals will be an admissible representation of the superset relation. For any two elements Y and Z of $\mathcal{P}(X)$ such that $Y \supseteq Z$, a μ such that $\mu(Y) = 5$ and $\mu(Z) = 3$ is an admissible measure function; another μ' such that $\mu'(Y) = 12,351$ and $\mu'(Z) = -0.0004$ also is; but a μ'' such that $\mu''(Y) = 2$ and $\mu''(Z) = 3$ will not do, because $2 \not\geq 3$. The crucial feature of ordinal scales is that nothing matters beyond order; which is why, if we define scales by its admissible measure functions, ordinal scales are very liberal.

Might the scale of *good/bad* be merely an ordinal scale? Lassiter's answer (Lassiter 2017, p. 177), with which we agree, is negative: the scale of *good/bad* cannot be merely ordinal, because in addition to order, the *distance* between elements also matters. The crucial data point here is the admissibility of quantificational adverbs such as *much*, *a little* or *a lot*.

Note that there is an interpretative difference between the following two sentences:

(4.78) Volunteering is better than donating.

(4.79) Volunteering is much better than donating.

However vague and variable the meaning of *much* may be, the fact that we can imagine a situation in which (4.78) would be true while (4.79) is false suggests that they do not mean the same thing.

(4.80) Volunteering is better than donating, but not much better.

Informally, the admissibility of such modifiers imposes the requirement on the scale of *good/bad* that the distance between measures be meaningful: (4.80) says that the degree of value of volunteering is higher than that of donating, but that the distance between the value of volunteering and the value of donating is not “much”. If the scale of *better* were merely ordinal, then it would not be possible to represent distances between degrees. But that sentence does represent something about the distance between volunteering and donating, and therefore the scale of *better* cannot be just ordinal.

More formally, the reasoning for this is the following: if *good/bad* had an ordinal scale, then for any two elements on that scale that are ordered with respect to each other, all measure functions that respect that ordering should be admissible. A sentence like (4.80) however, admits certain order-preserving measure functions but also rules out others, namely those that assign a value to each element that is at least as great as whatever quantity *much* stands for:

$$(4.81) \quad = \mu(V) > \mu(D) \ \& \ [\mu(V) - \mu(D)] \geq \mathbf{much}$$

The fact that the truth-conditions of (4.79) require ruling out certain order-preserving measure functions suggests that the scale of *good* must have more structure than that of an ordinal scale. In other words, regardless of how one defines *much*, there will be order-preserving measure functions for which the relation in (4.79) fails to hold—just think of any measure function assigning *some* difference in degree to volunteering and donating, although not *much*.

Based on these observations, Lassiter concludes that *good/bad* must have a stronger scale than ordinal (the reader can check that similar observations apply to all evaluative adjectives, as they are can all be modified by quantificational adverbs such as *much*, *a little* or *a lot*). The other two

salient options are interval and a ratio scales, in order of increasing strength. Lassiter eventually settles on an interval scale, but let us first consider his arguments against a ratio scale.

4.6.2 Ratio

Ratio scales are characterized by the fact that the relative “size” of elements does matter. In particular, difference in size between elements is measured in ratios, which means that only ordering-preserving measure functions that are obtained *via* a multiplication operation are admissible. In addition to this, ratio scales demand that concatenation be mapped onto the mathematical operation of addition. That is, the concatenation of two elements may only be mapped onto a measure function that assigns to such compound object the *sum* of the individual measures of the concatenated elements.

Scales like height and weight are familiar examples of ratio scales, where the relation between elements in the scale can be mapped onto measure functions that maintain a constant ratio between the numerical values assigned to them. More formally:

Theorem 2 (Ratio Scale) *If a scalar property $\mathcal{S} = \langle X, \geq, \circ \rangle$ is a RATIO scale, then the following representation theorem holds for every admissible measure function μ that maps \mathcal{S} onto $\langle \mathbb{R}, \geq, + \rangle$: for all $a, b \in \mathcal{S}$ and $a \geq b$, (i) $\mu(a) \geq \mu(b)$, (ii) $\mu(a \circ b) = \mu(a) + \mu(b)$ and (iii) For any μ' satisfying (1) and (2) and for any $x \in X$, there's an $n \in \mathbb{R}^+$ s.t. $\mu'(x) = n\mu(x)$.*

Note that an ordinal scale satisfies only the first of those conditions. A ratio scale imposes more conditions on the admissible measure function and is therefore stronger; or alternatively, reduces the size of the set of admissible measure functions.

In order to see how ratio scales constrain admissible measure functions, let us consider a familiar example, such as height. Seeing why the height scale \mathcal{S}_{height} is stronger than an ordinal scale is easy: suppose that Amir is taller than Mora. If height were an ordinal scale, we should be able to map Amir and Mora's heights to any pair of numerical values under the $>$ relation. But some of those values would radically misrepresent their heights. Suppose that Amir and Mora are respectively 182 and 165 centimeters tall. Consider a measure function μ' that assigns $\mu'(\text{Amir}) = 182$, $\mu'(\text{Mora}) = 165$, but such that their concatenated heights, $\mu'(\text{Amir} \circ \text{Mora})$, is equal to 17. μ' respects the ordering relation between them—i.e., complies with condition (i) in Theorem 2; but radically misrepresents the intuitive value of their concatenated heights. Or consider another measure function μ'' that assigns $\mu''(\text{Amir}) = 182$, $\mu''(\text{Mora}) = 181.9$ and $\mu''(\text{Amir} \circ \text{Mora}) = 363.9$. This measure function respects the ordering relation between Amir and Mora—complying with (i)—and the fact that their combined heights should be the arithmetical sum of their individual heights—complying with condition (ii). But it does not respect the intuitive relation that holds between Amir and Mora's heights, because it does not preserve the ratio between their heights. That is, does not respect condition (iii) in Theorem 2.

Conditions (ii) and (iii) in Theorem 2 impose more structure on the admissible measure functions for a ratio scale than mere preservation of order, and thereby rule out measure functions like μ' and μ'' . In particular, condition (ii) requires that only measure functions that assign to any concatenation of individuals the sum of the measures of their elements are acceptable, ruling out μ' . And condition (iii) demands that only measure functions that can be obtained by multiplying by the same positive real the values assigned to a pair of elements by some order-preserving (and additive with respect to concatenation) measure function are admissible. This

rules out measure functions such as μ'' . In other words, only *order*-, *sum*-, and *ratio*-preserving measure functions are admissible.

Going back to our example, if Amir measures 182 and Mora measures 165 (say, centimeters); then only measure functions that (a) preserve the order between the heights of Amir and Mora, (b) assign to their combined heights the sum of their individual heights, and (c) preserve the ratio 182 : 165, namely 1.1030, between the individual measures of Amir and Mora, are admissible. $182 + 165 \neq 17$, and therefore μ' is not an admissible measure function. And $182 : 181.9 \neq 1.1030$, and therefore μ'' is not admissible either.

Is the goodness scale a ratio scale? We have seen that it cannot be an ordinal scale, but Lassiter gives two arguments to support that it isn't a ratio scale either. First, since ratio scales make ratio comparisons interpretable, then it is at least *compatible* with sentences that admit ratio modifiers such as *1.1030x* that the relevant properties have a ratio scale. For example, *tall*, which we (independently) know to have a ratio scale, admits these modifiers:

(4.82) Amir is 1.1030x as tall as Mora.

By contrast, such ratio modifiers are not acceptable with *good/bad*:

(4.83) ?? Volunteering is 1.1030x as good as donating.

Note, again, that the acceptability of ratio modifiers does not prove that the relevant scales are ratio. Having a ratio scale is a sufficient condition for these modifiers to be acceptable—because a ratio scale makes ratio comparisons interpretable—but not the other way around, since ratio modifiers are also compatible with stronger scale types than ratio, as those stronger scale types will satisfy all conditions for a ratio scale as well as some more. Similarly, the fact that a sentence type featuring a ratio modifier is *not* acceptable is not enough to conclude that the corresponding scale is not ratio. It could be, for instance, that this type of modifier is ruled out because the *good/bad* scale is ordinal. However, since we have strong reasons to rule out an ordinal scale, this unacceptability pattern for ratio modifiers offers some reason to think that the *good/bad* scale might not be ratio.

Secondly, ratio scales are by definition POSITIVE with respect to concatenation, and the scale of *good/bad* is not, according to Lassiter. Being positive with respect to concatenation means that the concatenation of any two elements has a greater amount of the relevant property than either element. More formally, a scale $\mathcal{S} = \langle X, \geq, \circ \rangle$ is positive with respect to concatenation iff for any $a, b \in X$ that do not overlap, $a \circ b \geq a$ (except if \mathcal{S} is lower-bounded, and b has exactly the value of the lower-bound). Lassiter holds that the *good* scale lacks this property, based on the observation that it seems to respect the following inference pattern:

(4.84) a. $a \geq b$
 b. $a \geq c$
 $\therefore a \geq (b \circ c)$

If \mathcal{S} were positive with respect to concatenation, that inference should fail in many instances. But it does not fail for *good* (by contrast, it very clearly fails for *likely*, which is independently argued to have a ratio scale, see Lassiter 2017, p. 179). For an example, consider the following, intuitively valid inference (recall that concatenation for propositions is disjunction):

- (4.85) a. Volunteering is at least as good as donating.
 b. Volunteering is at least as good as subscribing to the mailing list.
 \therefore Volunteering is at least as good as donating or subscribing to the mailing list.

According to Lassiter, that this inference pattern is in general valid shows that the scale of *good/bad* has to be weaker than a ratio scale. We are thus left with the possibility of an interval scale.

4.6.3 Interval

Interval scales are stronger than ordinal scales, but weaker than ratio scales. They are stronger than ordinal scales because over and above mere order, the difference between elements on the scale, that is, their intervals, matters. However, they are weaker than ratio scales, because they do not determine a zero point, and therefore the positions of elements on the scale cannot be defined using ratios. Interval scales take into account the distance, or gaps, between elements—for this reason, the elements on an interval scale are not actually points, but intervals (although this will not matter for our purposes).

Temperature, clock time or danger are familiar examples of interval scales offered by Lassiter. Informally, what is crucial about those natural language cases is that the scales that those expressions lexicalize do not determine a zero point: a “zero” degree of temperature is a mere convention, and changes when we move from the Celsius to the Fahrenheit scale, for instance; similarly, it is intuitive to think that there is no zero point in clock time or in a scale of danger/safety.

Formally, this is cashed out by noting that ratio transformations are only meaningful relative to some reference point.

Theorem 3 (Interval Scale) *Where $S = \langle X, \geq, \circ \rangle$ is a scale, if S is an interval scale, then the following representation theorem holds for every admissible measure function μ that maps S onto $\langle \mathbb{R}, \geq, + \rangle$: for all $a, b \in S$ and $a \geq b$, (i) $\mu(a) \geq \mu(b)$, (ii) $\mu(a \circ b) = \mu(a) + \mu(b)$ and (iii) for any μ' satisfying conditions (i) and (ii) and for any $x \in X$, there are some a, b such that $a \in \mathbb{R}^+$ and $b \in \mathbb{R}$, $\mu'(x) = a\mu(x) + b$.*

That interval scales are stronger than ordinal scales is easily seen by considering that, just like ratio scales, they too restrict the set of admissible measure functions.

The linguistic counterpart of having an interval scale is that a possible explanation appears for why ratio modifiers are sometimes out. Consider the following examples:

- (4.86) ?? This bowl is 1.38x hotter than that one.
 (4.87) ?? Amir came 2x later than Mora.
 (4.88) ?? My neighbourhood is 4x safer than yours.

Such ratio comparisons are meaningless unless arbitrary zero points are defined on the relevant scales, and even then they are only acceptable with the appropriate qualification. So an attractive explanation for why (4.86) is odd becomes available: note that we measure temperature by scales such as Celsius, Fahrenheit or Kelvin. Now, (4.86) might be true in a certain scale (say, Celsius). But if we change to a Fahrenheit scale, the ratio 1.38 will fall apart because

the conversion between Celsius and Fahrenheit does not preserve ratios. For instance, if this bowl is 45°C and that bowl is 62.1°C, then we could say that this bowl is 1.38x hotter *on the Celsius scale*. But in a Fahrenheit scale, those temperatures are 113 and 143.78 respectively, and the ratio is no longer 1.38. If we introduce the qualification that we are using a Celsius scale, however, the sentence immediately improves:

(4.89) This bowl is 1.38x hotter than that one on the Celsius scale.

Note, in addition, that the comparative size of intervals *can* be measured using ratios. So even though it does not make sense to say that Amir came two times later than Mora, it does make sense to say that Amir arrived twice as late, or that he stayed for twice as long as Mora. This is because the scale of temporal *instants* lacks a natural zero point, although temporal *slices*, that is, time intervals between instants, do.

Finally, let us highlight again that simply attesting the unacceptability of a ratio modifier is not enough to determine that the relevant adjective has an interval scale (recall that it wasn't enough even to rule out a ratio scale). This explanation is suggestive in the case of temperature or time, but that is because we know independently how temperature and clock time are measured—and in particular, we know that zero points on the relevant scales are arbitrary.

Might the goodness scale be interval? This is Lassiter's conclusion, but his reasoning is essentially abductive: given that according to him an ordinal and a ratio scale can be ruled out, only interval scales remain as a candidate among the type of scales attested in natural language.

In order to explore this matter further, we need to turn to the properties of concatenation. As we saw, ratio scales are by definition positive with respect to concatenation. Interval scales are not, so this is already a point in favor of attributing an interval scale to *good/bad*, in light of the validity of (4.85). However, interval scales are compatible with different features of concatenation, and Lassiter considers the following two: maximality and intermediacy, and settles on the latter.

For an interval scale to be MAXIMAL with respect to concatenation, it must assign to any concatenation the value that it assigns to the greatest element in it. More precisely, a scale $\mathcal{S} = \langle X, \geq, \circ \rangle$ is maximal with respect to concatenation iff for any non-overlapping $a, b \in X$, if $a \geq b$, then $a \approx a \circ b$.

Lassiter claims that the behavior of *good* reveals that its scale cannot be maximal with respect to concatenation. The main example he uses to argue against maximality is the following (shortened from the original):

Juliet is considering whether to feign death by taking the drug that Friar Laurence has offered her. If she does, it will put her in a coma, and she will die unless Friar Laurence administers the antidote. If she takes it and the Friar does administer the antidote, she will be able to live happily ever after with Romeo. If she does not take the drug, she will live a long boring life without Romeo; this is much better than being dead, though. Unfortunately, the Friar is known for being cruel and capricious, and it is extremely likely (though not certain) that he will 'forget' to administer the antidote if she takes the drug (Lassiter p.183).

The possible outcomes are ordered by decreasing preference according to the story:

1. Juliet takes the drug and the Friar administers the antidote. **drug \wedge antidote**

2. Juliet does not take the drug. $\sim \text{drug}$
3. Juliet takes the drug and the Friar does not administer the antidote. $\text{drug} \wedge \sim \text{antidote}$

However, the following is also true about the story:

4. It is better for Juliet to not take the drug than to take it. $\sim \text{drug} > \text{drug}$

The problem for maximality is the following: since any sentence ϕ is logically equivalent to $(\phi \wedge \psi) \vee (\phi \wedge \sim \psi)$, then **drug** is logically equivalent to $(\text{drug} \wedge \text{antidote}) \vee (\text{drug} \wedge \sim \text{antidote})$. Now, maximality implies that **drug** is exactly as valuable as the best of its disjuncts, namely **drug** \wedge **antidote**. But then, taking the drug is exactly as good as taking the drug *and* getting the antidote, which implies that **drug** is better than $\sim \text{drug}$, contradicting 4.

Lassiter uses examples like this to argue that the goodness scale cannot be maximal with respect to concatenation.¹⁷

A readily alternative is to consider that it is intermediate with respect to concatenation, just like temperature. A scale is INTERMEDIATE with respect to concatenation if the value of a compound individual is somewhere *between* the values of its compounds. More formally, a scale $\mathcal{S} = \langle X, \geq, \circ \rangle$ is intermediate with respect to concatenation iff for any $a, b \in X$,

- if $a \geq b$, then $a \geq (a \circ b) \geq b$, and
- if $a > b$, then $a > (a \circ b) > b$ —unless a or b receive zero weight.

Attributing an interval scale that is intermediate with respect to concatenation immediately helps with the Juliet case, since if the goodness scale is intermediate, it is no longer the case that **drug** must be as good as **drug** \wedge **antidote**. In particular, as Lassiter argues, the actual value of **drug** must be some intermediate value calculated based on the *probabilities* of the outcomes **drug** \wedge **antidote** and **drug** \wedge \sim **antidote**, which delivers the intuitively right result that **drug** cannot be too good, given that the Friar is likely to not administer the antidote.

4.6.4 Discussion

Based on these observations, Lassiter preliminarily concludes that the goodness scale must be an interval scale that is intermediate with respect to concatenation, and offers a semantics that captures those facts (Lassiter 2017, 186 and ff). Our purpose at the end of this chapter is not to discuss that semantics, but to show that the hyperplan semantics presented in the previous sections can accommodate the pieces of data that Lassiter points our attention to, namely the inference pattern in (4.85) and the Juliet scenario. Moreover, we also raise a problem for Lassiter’s classification, based on the admissibility of “round” ratio modifiers. We restrict ourselves to *good* taking propositional arguments, but what we say here applies to other evaluative adjectives and objects of evaluation of other types (although perhaps these observations do not apply to negative evaluatives, see n.17).

¹⁷ Interestingly, the argument does not go through if we substitute *good* for *bad*. Assume that the previous ordering of outcomes is in increasing dis-preference, and that **drug** is worse for Juliet than $\sim \text{drug}$. If *bad* is maximal, the badness of **drug** is equivalent to the worst disjunct of $(\text{drug} \wedge \text{antidote}) \vee (\text{drug} \wedge \sim \text{antidote})$, namely **drug** \wedge \sim **antidote**. And then, **drug** is worse for Juliet than $\sim \text{drug}$. This is consistent with the ordering. Why is this so? It might be that the scale of goodness is not maximal, but the scale of badness is. After all, it might sound too optimistic to think that a bunch of things are as good as the best in them; but it is comparatively more reasonable to think that a bunch of things are as bad as the worst in them.

For simplicity, let us restate the inference pattern in (4.85) in the language of \mathcal{L}_V (see §3.4):

$$(4.90) \quad \begin{array}{ll} \text{a. } p \geq q \\ \text{b. } p \geq r \\ \therefore p \geq (q \vee r) \end{array}$$

In order to test whether this inference patterns holds in our semantics, we need to spell out the meaning of each premise and the conclusion. We already know what the premises mean (see §3.4.2): relative to a model M with a structure $S = \langle W, H \rangle$,

$$(4.91) \quad \llbracket p \geq q \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) : \llbracket \uparrow q \rrbracket_{\langle w_i, h_i, X \rangle} = 1 \rightarrow \llbracket \uparrow p \rrbracket_{\langle w_i, h_i, X \rangle} = 1$$

$$(4.92) \quad \llbracket p \geq r \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) : \llbracket \uparrow r \rrbracket_{\langle w_i, h_i, X \rangle} = 1 \rightarrow \llbracket \uparrow p \rrbracket_{\langle w_i, h_i, X \rangle} = 1$$

In words: p is at least as valuable as q/r just in case every subset of $\mathcal{P}(W)$ is such that, if q/r is valuable relative to it, then p is too. In turn, this formulae can be spelled out as follows:

$$(4.93) \quad \llbracket p \geq q \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) : \\ \{w' : \llbracket q \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X) \rightarrow \{w' : \llbracket p \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X).$$

$$(4.94) \quad \llbracket p \geq r \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) : \\ \{w' : \llbracket r \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X) \rightarrow \{w' : \llbracket p \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X).$$

In words: p is at least as valuable as q/r just in case every subset X of $\mathcal{P}(W)$ is such that, if the worldly component of q/r is among the alternatives preferred by the hyperplan of the index relative to X , then the worldly component of p is among the preferred alternatives as well.

In virtue of the disjunction rule in (3.21), the set of worlds at which a disjunction is true is the set of worlds at which either the first disjunct is true or the second disjunct is true:

$$(4.95) \quad \{w' : \llbracket q \vee r \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} = \{w' : \llbracket q \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1 \text{ or } \llbracket r \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\}.$$

Thus, the conclusion has the following meaning:

$$(4.96) \quad \llbracket p \geq (q \vee r) \rrbracket_{\langle w_i, h_i, \mathcal{P}(W) \rangle} = 1 \text{ iff } \forall X \subseteq \mathcal{P}(W) : \\ \{w' : \llbracket q \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1 \text{ or } \llbracket r \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X) \rightarrow \{w' : \llbracket p \rrbracket_{\langle w', h_i, \mathcal{P}(W) \rangle} = 1\} \in h_i(X).$$

That is, p is at least as good as q or r just in case every subset of $\mathcal{P}(W)$ relative to which q or r is preferred, is also one relative to which p is also preferred by the hyperplan of the index.

In order to show that our hyperplan semantics guarantees the validity of (4.90), all we need to show is that any index of evaluation that makes both premises true makes the conclusion true as well. Suppose that that were not the case: then, there ought to be some subset X of $\mathcal{P}(W)$ such that, if q is selected, then p also is (making the first premise true). And if r is selected, then p also is (making the second premise true). But it also ought to be false that, if q or r is selected, then p also is (making the conclusion false). For that to be the case, then at least q or r have to be selected but not p . By the first premise, we know that X is such that, if q is selected, then p also is. So in order to make the conclusion false, it would have to be the case that, if r is selected, then p is not. But by the second premise we know that X is such that, if r is selected, then p also is. By *reductio*, there can be no such X , and (4.90) is valid in \mathcal{L}_V .

The Juliet case, on the other hand, is easier to tackle. The possible outcomes are repeated here (in decreasing order of preference):

1. **drug** \wedge **antidote**
2. \sim **drug**
3. **drug** $\wedge \sim$ **antidote**

Together with the conflicting observation:

4. \sim **drug** $>$ **drug**

In a similar vein to Lassiter's treatment of the case, the point here is that all the outcomes of the case ought *not* to be evaluated relative to the same index of evaluation, because the relevant set of alternatives are not always the same. Given how the case is set up, outcomes (2) and (3) as well as observation (4) all assume that the Friar will not administer the antidote—either as a background condition on alternatives or as part of the alternatives under evaluation. But crucially, (1) cannot be evaluated relative to such an index of evaluation, because it considers as an option that the Friar does administer the antidote. In order to establish an ordering between outcomes (1)–(3), those sentences cannot be evaluated relative to an index of evaluation that assumes that the Friar will not administer the antidote, which is what the observation in (4) is based on. There is no contradiction, rather, there is a necessary shift in the index of evaluation of those sentences, which pose no problem for our hyperplan semantics.

Should we conclude with Lassiter that *good* has an interval scale that is intermediate with respect to concatenation? We would, if not for the following: recall that ratio modifiers were acceptable with ratio-scale adjectives and unacceptable with interval-scale adjectives:

(4.82) Amir is 1.1030x as tall as Mora.

(4.86) ?? This bowl is 1.38x hotter than that one.

As we noted, evaluative adjectives pattern like interval adjectives when we look at precise ratio modifiers like the ones above:

(4.83) ?? Volunteering is 1.1030x as good as donating.

But surprisingly, when we turn from precise to “round” ratio modifiers, evaluative adjectives behave again like ratio-scale adjectives (Faroldi and Soria Ruiz 2017; Solt 2018; also Cresswell 1976, p. 283):

(4.97) Your daughter is like, four times more beautiful.¹⁸

(4.98) He'd have to be ten times more charming than Arnold.¹⁹

So we have arrived at a puzzle: data about concatenation and the non-acceptability of precise ratio modifiers suggest that evaluatives have interval scales. But the acceptability of round ratio modifiers speaks against an interval scale.

¹⁸Adapted from the series *Fresh Off the Boat*, season 5 chapter 5, 2018.

¹⁹Adapted from the movie *Pulp Fiction*, 1994. Admittedly, two examples taken out of movies are not much, but the reader surely can come up with many more natural sounding examples.

One possible way out—Lassiter’s solution (2017, 89 and ff)—would be to dispel the data about round ratio modifiers by arguing that they are hyperbolic uses. Lassiter says that the fact that those sentences become unacceptable when we add an adverb like *exactly* points in this direction:

(4.99) Your daughter is like, (# exactly) four times more beautiful.

(4.100) He’d have to be (# exactly) ten times more charming than Arnold.

This is not, however, an argument for hyperbolic uses—rather, it is a consequence of the fact that precise measures (i.e. 3.9 times, 4.154 times,... more beautiful) are not interpretable, which we already knew. This is so because in general, *exactly* n reduces *pragmatic slack* in the value of n (Lasersohn 1999). An intuitive way of characterizing this contrast (not necessarily Lasersohn’s view) is the following: *Bill is twice as tall as Ann* can mean that Bill is 1.99x or 2.01x as tall as Ann. But crucially, *Bill is exactly twice as tall as Ann* means that he is *not* 1.99x nor 2.01x as tall. However, if those ratios are antecedently uninterpretable, then it is to be expected that adding *exactly* makes those sentences bad.

A different way out of this dilemma would be to hold on to the possibility of a ratio scale, albeit a very restricted one. Evaluative scales would then have a kind of ratio scale that only admit ratios that result from multiplication by very round numbers, like 2, 4, 10, 100, 1000... More formally, the admissible measure functions obtained following condition (2) in Theorem 2 above should be restricted to those that result from multiplying admissible measure functions by numbers in some subset of very round members of the natural numbers. This would technically capture the data about the loose ratio modifiers, but it would not account for the point about concatenation.

We leave this issue for future work, but note that this is a genuine puzzle that concerns the scale of all evaluative adjectives, since they all (roughly) present the kind of sensitivity to precise and round ratio modifiers that we have presented here.

4.7 Recap & conclusion

In this long chapter, we turned to apply the semantics of \mathcal{L}_V to natural language expressions: first, the pair of thin evaluative adjectives *good* and *bad* in their guises as propositional operators (that is, taking propositions as their arguments). This move required complicating slightly the semantics of \mathcal{L}_V by introducing a parameter that can supply sets of alternatives to the hyperplan of the index of evaluation. After doing this, we were in a position to solve what we dubbed the “sub-sentential Frege-Geach” problem in Chapter 2, namely the problem of why evaluative adjectives in comparative form do not express outright practical commitments, even though their meaning is compositionally derived from the positive form of those adjectives, which do express such attitudes. In short, our answer was that, when a speaker uses the positive form of an evaluative adjective, she adopts a commitment to act in a certain way in the very situation of speech, in virtue of the fact that she is expressing a certain choice among a set of practical alternatives supplied by the index of evaluation; while when she uses the comparative she does not. We then went on to discuss other bearers of value, such as action-types and individuals, and we saw how to apply this proposal to evaluative adjectives with increasing degrees of thickness or (stable) descriptive meaning.

In the last section of this chapter, we discussed further semantic properties of evaluative adjectives, first in relation to the relative/absolute distinction, and then in connection to features of their scales. We concluded—somewhat preliminarily—that evaluative adjectives appear to be relative-standard, although they are definitely not paradigmatic examples thereof. We then turned to the question of what type of scale they have—whether ordinal, ratio or interval. Following Lassiter 2017, we have argued that there are strong reasons to think that their scale is not ordinal and not ratio. The remaining option is that it is an interval scale that is intermediate with respect to concatenation, but that possibility clashes with the acceptability of “round” ratio modifiers.

In the following chapters, we will explore further the notion of evaluative and subjective properties, and we will then see how evaluative sentences work in conversation—which will require us to move from the static semantics given here to a dynamic model of communication.

Chapter 5

Subjectivity & evaluativity: an experimental approach to ordering subjectivity

Summary¹

This chapter discusses two notions that have been the focus of much discussion in recent linguistic and philosophical literature, namely the notion of *judge-dependency* and *ordering subjectivity*. In recent experimental work, Stephanie Solt has argued that the spectrum-like nature of the phenomenon of ordering subjectivity calls for a distinction, within the class of subjective adjectives, between multi-dimensional and judge-dependent ones. We agree with her, but we supplement her conclusions by arguing that applying her experimental paradigm to moral adjectives suggests a further distinction, within the class of judge-dependent adjectives, between (at least) experiential and evaluative adjectives.

5.1 Introduction

In Chapter 2, evaluative adjectives were characterized by the possession of two features: ORDERING SUBJECTIVITY and ACTION-GUIDANCE. Up to this point, the latter feature has been the main focus of this dissertation: indeed, most of the long two previous chapters consists in an exposition of a scalar semantics for evaluative adjectives that builds action-guidance onto the very meaning of evaluative adjectives by—roughly—assigning plans of action as their denotation. This chapter, by contrast, returns to discussing ordering subjectivity.

Consider the contrast between the following two dialogues:

- (5.1) a. Carmen is taller than Marieke. *≈ not subjective*
 b. No, Marieke is taller.
- (5.2) a. The cake is tastier than the chocolate cookies. *≈ subjective*
 b. No, the chocolate cookies are tastier.

¹This chapter stems from joint work with Federico Faroldi. Mora Maldonado has provide invaluable help with the experimental design, data analysis and comments on a final draft. I thank both, but any errors are my sole responsibility.

As we saw in §2.2.1, the speakers in both these dialogues are having a disagreement, but whereas the disagreement in (5.1) seems to turn on a completely objective and factual matter—i.e. the respective heights of Carmen and Marieke, the disagreement in (5.2) appears to be largely subjective, as it relies on the interlocutors’ food preferences.

This occurs when these adjectives appear in comparative form—if they were in the positive form, both dialogues would appear subjective. The property in virtue of which the comparative form of an adjective gives rise to disagreements that appear to be subjective is called ORDERING SUBJECTIVITY (Bylinina 2017; Solt 2018). Ordering subjectivity distinguishes evaluative (among other) adjectives like *good*, *beautiful* or *tasty* from dimensional adjectives like *tall* or *wide* (§2.2.1).

Solt (2018) presents novel experimental data about the distribution of this property.² In the literature, it is assumed (often implicitly) that adjectives are clearly divided between those that are subject to ordering subjectivity and those that are not. The main empirical discovery of Solt 2018 is that this is not the case: when one considers a sufficiently broad class of gradable adjectives, it turns out that speakers’ intuitions about the distribution of ordering subjectivity are not uniform. Participants generally agree that dimensional adjectives like *tall* clearly lack this property, and that predicates of personal taste (PPTs), such as *fun* and *tasty*, have it. But her results show that many other adjectives fall in a mixed class, in that there is a large variability in participants’ opinions as to whether disagreement dialogues involving those adjectives are subjective.³ Examples of adjectives in this mixed class are: *dull*, *curved* or *rough*.

Solt offers an account for why certain adjectives might fall in the mixed class, and for why certain adjectives fall in the clearly subjective class. Roughly, her view is that adjectives that fall in the mixed class are such that comparative uses can be given an objective or a subjective interpretation, for a variety of reasons. Some adjectives, such as *heavy*, are ambiguous between a qualitative or subjective and an quantitative or objective interpretation (*heavy* can mean how much an object weighs and how much its weight *feels* to, e.g., someone carrying it; Kennedy 2013b). Others, such as *curved*, can be broken down into different dimensions or aspects (frequency and number of curves, for example), some of which may be objectively measurable and some of which may not be.

On the other hand, for adjectives that fall in the clearly subjective class there is no readily available objective interpretation. Moreover, all these adjectives can be characterized, according to Solt, as JUDGE-DEPENDENT. That is, their semantics involves the notion of a semantic judge (see also §3.2.1). The difference between adjectives in the mixed class and those that are in the clearly subjective class is, therefore, modelled by Solt by introducing a judge in the semantics of the latter. In this way, Solt aims to represent the fact that clearly subjective adjectives necessarily involve an ‘inherent human element’ (2018, p. 76, after McNally and Stojanovic 2017).

Our purpose is to extend this experimental paradigm to moral adjectives, which Solt does not consider. Moral adjectives are interesting *vis-à-vis* Solt’s experimental paradigm for two reasons: first, moral adjectives arguably involve a judge, but the judge of a moral judgment hardly

² See also E. A. Smith et al. 2015.

³ What we mean by there being a large variability in participants’ opinion is that some participants took the relevant dialogues to be subjective disagreements, while some did not. But in both Solt’s experiment and ours, participants were (i) given a binary choice between describing a disagreement as a matter of opinion or fact, and (ii) presented with dialogues containing the relevant adjectives only once. So we cannot conclude whether participants are consistent in their judgments. The variability measure is across, and not within, participants.

falls under the usual notions of judge discussed in the literature on perspectival expressions. To be more specific, the usual notions of a judge in judge-dependent semantics represent it as a subject of experiences (for PPTs), qualitative perceptions (color adjectives) or emotions (emotional predicates); but moral judgment does not in general require any particular experience, perception or emotion. Consequently, the semantics of moral vocabulary does not require invoking experiencers nor perceptual or emotional agents. Hence, if moral adjectives turn out to pattern with the clearly subjective class of adjectives, there is reason to enrich Solt’s notion of a judge, which is markedly experiential.

And secondly, one might suspect at least some people to have strongly objective intuitions about morality: a disagreement about what course of action was morally more correct might seem *prima facie* more objective than a disagreement about what movie was more fun.⁴ This is a reason to expect moral adjectives to pattern with the adjectives in the mixed class. If they do, then there is reason to challenge Solt’s claim that adjectives in the mixed class do not have a judge dependent semantics.

As we will see, our results show that moral adjectives fall under the clearly subjective class. Thus, we defend that the notion of judge-dependency has to be enriched to make room for the type of judge that is involved in moral judgment. In line with what has been defended in previous chapters, we defend that the role of the judge that is required for an appropriate semantics for moral adjectives is a practical agent, that is, an agent with intentions to act in certain ways. This is better seen considering an example:

- (5.3) a. Bill leads a more ethical lifestyle than Amy.
b. No, Amy does.

As our results show, participants take dialogues like (5.3) to be subjective more or less as often as dialogues involving aesthetic adjectives, PPTs or emotional adjectives. But in contrast to those cases, the impression that the disagreement between the interlocutors in (5.3) is subjective need not be due to the fact that the speakers have different experiences, perceptions or emotions. What we defend is that the “clash” in disagreements like (5.3) is due to speakers having conflicting *plans*. This is already captured in the hyperplan semantics offered in previous chapters, so we will leave all formalism aside.

The chapter is structured as follows: §5.2 presents Solt’s 2018 study. Her experimental design (§5.2.1) and results are laid out (§5.2.2), and her discussion of the results is summarized (§5.2.3). Subsequently, we present our study, which uses essentially Solt’s paradigm to test a set of moral adjectives together with a sample of the adjectives that she tested (§5.3). We finish this chapter with a discussion of our results (§5.4). There, we argue for the need to enrich the notion of judge-dependency in the way just discussed. §5.5 concludes.

5.2 Solt’s 2018 study

In her 2018, Solt tackles the question of the distribution of ordering subjectivity across gradable adjectives. Let us define ordering subjectivity as follows (see also §2.2.1):

⁴ See Stojanovic 2019 for the distinction between disagreement about morality and about taste.

Definition 20 (Ordering subjectivity) A gradable adjective G is ORDERING SUBJECTIVE iff whether a sentence containing the comparative form of G is true or false is perceived to be a subjective matter.

That definition calls for at least two clarifications: first, what is meant by *subjective*? And secondly, why is the comparative relevant?

Different authors give different characterizations of *subjectivity*. For our purposes however, it's easiest to simply follow Solt's characterization. As we will see, Solt presents participants with disagreement dialogues and asks them whether they consider the issue to be a *matter of opinion*. This will be our criterion as well. If the answer to a certain question or issue is perceived to be a matter of opinion, it is a subjective matter; if it is not, then it is not a subjective matter. Just to illustrate: whether one can have red wine with tuna is a matter of opinion; whether the sun sets at 4:36:32pm on January 4th, 2019 at coordinates 55.953251,-3.188267 is not. Note that this a binary distinction: an issue or a question either is or is not a matter of opinion.⁵

To the second question there is a short and a long answer. The short answer is that the class of gradable adjectives that give rise to subjective disagreement in the comparative form is a subclass of those that give rise to subjective disagreement in the positive form: *tall* and *tasty* belong to the latter class, but only *tasty* belongs to the former. So by considering the comparative form, we are looking at a more localized phenomenon than by looking at the positive form.

- (5.4) a. Carmen is tall. \approx subjective
b. No, she is not.
- (5.5) a. Haggis is tasty. \approx subjective
b. No, it is not.
- (5.1) a. Carmen is taller than Marieke. \approx not subjective
b. No, Marieke is taller.
- (5.2) a. Haggis is tastier than andouillette. \approx subjective
b. No, andouillette is tastier!

The longer answer is as follows: in standard semantics for gradable adjectives, the question of whether an adjective is subjective boils down to the question of whether its meaning is determined with the aid of some contextual parameter (§2.2.2). Now, there is an important difference between the positive and the comparative form of a gradable adjective: the positive form mobilizes at least one contextual parameter more than the comparative form. So if the positive form is perceived to be subjective, it could be due to that parameter.⁶ Hence, the subjectivity of a disagreement dialogue involving the positive form of a gradable adjective can

⁵We should not presuppose that speakers are perfectly capable of drawing this distinction: as this recent Pew Research Center poll shows, people are surprisingly bad at distinguishing statements of opinion and statements of fact taken from the news (<https://www.journalism.org/2018/06/18/distinguishing-between-factual-and-opinion-statements-in-the-news/>). Moreover, people appear to consider more factual those statements that lean towards their values and opinions. Solt's and the present study abstract away from this latter confound by avoiding issues that participants could be antecedently opinionated about. For example, Solt does not present subjects with disagreement dialogs that compare, say, the taste of cookies to the taste of stew. Rather, their dialogs contain items like *this is tastier than that*. We do the same.

⁶In degree semantics (Kennedy 2007), the positive form requires determining a specialized degree on a scale, called a STANDARD or THRESHOLD, such that the positive form of the adjective can be truthfully predicated of any

be due to speakers determining a different threshold or a different comparison class for the positive form.

The comparative, by contrast, does not mobilize contextually determined thresholds or comparison classes. If a comparison is perceived as subjective, it has to be because *the very ordering* that forms the extension of the adjective is context-sensitive. Therefore, asking whether a disagreement such as (5.1) and (5.2) is a matter of opinion is a way of investigating whether the ordering induced by the relevant adjective is context-sensitive. That is the question that Solt 2018 tackles.

5.2.1 Experimental design

To answer that question, Solt tests a set of gradable adjectives, classified along traditional categories:

1. **RELNUM**: Relative-standard with numerical measures:⁷ e.g. *tall, wide*
2. **ABS2**: Absolute-standard with closed scales on both ends:⁸ e.g. *empty, full*
3. **ABS1**: Absolute-standard with closed scale on one end: e.g. *salty, wet*
4. **RELNO**: Relative-standard without numerical measures: e.g. *light, hard*
5. **EVAL**: Evaluative adjectives, and among them, PPTs (*tasty*), aesthetic (*beautiful*) and emotional adjectives (*happy*).

She presents subjects with disagreement dialogues such as the ones we have just seen, and asks subjects whether the disagreement they just saw is a matter of opinion. Here are a few of her test items:

- (5.6) a. Look – Tommy’s shirt is dirtier than the one his little brother Billy is wearing.
b. No, Billy’s shirt is dirtier than Tommy’s.
- (5.7) a. The lecture we heard last week was more boring than today’s lecture.
b. No, today’s lecture was more boring.
- (5.8) a. The program we watched about India was more interesting than the one about Japan.
b. No, the program about Japan was the more interesting of the two.

Below each dialogue, she presents participants with a choice between the following two options: ‘*only one can be right, the other must be wrong*’ and ‘*it’s a matter of opinion*’.

object whose degree is above the threshold. In a delineation approach (C. Barker 2002; Burnett 2017; Klein 1980), the positive form requires introducing a contextually determined comparison class relative to which the extension, anti-extension and extension gap of the relevant adjective are determined. See the discussion in Chapter 2.

⁷Adjectives have numerical measures just in case they admit numerical modifiers, e.g. ‘*Carmen is 168cm tall*’.

⁸As Solt notes, ABS2 adjectives may be thought to have numerical measures as well, since they admit of proportional modifiers, e.g. ‘*55% full*’.

5.2.2 Results

The main result of Solt's experiment is that participants' perception of ordering subjectivity is largely variable. That is, for a large class of adjectives there is variability across participants as to whether the relevant dialogues are perceived to be subjective. While RELNUM and ABS2 adjectives were perceived as objective and EVAL adjectives were perceived as subjective for most participants, there were a surprisingly high number of adjectives for which participants' judgments diverged, namely all adjectives in the ABS1 and RELNO classes.

5.2.3 Solt's account

In the remainder of her paper, Solt offers an account of her results, that is, of why there is a class of clearly objective and subjective adjectives, as well as a big mixed class of adjectives. In general terms, her account of why adjectives fall into a clearly objective, mixed and clearly subjective class relies on the availability of 'precise, quantitative measurement' (Solt 2018, p. 67). Adjectives in the clearly objective class (RELNUM and ABS2) admit numerical modifiers, so it is reasonable for speakers to assume that those disagreements can be resolved objectively, by measuring the degree of the relevant property that each *relatum* has, and thus that only one of the speakers in those dialogues can be right. On the other hand, adjectives in the clearly subjective class (EVAL) are such that numerical measures are not available at all. Finally, adjectives in the mixed class (RELNO and ABS1) *can* receive an objective interpretation. Participants who interpreted adjectives in the mixed class subjectively would have answered that the disagreements were a matter of opinion, whereas speakers who interpreted those same adjectives objectively, would have answered that only one speaker could be right.

The mixed class

Why would participants interpret adjectives in the mixed class sometimes subjectively, sometimes not? Solt offers a few diagnostics. First, some adjectives in the mixed class, such as *hard*, are ambiguous between a subjective/qualitative and an objective/quantitative interpretation (Kennedy 2013b). This is likely the case for *hard*: whereas the most common interpretation of hardness is in relation to a certain perception (how hard a chair *feels*), there exist scientific measures of hardness for materials.

Other adjectives are MULTI-DIMENSIONAL, which means that their orderings depend on different dimensions or aspects.⁹ For example, the degree of dirt on a towel can be broken down into several dimensions: stain distribution, stubbornness and nastiness. This opens up two very salient aspects for cross-subject variability, namely the exact dimensions that are taken into account and the weight that each dimension receives in the overall dirt-ordering. So a disagreement about what towel is dirtier could be deemed subjective by someone who thinks that speakers engaging in that disagreement had different dimensions of dirtiness and/or were giving different weight to each dimension of dirtiness. Importantly however, it is possible too that some speakers consider that there is an objective measure of dirtiness—for example, how much

⁹ The putative multi-dimensionality of evaluative adjectives was already discussed briefly in Chapter 2 (see §2.2.2).

extraneous material is attached to the towel's surface. This would allow the needed variability for an adjective like *dirty* to fall in the mixed class.

Multi-dimensionality has been the subject of much work in recent literature on gradability. Especially relevant is the work of Sassoon (2013a, 2016), (although see as well McNally and Stojanovic 2017). Sassoon uses as a grammatical diagnostic for multi-dimensional adjectives the admissibility of “respect-denoting” PPs and quantificational phrases, such as *with respect to* or *except for*. To see this, contrast a dimensional adjective like *tall* with a multi-dimensional adjective like *healthy*:

- (5.9) a. Carmen is healthy with respect to blood pressure.
b. # Marieke is tall with respect to...
- (5.10) a. In what respects is Carmen healthy?
b. # In what respects is Marieke tall?
- (5.11) a. Carmen is healthy except for her sugar level.
b. # Marieke is tall except for...

Interestingly however, many of the adjectives tested by Solt do not fare too well in these tests. Aside from the dimensional ones, which are predicted to be unacceptable, consider the following examples from the mixed class (the examples are taken from her paper, p.74):

- (5.12) a. The line was(n't) straight/curved except for ...
b. The leather was(n't) smooth/rough except for ...
c. The knife was(n't) sharp/dull except for ...
d. The soup was(n't) salty except for ...

For all those adjectives, it is difficult to fill the blank with a dimension of the relevant property.

Nonetheless, suggests Solt, even if these adjectives do not clearly pass Sassoon's tests for multi-dimensionality, we can distinguish different dimensions or aspects that can play a role in determining an ordering. For instance, when considering a curved road, we can take into account the frequency and sharpness of the curves it has, even though it is very odd to say something like:

- (5.13) The road was curved with respect to frequency, but not with respect to sharpness.

The fact that we can conceptually distinguish different dimensions that may affect the ordering induced by an adjective even though the adjective may not appear multi-dimensional under Sassoon's tests, suggest a distinction between two types of multi-dimensionality:

1. Quantificational multi-dimensionality: The type of multi-dimensionality detectable *via* respect-denoting PPs and quantification (e.g., *healthy*).
2. Conceptual multi-dimensionality: not detectable *via* respect-denoting PPs and quantification, but conceptually distinguishable (e.g., *curved*).

Solt's claim is that many adjectives in the mixed class count as conceptually multi-dimensional, if not quantificationally so.

The clearly subjective

Adjectives that clearly fall in the subjective class (i.e. those antecedently classed as EVAL), according to Solt, are those for which no numerical measure can be given. This correlates with the fact that they involve an ‘inherent human element’ (2018, p. 76), or as McNally and Stojanovic put it, ‘the necessary mediation of some sentient individual’ (2017, p. 28). In the literature on PPTs, this inherent human element has usually been cashed out in the introduction of a context-dependent JUDGE-parameter, either as part of the proposition expressed by sentences containing the relevant expressions—so called CONTEXTUALISM about PPTs, cf. Sæbø 2009; Stojanovic 2007 *a.m.o.*—or as a parameter in the index of evaluation—the RELATIVIST proposal, cf. Lasersohn 2005; Stephenson 2007a *a.m.o.*

For our present purposes, more than the formal implementation of a judge-semantics it matters what “role” the judge plays in determining the ordering of individuals that form the extension of an adjective. Solt discusses mainly three possible roles that a judge can play in determining the extension of adjectives in the EVAL class:

1. An experiencer:

(5.14) Chocolate is tastier than cheese.

2. A subject of emotions:

(5.15) The funeral was happier than the party.

3. A “perceiver”:

(5.16) This painting is uglier than that one.

Different dimensions may also play a role in determining the extension of these adjectives. Following Solt however, those dimensions do not determine directly the orderings of the adjectives in EVAL, but rather act as factors that contribute and alter the judge’s overall assessment. This would block putative objective or quantitative interpretations: for example, taste can presumably be broken into (at least) saltiness and texture, which are potentially measurable properties. However, if the saltiness of a food is inevitably no more than a factor in a judge’s subjective experience of a food’s taste, then it is difficult to see how saltiness could impact different judge’s experiences of that food in an uniform way.

In sum: Solt’s main result is that there is great variability in people’s assessment of ordering subjectivity across adjectives, ranging from adjectives that clearly lack it (RELNUM and ABS2) to adjectives that clearly have it (EVAL) through a rather big mixed class (RELNO and ABS1). The mixed adjectives can be given subjective and objective interpretations, which is why participants diverge in their judgments. This, in its turn, is most likely due to either a qualitative/quantitative ambiguity in their meaning or (quantificational or conceptual) multi-dimensionality. On the other hand, adjectives in the clearly subjective class are broadly judge-dependent, which means that their orderings are determined relative to a judge, which may play the role of a subject of experiences, perceptions or emotions.

5.3 Our study

Given Solt’s conclusions, it would be suggestive to see how moral adjectives fare in this experimental paradigm. As we said in the introduction, there are two main reasons for this, which were already summarized in the introduction: first, even though moral adjectives arguably involve a judge, it hardly falls under the judge “roles” discussed by Solt—that is, the judge of a moral statement need not, in general, be an experiencer, perceiver or have any particular emotion. Secondly, one might suspect that some people have strongly objective intuitions about morality. This is a reason to expect moral adjectives to pattern with the mixed class.

In order to further motivate the first claim, we may consider the acceptability of *for/to* PPs, which is different for moral adjectives *vis-à-vis* adjectives in Solt’s EVAL class. According to Bylinina 2017, these PPs denote experiencers and fill an argument slot for certain adjectives, such as PPTs. Indeed, PPTs (e.g. *tasty, fun*) take *for/to* PPs, while some moral adjectives do not. However, this is not visible at first sight, because *for/to* PPs can also be adjuncts, semantically denoting a doxastic operator (something along the lines of *in x’s opinion* or *according to x*). Consider the following sentences:

- (5.17) The test was fun for me.
- (5.18) The soup was tasty to Alice.
- (5.19) Torture is unethical for me.

Contrary to appearances, the PP *for me* in (5.19) does not have the same syntactic and semantic role that the PPs in (5.17) and (5.18) (Glanzberg 2007).¹⁰ Firstly, the choice of preposition in (5.17)-(5.18) is idiosyncratic, but not in (5.19) (Stephenson 2007a, p. 520): swapping *for* for *to* in (5.17)-(5.18) results in bad sentences, while it is acceptable with (5.19):

- (5.20) ?? The test was fun to me.
- (5.21) ?? The soup was tasty for Alice.
- (5.22) Torture is unethical to me.

Secondly, adjuncts are separable from the main verb by other adjuncts, while arguments are not; an adjunct like *yesterday* cannot be inserted between a verb and its syntactic argument. But adjuncts can be inserted in such way:

- (5.23) I ate an [_{arg} apple] [_{adj} yesterday]
- (5.24) # I ate [_{adj} yesterday] [_{arg} an apple]
- (5.25) I ate an apple [_{adj} yesterday]
- (5.26) I ate an apple [_{adj} quickly] [_{adj} yesterday]

Predicates like *tasty* and *fun* display this pattern: (5.27) and (5.28) are fine, but inserting the adjuncts between the adjective and the judge PP makes the sentences bad:

- (5.27) The test was fun [_{j-PP} for me] [_{adj} in spite of the baby logic bit]

¹⁰The distinction between arguments and adjuncts isn’t clear cut in general, and even less so with respect to *for/to* PPs. In particular, the *for/to* PPs that we observe here are all optional, which is already a feature that syntactic arguments generally lack. We thank Manuel Križ for useful comments.

- (5.28) The soup was tasty [_{j-PP} to Alice] [_{adj} in spite of the celery]
 (5.29) # The test was fun [_{adj} in spite of the baby logic bit] [_{j-PP} for me]
 (5.30) # The soup was tasty [_{adj} in spite of the celery] [_{j-PP} to Alice]

By contrast, both constructions are fine with *unethical*:

- (5.31) Torture is unethical [_{j-PP} for me] [_{adj} under any circumstance]
 (5.32) Torture is unethical [_{adj} under any circumstance] [_{j-PP} for me]

We side with Bylinina in suggesting that the behavior of the PPs of (5.17)-(5.18) (idiosyncrasy; inseparability) points to a thematic relation between those phrases and the relevant predicates. By contrast, the lack of prepositional idiosyncrasy and the separability data suggests that the *for/to* PPs that we see in sentences like (5.19) are adjuncts, and in particular, doxastic operators. What (5.19) says is that, in the speaker's opinion, torture is unethical. In sum, moral adjectives like *unethical* do not admit thematic *for/to*-PPs.

We take these observations to provide further evidence that moral adjectives lack an experiential semantics. Thus, in this experimental study we set out to test how moral adjectives fare with respect to Solt's experimental paradigm. To do that, we took a sample of adjectives from Solt's experiment, added an approximately equal number of moral adjectives, and tested them under Solt's disagreement paradigm.

5.3.1 Participants

Participants were 40 native speakers of English, recruited *via* Amazon Mechanical Turk (MTurk). They were paid \$0.50 for their participation (the task took approx. 5 minutes). Recruiting was limited to MTurk workers with U.S. IP addresses. No participant was excluded.

5.3.2 Materials

Test items were based on 24 adjectives. From these, 13 were sampled from Solt's study:

1. **RELNUM**: *tall, expensive*
2. **ABS2**: *empty, full*
3. **ABS1**: *salty, wet*
4. **RELNO**: *light, hard*
5. **EVAL**: *tasty, ugly, happy, intelligent*

These were supplemented with 11 moral adjectives, classified along the following axes:

1. **THIN moral adjectives**: *moral, ethical, virtuous*
2. **THICK moral adjectives**: *coward, generous, loyal, honest*
3. **NORM/VALUE**: *important, justified, rational, valuable*

The difference between thin and thick adjectives has been discussed, so we will not insist on it (see §4.3.3). The class called NORM/VALUE is an heterogeneous class of all-purpose value and normative adjectives: *important* is usually applied to things that range from subjective to objective value; *justified* or *rational* are used indistinctly for moral, practical or epistemic praise; and *valuable* is an umbrella evaluative adjective that can be use to evaluate all kinds of things in all kinds of ways.

For each adjective, a disagreement dialogue was created. For the set of adjectives taken from Solt's study, the same dialogues were used; for the class of moral adjectives, new dialogues were devised. Here is a sample of the new dialogues:

- (5.33) a. General McAdam was more coward than General Smith in that particular battle.
b. No, Smith was more coward.
- (5.34) a. Ann's work is more valuable than Jim's
b. I disagree, Jim's work is more valuable.
- (5.35) a. It is more important to follow Mary's than Bill's advice.
b. I disagree, Bill's advice is more important.

Participants were presented with 11 test items and 13 control items, as well as 12 filler dialogues split between factual (A: *Sharks are mammals*; B: *No they are not!*) and subjective disagreements (A: *This restaurant has wonderful service, I love it*; B: *No, it's awful*). See the Appendix for the full list of critical items.

5.3.3 Procedure

Participants were presented with the following set of instructions:

This study is about disagreements between people. Sometimes when two people disagree, only one of them can be right, and the other must be wrong. For example, in this short dialogue, Speaker A and Speaker B can't both be right, because Rosa can't have been born in both July and April.

- Speaker A: Rosa was born in July.
- Speaker B: No, Rosa was born in April.

But sometimes when people disagree, there is no right or wrong answer - it's just a matter of opinion. Here's an example:

- Speaker A: Susan looks a lot like her sister.
- Speaker B: No, they don't look alike at all!

In this HIT, you will see a series of short dialogues between two speakers, A and B. Your task is to say whether there is a right or wrong answer, or whether it's a matter of opinion.

Please answer based on your intuitions; do not think too long about each question. **Do not proceed with this experiment if you are not a native English speaker.**

Participants were then presented with dialogues like (5.33)-(5.35) above, and told to choose among the following two options:

- What do you think?

1. Only one can be right, the other must be wrong.
2. It's a matter of opinion.

Answering 1 was classified as a FACT answer; answering 2 was classified as OPINION. At the end, participants were asked for their country, age, biological sex and were given the opportunity to comment.

5.3.4 Results

The proportion of FACT choices per Adjective is illustrated in Figure 5.1; the proportion of FACT choices per Adjective Class is illustrated in Figure 5.2.

Following Solt's analyses, we analysed the responses (FACT vs. OPINION) by modelling response-type likelihood using a logit mixed-effect models (Jaeger 2008), with the factor Adjective Class as fixed effect (8 levels), and random intercept per subject.

The reference level for this omnibus model (aka the baseline in treatment contrast) was the class RELNUM. The z-scores and p-values reported are those calculated by the *lme5* package by a Wald III test.

The results of this omnibus model indicate that all classes are significantly different from RELNUM, with the exception of ABS2. Model output is provided in Table 5.1.

	RelNum (intercept)
Abs2	$z = -0.62, p = 0.53$
Abs1	$z = -2.82, p < .01$
RelNo	$z = -2.31, p < .05$
Eval	$z = -8.13, p < .001$
N/V	$z = -8.958, p < .001$
Thin	$z = -7.95, p < .001$
Thick	$z = -7.3, p < .001$

Table 5.1: Results of Omnibus model

These results roughly replicate Solt's original findings, modulo the new classes added in our experiment.¹¹

In order to address our main question (i.e., where do moral adjectives fall in the subjectivity spectrum), we fit a second set of models where we compared each of the classes included in Solt (2018) to the new sets of adjectives (THIN, THICK and NORM/VALUE). These models were constructed in the same way as before (fixed and random structure) differing only in how the baseline was encoded.

- ABS1 (reference level) is significantly different from N/V ($z = -7.240, p < .001$), THICK ($z = -5.3, p < .001$) & THIN ($z = -6.054, p < .001$).
- RELNO (reference level) is significantly different from N/V ($z = -7.71, p < .001$), THICK ($z = -5.9, p < .001$) & THIN ($z = -6.42, p < .001$).

¹¹We have also reproduced the posthoc pairwise comparisons done by Solt(2018)(ABS1 vsEVAL: $z = -610, p < .001$, ABS1 vsRELNO: $z = .5, p = .5$, EVAL vsRELNO: $z = -5.8, p < .001$)

- EVAL (reference level) is *not* significantly different from N/V ($z = -1.78, p = .07$), THICK ($z = 1.71, p = .08$) & THIN ($z = .35, p = .072$).

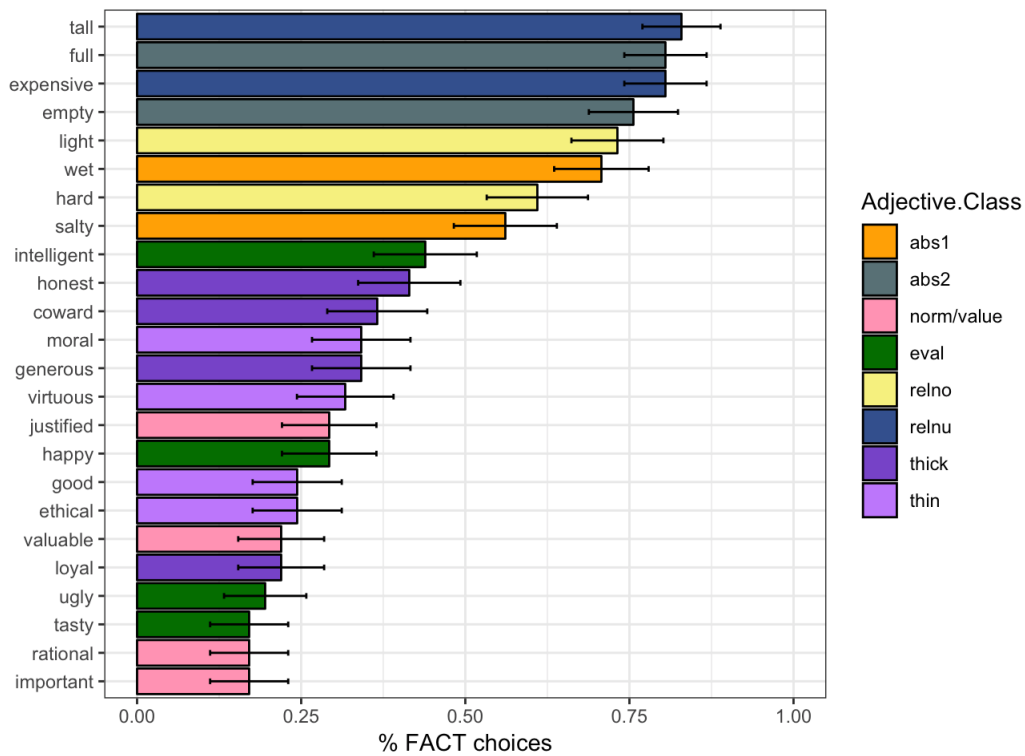


Figure 5.1: FACT choices per adjective




Figure 5.2: FACT choices per adjective class

5.4 Breaking down judge-dependency

Our results can be summarized in two claims: on the one hand, we replicated Solt’s result for the classes that she tested, which is good. On the other hand, the adjectives that we tested (which we may simply call moral adjectives) were found to pattern with the adjectives in the EVAL class. Our initial hypothesis that speakers might have stronger realist or objectivist intuitions about morality than say, about taste or aesthetics was not confirmed by these results. However, introducing moral adjectives into the picture ought to change slightly Solt’s characterization of evaluative adjectives. In this last section, we discuss the possibility that these results about moral adjectives call for a refinement of the notion of judge-dependency.

These results strongly suggest counting moral adjectives within the EVAL class. This is supported by Solt’s general characterization of the EVAL class as those adjectives whose orderings are determined, or mediated by some ‘inherent human element’. So much seems true of moral statements—that their truth depends on a human subject.

However, Solt’s characterization of the different “roles” that judges can play does not readily apply to moral thought: recall that Solt claims that the extension of judge-dependent adjectives is mediated by a judge, who is the subject of some experience, perception or emotion. Moral judgment, however, does not require any of those things. Consider the contrast between the following dialogues and (5.3):

- (5.2) a. The cake is tastier than the chocolate cookies.
b. No, the chocolate cookies are tastier.
- (5.36) a. Those sneakers are uglier than the sneakers you tried on earlier.
b. No, the earlier sneakers were uglier.
- (5.37) a. This cat is happier than that dog.
b. No, the dog is the happier one of them.
- (5.3) a. Bill leads a more ethical lifestyle than Amy.
b. No, Amy does.

In all these disagreements, an OPINION answer (i.e., thinking that the issue is a matter of opinion) could be driven by considering that each speaker is a different judge with a different—and valid—opinion about the matter. Moreover, this can be rooted on these speakers having different experiences, perceptions or feelings about the matter: in (5.2), speakers have both tried both foods and their tastes diverge; in (5.36) they have different aesthetic appreciation of the sneakers; and in (5.37) they have a different “emotional appraisal”, about the pets. But none of this applies straightforwardly to (5.3). To have a moral disagreement, it is not necessary to have different experiences or perceptions. Speakers in (5.3) need not have experienced in any way how Bill and Amy behave, nor do they need to see what emotions they appreciate in their behavior, nor what emotions Bill and Amy’s behavior causes in them. They are judging the matter differently, and people take them to be entitled to their divergent opinion (as witness our results). But it is not clear *what* exactly could their divergence be about.¹²

¹²Perhaps this applies to (5.36) as well: after all, it isn’t clear that two people looking at the same two pairs of sneakers could have a different perception of them. Consider, in this respect, the contrast with *tasty*: in (5.2), the experience upon which the evaluation rests can differ from speaker to speaker. You and I may perceive the taste of sugar differently, and that can interfere with our perception of the taste of chocolate in different ways, let’s say.

We want to close this chapter by pointing out two possible avenues to explore. First, recall that in previous chapters, we have defended that evaluative (including moral) adjectives can mobilize different priorities or preferences for action, which we modelled in terms of PLANS (i.e., sets of hyperplans). We can understand the disagreement between the speakers in (5.3), which is a disagreement about *ethical*, in terms of speakers advocating for different moral norms. And in turn, we can understand different moral norms in terms of plans that each speaker advocates. It need not be a disagreement about their experiences, emotions or perceptions, but simply a disagreement about how to conduct themselves. This type of disagreement would fall under what is sometimes known as DISAGREEMENT IN ATTITUDE (Gibbard 2003; Ridge 2013, 2014; Stevenson 1963).

A disagreement in attitude is a disagreement about *what to do*, rather than a disagreement about *how things stand*. Applied to dialogues like (5.3), the idea is that speakers in these dialogues are expressing conflicting practical attitudes (see §2.3). By uttering *Bill leads a more ethical lifestyle than Amy*, a speaker *A* expresses their support of Bill’s lifestyle over Amy’s; by replying things like *no, I disagree*, etc., a speaker *B* can either express rejection of Bill’s lifestyle over Amy’s or simply reject *A*’s support of it. Speakers in this type of dialogue disagree about what attitude to adopt, not about what is the case. Furthermore, under this conception of the disagreement in (5.3), the subjectivity of these dialogues is preserved since, even though speakers are disputing what lifestyle to support or reject, participants may be eschewing the idea that there is a single correct answer to that question.

If view of this, Solt’s characterization of judge-dependency should be refined in order to make room for the possibility that moral disagreement depends on advocating different plans for action. Judge-dependency cannot be reduced to experiential-semantics (Bylinina 2017).

Often times however, the different judge “roles” appear together. For instance, PPTs have both an experiential semantics and carry evaluative meaning. *Tasty*, for instance, both requires an experience of and positively evaluates whatever falls in its extension. But for some adjectives, that is not the case: *easy* (one of the EVAL adjectives tested by Solt) arguably requires an experiencer, but not an evaluator—since calling something easy can be neutral with respect to its value. By contrast, some moral adjectives, such as *ethical*, do not require an experiencer even though they carry evaluative meaning.

Secondly, note that at least some moral adjectives also admit *beneficiary* arguments (which are also *for/to* PPs):

- (5.38) a. The members of the committee were generous to Ann.
- b. Bill’s friends will be loyal to him.
- c. Mary’s advice is important for Sue.

Just like experiencer *for/to*-PPs, beneficiary arguments seem to be inseparable from the corresponding predicate:

- (5.39) a. Bill’s friends will be loyal to him to the end
- b. # Bill’s friends will be loyal to the end to him.

But in the case of an aesthetic evaluation, the perceptual experience upon which the evaluation rests is shared. In (5.36), there is a clear sense in which people are *seeing the same thing*. In light of this, there is some reason to drop the assumption that aesthetic adjectives belong in the ‘experiencer’ class, and perhaps we should consider them closer to moral adjectives than to PPTs and emotional predicates. We leave this for future consideration.

Nonetheless, beneficiary arguments are distinguishable from experiencer *for/to*-PPs by the fact that the sentences in which they appear do not entail that the person referenced in the PP has any experience or perception whatsoever. In (5.38a) for instance, there is no inference that Ann experienced the generosity of the committee members in any way—all that the sentence requires is that she benefits from it.

Beneficiary arguments are also distinguishable from the adjunct-like doxastic *for/to*-PPs that we discussed earlier. Recall that those PPs were doxastic operators, equivalent to *according to x* or *in x's opinion*. If beneficiary arguments were adjunct-like doxastic operators, then the following constructions should sound contradictory (contrast (5.40) with (5.41)—*unethical* does not take this type of argument):

- (5.40) a. The members of the committee were generous to Ann although she doesn't think so.
 b. Bill's friends will be loyal to him although he doesn't realize it.
 c. Mary's advice is important for Sue although she disagrees.
- (5.41) # Torture is unethical for me although I don't think so.

Beneficiary arguments have hardly been noticed in the literature on evaluatives (but see Stojanovic 2016, pp. 19–20 and Schaffer 2011, p.196, n.24). A divergence in (implicit) beneficiaries could be another reason why participants might be judging that at least some of the dialogues in our experimental items are subjective disagreements. This is different from a divergence in plans, or moral standards, or anything like this. Because two speakers could have exactly the same set of norms in mind when they are making a value judgment and yet be thinking of different beneficiaries. This might be especially pertinent when the disagreement involves thin evaluatives like *good* and *bad*. We leave further considerations about beneficiary arguments for future work.

5.5 Recap & conclusion

In this chapter, we have discussed Solt's 2018 experimental paradigm to test the ordering subjectivity of adjectives. We noted that Solt did not test for broadly moral adjectives, which were an interesting class *vis-à-vis* her account of why certain adjectives are clearly subject to ordering subjectivity. Our results confirmed that moral adjectives fall in this class as well. This called for an enrichment of the notion of judge-dependency in order to make room for the possibility that, beyond having conflicting experiences, perceptions or emotions, speakers can have conflicting priorities or preferences for action. We refer the reader back to Chapter 3-4 for how this can be formally modelled, and in the following chapter we will see how a hyperplan-based semantics can help explain how evaluative sentences behave in conversation.

Appendix: experiment materials

This appendix provides the list of critical and control items used in the experiment.

- Critical items:

- (1) a. What Rose did is morally better than what her boss did.
b. No, I disagree, what her boss did was better.
- (2) a. The previous regime was more moral than the government that was established after the coup.
b. I disagree, the new government was more moral.
- (3) a. Bill leads a more ethical lifestyle than Amy.
b. I disagree, Amy's lifestyle is more ethical
- (4) a. The current CEO of that company behaved more virtuously than the previous one.
b. I disagree, the previous CEO was more virtuous.
- (5) a. General McAdam was more coward than General Smith in that particular battle.
b. No, General Smith was more coward.
- (6) a. The members of the committee were more generous than the invited participants.
b. No, the invited participants were more generous!
- (7) a. Bill's friends are more loyal than Mary's.
b. No, Mary's friends are more loyal.
- (8) a. The treasurer turned out to be more honest than the president.
b. I don't agree with you, the president was more honest.
- (9) a. It's more important to follow Mary's than Bill's advice.
b. I don't agree with you, Bill's advice is more important.
- (10) a. What she did is more justified than what he did to her.
b. I don't agree, what he did is more justified.
- (11) a. Jane acted more rationally than Joe in the aftermath of the tragedy.
b. I disagree, Joe was more rational than Jane.
- (12) a. Everyone knows that Ann's work is more valuable than Jim's.
b. I don't agree with you, Jim's work is of far more value.

• Control items (taken from Solt 2018):

- (1) a. Frank is taller than his friend Jimmy.
b. No, Jimmy is the taller one.
- (2) a. The old Ipad is more expensive than the new one.
b. No, the new one is the more expensive device.
- (3) a. The movie theater is emptier today than it was yesterday.
b. No it isn't. It was emptier yesterday.
- (4) a. The wine bottle is fuller than the champagne bottle.
b. No, the champagne bottle is fuller.
- (5) a. The vegetable soup is saltier than the chicken soup.
b. No, the chicken soup is saltier.

- (6) a. The green towel is wetter than the red one.
b. No, the red towel is wetter.
- (7) a. Caryl and Tina both have blond hair, but Caryl's is lighter than Tina's.
b. No, Tina's hair is definitely lighter than Caryl's.
- (8) a. I would rather use the yellow pillow—it's harder than the white one.
b. No, the white pillow is the harder one of the two.
- (9) a. The cake is tastier than the chocolate cookies.
b. No, the chocolate cookies are tastier.
- (10) a. Those sneakers are uglier than the sneakers you tried on earlier.
b. No, the earlier sneakers were uglier.
- (11) a. This cat is happier than that dog.
b. No, the dog is the happier one of them.
- (12) a. The math professor is more intelligent than the physics professor.
b. No, I disagree. The physics professor is the more intelligent one.

Chapter 6

The Dynamics of value judgment

Summary¹

Evaluative sentences are by default used to make attributions of value. In the appropriate setting however, some evaluative sentences can also be used to express judgments of fact. In this chapter, I argue for a *dynamic* implementation of the non-factualist framework presented in previous chapters as offering the best account of such double uses of evaluative sentences.

6.1 Introduction

In previous chapters, I have defended that evaluative expressions in general, and adjectives in particular, have a non-factualist semantics. Roughly, this means that when we use evaluative language we are not describing the world, but rather, we are engaging in a kind of practical deliberation; we are not discussing how the world is, but how we want it to be and how we plan to behave, in order to make it so.

I have not, however, been very explicit about how using evaluative language can do this; that is, how asserting an evaluative sentence at some context impacts the conversation. In Chapter 3, we gave a semantics for a propositional language, \mathcal{L}_V , containing evaluative operators \uparrow and \downarrow (as well as the binary operators \geq , $>$ and \approx) such that they operated on hyperplans, which were assumed to be part of the intensional models relative to which expressions are evaluated for truth. And in Chapter 4 we saw how to apply this type of model to natural language evaluative adjectives. But we have not really said much about the philosophical significance of these models, nor of how evaluatives behave in conversation. The purpose of this chapter is to amend this. We start from the idea—already hinted at in §3.2—that *hyperplans represent practical commitments*; and that we can model the conversational impact of uttering evaluative sentences as effecting changes in a parameter on the conversational common ground that stands for the practical commitments shared by speakers.

That is the general picture that will be defended; but this chapter has also a more specific empirical purpose, namely to show that certain evaluative sentences can also be used to impart descriptive information, just like non-evaluative sentences. It follows that the full semantic story for evaluative language needs to combine the hyperplans semantics explored in previous

¹This chapter has benefitted greatly from extensive discussions with Carla Umbach and Mora Maldonado. Any errors are of course mine.

chapters with a traditional intensional semantics, in which evaluative sentences are associated with sets of possible worlds. We will see how to do this in the coming sections.

Let us start, then, by challenging the following two intuitively attractive pieces of common wisdom about evaluative sentences:

- Piece of common wisdom #1: evaluative sentences are used by speakers to make *attributions of value*, that is, to make statements whereby they express their sentiment or opinion (which, remember, need not be outright positive or negative, cf. §4.4.2).

This is an example:

(6.1) Tina's is a good car.

- Piece of common wisdom #2: speakers uttering evaluative sentences *fail* to make value attributions only in very special circumstances. To wit: (i) when those sentences are embedded under certain operators; (ii) when they are used *exocentrically* (that is, when the speaker is taking someone else's perspective, Lasersohn 2005) or (iii) when used in free indirect speech.

These are some examples:

(6.2) According to Patti, Tina's is a good car.

(6.3) (*Context: Polly, talking about her daughter Tina's new toy car*) Tina's a good car, the kids in her class are all super jealous.

(6.4) Nina thought about it for a minute and took her phone out to call Tina. After all, Tina's was a good car, and the price was just right.

In contrast with (6.1), the speakers in (6.2)-(6.4) do not express their own sentiment or opinion about Tina's car by using those sentences. Rather, they explicitly report or convey the opinion or stance of someone else.

Consider now the following quote from Hare's *The Language of Morals*:

If I say to someone '*M* is a good motor-car', and he himself has not seen, and knows nothing of *M*, but does on the other hand know what sorts of motor-car we are accustomed to call 'good' (knows what is the accepted standard of goodness in motor-cars), he undoubtedly receives information from my remark about what sort of motor-car it is. He will complain that I have misled him, if he subsequently discovers that *M* will not go over 30 m.p.h., or uses as much oil as petrol, or is covered with rust, or has large holes in the roof. His reason for complaining will be the same as it would have been if I had said that the car was red and he subsequently discovered that it was black. I should have led him to expect the motor-car to be of a certain description when in fact it was of a quite different description (Hare 1952, p. 113, quoted relevantly in Umbach 2016, p. 147).

Hare is considering the possibility of using (6.1) in a purely descriptive way, that is, to make a statement of fact (similar to saying *M is a red motor-car*) and not a judgment of value—against piece of common wisdom #1. And he is saying that one can do this without embedding (6.1), adopting someone else's perspective or using free indirect speech—against piece of common wisdom #2.

Moreover, Hare claims that such a descriptive use of an evaluative adjective like *good* arises in a particular contextual setup: a situation in which the audience ‘has not seen, and knows nothing of *M*, but... ..knows what is the accepted standard of goodness in motor-cars’. In contemporary parlance, a descriptive use requires that the properties of car *M* are not *common ground*, but a certain standard of goodness in motor-cars is.² Hare, however, does not describe what the common ground would have to be like in order for that sentence to have its regular, evaluative use. Nonetheless, it seems that if we “flip” the common ground, so that the properties of *M* are shared while the accepted standard is not, we obtain the usual situation in which an utterance of a sentence like *M is a good motor-car* expresses a judgment of value.

- (6.5) a. (The standard for motor-cars is common ground; the properties of *M* are not):
M is a good motor-car. *≈ descriptive*
 b. (The properties of *M* are common ground; the standard for motor-cars is not):
M is a good motor-car. *≈ evaluative*

That contrast is structurally identical to that observed in C. Barker 2002 between *metalinguistic* and *descriptive* uses of dimensional adjectives like *tall*.³ As Barker notes, while dimensional adjectives are most commonly used to impart information about the world (in this case, about people’s heights), they can also be used to impart *metalinguistic* information, that is, information about the how to use the word *tall*.

- (6.6) a. (The threshold for tallness is common ground; Feynman’s height is not):
 Feynman is tall. *≈ descriptive*
 b. (Feynman’s height is common ground, the threshold for tallness is not):
 Feynman is tall. *≈ metalinguistic*

The contextual setup that gives rise to Barker’s descriptive use of *tall* is almost identical to the setup that gives rise to Hare’s descriptive use of *good*. Conversely, the contextual setup that gives rise to what Barker’s *metalinguistic* use of *tall* is identical to the kind of context that, we hypothesize, gives rise to the standard, evaluative use of *good*. It would appear suggestive, then, to treat evaluative uses of *good* as a metalinguistic use of language. Nonetheless, I will argue that metalinguistic uses of dimensional adjectives and evaluative uses of evaluative adjectives, while differing in a similar way from descriptive uses, cannot be assimilated to one another, for the following reason: evaluative uses of evaluative adjectives update the practical commitments of speakers, while metalinguistic uses do not—they simply introduce constraints on language use. And while a constraint on the *use* of language can be considered a type of practical commitment, it does not go beyond language. The practical commitments mobilized by evaluative uses of evaluative expressions, by contrast, are commitments about how to behave, not just about how to speak.

To sum up, I set myself two tasks in this final chapter. The first and foremost is follow Hare and Barker’s lead and explore the possibility of using evaluative sentences to make statements of

² To prevent a possible confusion: I use the term *standard* to mean the variable orderings of objects that constitute the extension of an evaluative adjective at some index of evaluation; and I will reserve *threshold* for the variable point on a scale that licenses the positive form of a gradable adjective. In view of the distinction between ordering- and positive form-subjective adjectives from §2.2, one could say the following: ordering subjective predicates (like *good*) are so in virtue of allowing for variable, negotiable standards; while positive-form subjective predicates (like *tall*) are so in virtue of allowing for variable, negotiable thresholds.

³ C. Barker 2013 applies the paradigm in C. Barker 2002 to predicates of personal taste.

fact. I will argue that the descriptive usage of evaluative adjectives, far from exotic or derivative, is perfectly natural. This is seen clearly once the mechanics of the contexts in which such uses arise are spelled out. It follows that—and this is the main empirical claim of this chapter—evaluative sentences have evaluative, descriptive and mixed uses. That is to say, such sentences can be used by speakers to express *judgments of value* as well as *judgments of fact*.⁴

My second purpose is to offer a dynamic model of communication that can accommodate evaluative, descriptive and mixed uses of evaluative sentences. The model is essentially an extension of the standard Stalnaker-Heim-Barker framework of dynamic semantics, enriched with a *hyperplan parameter* representing the common practical commitments adopted by interlocutors (Yalcin 2012, 2019). The proposal, in a nutshell, is to take the semantic structure discussed in Chapter 3 and Chapter 4 and interpret it dynamically.

In that structure, there was a set of possible worlds and a set of hyperplans. The idea is, simply, to interpret those parameters as representing, respectively, the shared factual information and the shared practical commitments of participants in the conversation. Informally: evaluative and descriptive uses of evaluative adjectives are understood as different types of update on the *common ground*, understood in turn as containing at least a set of open epistemic possibilities (the familiar *context set*) and the aforementioned hyperplan parameter. In their evaluative use, sentences containing evaluative adjectives update the latter parameter, changing the plans adopted by interlocutors. In their descriptive use however, those same sentences operate exclusively on the former parameter, just like any other descriptive claim.

The chapter is structured in the following way: §6.2 presents and defends the basic contrast between descriptive and evaluative uses of evaluative adjectives. It is argued that descriptive uses of evaluative sentences can be challenged by targeting speaker's evidence (§6.2.1) and can be the subject of lies and misleading statements (§6.2.2); moreover, they do not appear to be subject to subjective or faultless disagreement (§6.2.3) and cannot alter the practical commitments of interlocutors (§6.2.4)—all of which make them more similar to traditional ascriptions of fact than ascriptions of value. §6.3 discusses the hypothesis that evaluative uses of evaluative adjectives are metalinguistic uses of language. §6.4 presents the main ingredients of the dynamic model that can accommodate the linguistic phenomenon observed in §6.2. §6.5 makes a first attempt at offering an account, but a refinement is necessary, as the initial approach does not accommodate descriptive uses of evaluative sentences. In §6.6 we fix this, and introduce a general dynamic entry for evaluative adjectives (§6.6.1) that can account for evaluative (§6.6.2) and descriptive (§6.6.3) uses of evaluative adjectives. §6.7 concludes.

6.2 Evaluative and descriptive uses of evaluative sentences

Suppose that two people, Mora and Amir, are having a conversation about cars in an auto-show. Now consider the following four variations on the context in which Amir could use roughly the same evaluative sentence (containing the evaluative adjective *better*):

- (6.7) a. (**Strangers**): Amir and Mora have never talked to each other and have absolutely no idea about each other's taste in cars; Amir is on the phone telling Mora about

⁴ As in previous chapters, I speak interchangeably of making a value/fact *attribution* or *ascription*, which is a speech act, and expressing a *judgment* of value/fact, which is a mental act.

the cars that he is seeing at an auto show.⁵

Amir: The Mercedes is better than the Audi.

- b. (*Vis-à-vis*):⁶ Amir and Mora have never talked to each other and have absolutely no idea about each other's taste in cars; they are standing in front of a set of cars from the auto-show, whose specs they have carefully studied.

Amir: The Mercedes is better than the Audi.

- c. (**Experts**):⁷ Amir and Mora are car experts, they have a very similar standard for cars and they know that they do; Amir is telling Mora on the phone about the cars he is seeing at the auto show.

Amir: The Mercedes is better than the Audi.

- d. (**Sharing-is-caring**): Amir and Mora are car experts, they have a very similar standard for cars and they know that they do; they are standing in front of a set of cars from the auto-show, whose specs they have carefully studied.

Amir: The Mercedes is better than the Audi.

A comment about each of these contexts: first, in (**Strangers**), Amir's utterance might be thought to be somewhat odd. After all, there is a sense in which there's hardly any common ground between Amir and Mora. When Amir speaks, Mora doesn't know what kind of cars Amir likes nor what the cars that he is referring to are like. So it is uncooperative for Amir to use a massively underspecified term like *better*, knowing that his interlocutor will gain very little from it. In that context, it would be natural for Mora to either reject Amir's utterance or ask for some kind of clarification (*What do you mean, better?*). On the other hand, it is natural to think that even though Amir and Mora do not actively share information about their tastes and the cars, *world* or *background knowledge* restricts the common ground: on the one hand, cars in auto-shows tend to have certain standard qualities: they're (e.g.,) fast, beautiful, reliable and new. On the other hand, those are also the type of properties that people tend to appreciate in cars. Of course, there is a lot of uncertainty, but as Hare says, it is unlikely that people value large holes on the roof as a desirable property of cars. So even when there is no previous common ground, background knowledge restricts the available options. As I will argue, in (**Strangers**) Amir makes a mixed use of an evaluative adjective. Mixed uses expresses neither purely evaluative nor purely descriptive information. Moreover—and perhaps more surprisingly—Amir's utterance does not offer *any* information about the cars themselves nor about the car standards of Amir. It merely offers information about how Amir will behave *vis-à-vis* the cars. Put differently: in (**Strangers**), Mora will not learn anything new about the cars nor about Amir's taste in cars, but she can learn something about Amir, namely, that he will choose the Mercedes over the Audi.

When that same sentence is uttered in any of the other contexts, it sounds more natural. Nonetheless, there are important differences between each of these contexts: most relevantly, I'm going to argue that in (*Vis-à-vis*) Amir is expressing a *judgment of value*; while in (**Experts**) Amir is expressing a *judgment of fact*⁸

⁵ Why would Amir be describing the scene on the phone to a stranger? This is not something that we need to worry excessively about, but we can think of possible situations: suppose that Amir is reporting for a newspaper on the auto-show, and happens to be talking to a newly hired editor. Or perhaps he is having a panic attack and is calling his therapist, but there is someone else on the emergency line and she is asking Amir to describe his surroundings, in order to calm him down.

⁶After Umbach 2016, p. 144's dialogue (15).

⁷After Umbach 2016, p. 144's dialogue (16).

⁸ Note too that in (*Vis-à-vis*), Amir could have used that sentence to express a purely *subjective* judgment about

Finally, it is interesting to note that in (**Sharing-is-caring**) Amir’s utterance feels somewhat redundant. Given that speakers share all the relevant information, Amir’s utterance cannot be taken as an attempt to offer more information about the cars nor to establish a new standard for goodness in cars. It feels as though Amir’s utterance in (**Sharing-is-caring**) is a way for him to confirm that him and Mora are *on the same page* (this is reminiscent of Veltman’s ‘test’-semantics for epistemic modals, Veltman 1996; Yalcin 2007).

In what follows, we will focus on the contrast between (**Vis-à-vis**) and (**Experts**). (**Vis-à-vis**) mimics Barker’s setup for a metalinguistic use of *tall*, while (**Experts**) is similar to Hare’s setup for a descriptive use of *good*. The two axis of variation that this contrast relies on is (i) whether Amir and Mora share a standard for good cars and (ii) whether they share information about the cars (whether they have both studied the cars’ specs). When neither a standard nor factual information is shared, speakers make a MIXED use of evaluatives. When factual information is shared, but a standard is not, speakers make an EVALUATIVE use of evaluatives. When a standard is shared but factual information is not available, speakers make a DESCRIPTIVE use. And finally, when speakers share a standard as well as all relevant factual information, evaluatives are deployed as a TEST, to make sure that speakers are on the same page. We represent this in Table 6.1:

	Different information	Same information
Different standard	(Strangers): MIXED	(Vis-à-vis): EVAL
Same standard	(Experts): DESC	(Sharing-is-caring): TEST

Table 6.1: The different contextual setups in (6.7).

The difference between expressing a judgment of fact and a judgment of value might not be obvious, and therefore the contrast that I aim to draw between the (**Vis-à-vis**) and (**Experts**) contexts might not be immediately apparent. Let us try to establish it more clearly by considering four features that tear apart the different “conversational profile” of utterances expressing judgments of fact and utterances expressing judgments of value. As we will see, those features do in fact tell (**Vis-à-vis**) and (**Experts**) apart (see Ruiz and Stojanovic 2019 for a more extensive discussion of this phenomenon).

6.2.1 Epistemic status

Statements of fact can be challenged and supported by evidence. Attributions of value cannot, as they express a speaker’s more or less direct appreciation, and it would be *prima facie* strange for a hearer to question the speaker’s evidence for them. We can test whether an utterance has this property by considering the appropriateness of challenging that utterance with a claim that targets the speaker’s evidence. When a sentence clearly expresses a judgment of value, that kind of reply is odd, but when a sentence clearly expresses a judgment of fact, it is natural.

his standard for cars, without aiming to obtain Mora’s agreement. Under that interpretation, Amir’s statement would have been equivalent to an utterance of that same sentence embedded under a doxastic operator making explicit reference to himself, his point of view or his opinion:

- (i) To me, / In my view, / In my opinion, the Mercedes is better than the Audi.

In what remains of this paper, we will not focus on this “subjective” interpretation of Amir’s utterance.

- (6.8) a. Milica: The new Star Wars movie is soo boring.
 Camila: ?? Are you sure? / ?? I doubt it. / ?? That cannot be.
 b. Milica: The new Star Wars movie is 140' long.
 Camila: Are you sure? / I doubt it. / That cannot be.

Let us formulate that property thus:

Definition 21 (Epistemic status) *An utterance's status is epistemic just in case it is natural to challenge it by targeting the speaker's evidence.*

The utterance in (*Vis-à-vis*) does not have epistemic status, while the utterance in (**Experts**) does. To see this, consider the following continuations:

- (6.7b) (*Vis-à-vis*):
 Amir: The Mercedes is better than the Audi.
 Mora: ?? Are you sure? / ?? I doubt it. / ?? That cannot be.⁹
 (6.7c) (**Experts**):
 Amir: The Mercedes is better than the Audi.
 Mora: Are you sure? / I doubt it. / That cannot be.

In (*Vis-à-vis*), in light of the fact that Mora and Amir have different evaluative standards for cars and the fact that they both have seen the cars, it is odd for Mora to challenge Amir's evidence for evincing what looks like Amir's opinion. But that reply is significantly more natural in (**Experts**), where they both share a standard for cars. Of course, we need to suppose that Mora has independent reason to distrust Amir's statement about the Mercedes. If that is so, then it makes sense for her to challenge his utterance by targeting his evidence. We submit that Amir's utterance in (**Experts**) has epistemic status, while his utterance in (*Vis-à-vis*) does not.¹⁰

6.2.2 Deceit potential

Secondly, statements of fact are such that external evidence may reveal the speaker to have lied or misled their audience. Utterances expressing judgments of value do not have this property, as again, external evidence can hardly speak against a personal appreciation. The presence of this property can be tested by considering the appropriateness of accusing the speaker of lying in light of evidence that makes their claim false from the hearer's point of view. If that is the case, then it makes sense for the hearer to challenge the speaker's claim by accusing them of lying. If the utterance does not have deceit potential, the accusation is off:

⁹ This test is used in Umbach 2016 to distinguish descriptive and metalinguistic uses of aesthetic adjectives, i.e. *beautiful*.

¹⁰ As Neftalí Villanueva points out (p.c.), this feature of evaluatives places them close to avowals, that is, sentences like *I am hungry* or *I'm wondering if John will arrive on time*. One of the main features of these sentences is that they are considered *groundless*, in that the speaker is never required to offer further evidence for them. Evaluatives uses, I argue, are similar in this respect. However, this does not mean that one cannot discuss *why* one makes a certain value judgment. We do that all the time. But to go into the reasons for an evaluation is not to go into the evidence behind it. In other words, the reasons for an evaluation are not about how the world is (or how we know it to be). Most often, evaluative reasons link particular evaluations with general evaluative standards (e.g., *This movie is bad because it is a propaganda movie (and propaganda movies are always bad)*). Evaluative reasons do not seem to leave the evaluative realm.

- (6.9) a. Milica: Eating animals is wrong.
 Camila: ?? That's a lie / ?? That's misleading.
 b. Milica: Eating animals is widespread.
 Camila: That's a lie / That's misleading.

We can formulate this property as follows:

Definition 22 (Deceit potential) *An utterance has deceit potential just in case new information about the world can reveal the speaker to have lied or misled her audience.*

This was part of Hare's argument for saying that the predicate *good motor-car* could have a descriptive use: in a context in which speaker and hearer have roughly the same standard about cars and know that they do—a context like (**Experts**)—if the speaker says that a car is good and then the hearer finds out that the car fails to meet the 'accepted' standard, the hearer can accuse the speaker of lying or being misleading. On the other hand, if speaker and hearer do not share a standard about good cars—as in (*Vis-à-vis*)—it would be strange for the hearer to make that accusation.

To see this, suppose that Mora disagrees with Amir in (*Vis-à-vis*). It would be weird for her to accuse him of lying or being misleading. But now suppose that, in (**Experts**), Mora eventually gets to study the cars and turns out to disagree with Amir's claim that the Mercedes is better. In that case, it does make sense for her to accuse Amir of lying or being misleading:

- (6.7b) (*Vis-à-vis*):
 Amir: The Mercedes is better than the Audi.
 Mora: ?? That's a lie. / ?? That's misleading.
 (6.7c) (**Experts** + Mora has studied the cars):
 Amir: The Mercedes is better than the Audi.
 Mora: That's a lie. / That's misleading.

We conclude that Amir's utterance in (**Experts**) has deceit potential, while his utterance in (*Vis-à-vis*) does not.¹¹ It is important to stress that the point here is *not* that the utterances in (*Vis-à-vis*) cannot be replied to by saying things like 'that's a lie'. The point is rather that such accusation could be justified by appealing to external information about the object under evaluation only in (**Experts**) and not in (*Vis-à-vis*).

6.2.3 Subjective disagreement

Thirdly, as discussed at more length in §2.2.1 and in Chapter 5, attributions of value easily give rise to so-called FAULTLESS or SUBJECTIVE DISAGREEMENTS, that is, disagreements in which no party appears to be making any kind of mistake and where the matter under discussion does not depend, or at least not obviously, on some objective fact of the matter.¹² By contrast,

¹¹By way of curiosity, the oddness of accusing someone who is voicing their opinion of lying is immediately reminiscent of Bob Dylan's famous [reply](#) to a heckler calling him 'Judas' after "going electric":

Heckler: Judas!
 Dylan: I don't believe you... you're a liar!

¹² See n.2 on Chapter 2 for references on faultless/subjective disagreement.

statements of fact do not in general give rise to subjective disagreements; that is, in most cases it is assumed that either speaker is mistaken and there is a determinate truth of the matter that decides who is right. The disagreement between Milica and Camila in (6.10a) does not seem subjective, while the disagreement in (6.10b) does:

- (6.10) a. Milica: The Eiffel Tower is taller than the Shard in London. *≈ not subjective*
 Camila: No, the Shard is taller.
 b. Milica: Andouillette is tastier than haggis. *≈ subjective*
 Camila: No, haggis is tastier.

We can formulate this property in the following way:

Definition 23 (Subjectivity) *An utterance is subjective just in case a disagreement with it appears prima facie subjective.*

To see this, consider a disagreement dialogue in the first context:

- (6.7b) (*Vis-à-vis*):
 Amir: The Mercedes is better than the Audi. *≈ subjective*
 Mora: I disagree.

On what grounds might Mora be disagreeing with Amir? Assuming that both her and Amir have access to the same information about both cars, it is most natural to think that their disagreement turns on what standard to adopt, Mora's or Amir's. And in principle, both of them seem equally justified in evaluating relative to their standard (see C. Barker 2013, 242, for a similar point), and for this reason, the disagreement between them appears subjective.

By contrast, consider a similar disagreement in (**Experts**), once that Mora learns about the cars:

- (6.7c) (**Experts** + Mora has studied the cars):
 Amir: The Mercedes is better than the Audi. *≈ not subjective*
 Mora: I disagree.

Again, what might ground Mora's dissent in this context? It is reasonable to think that Mora is *either* challenging her and Amir's common standard, *or* she has independent information about the cars, and she is using that information to challenge Amir's claim. In the first case, the disagreement would appear subjective, but the context would no longer be one in which Mora and Amir share a standard, since Mora would be trying to change the standard with her utterance. The second option seems more natural; but then, it seems that their disagreement would turn on the properties of the cars and not on what standard to adopt. Therefore, their disagreement would not be subjective.¹³

¹³ Note that, up to here, we have been using *subjective* to pick out a class of adjectives (those that we called lexical or ordering-subjective in §2.2.1; Chapter 5) or even more specifically, particular forms of certain adjectives (cf. positive form subjectivity discussed in §2.2.1). Subjectivity might have appeared to be a property of certain adjectives, or of particular forms of those adjectives. Now, by contrast, we speak of utterances, not words, being subjective. Are we being equivocal? No: as we said before, for *difficult* to be subjective is for unembedded sentences containing the word *difficult* to be *liable* to appear in disagreement dialogues that would naturally be characterized as subjective (perhaps by one of our participants in Chapter 5). Therefore, *pace* our previous usage, subjectivity is not a property of words—it is a property of certain standardized situations in which people disagree, or seem to disagree about the extension of those words, such as (*Vis-à-vis*). When we turn to non-standard situations containing the same dialogues, such as (**Experts**), the subjectivity fades away, which suggests that subjectivity was not, after all, a property of the words that we have been calling “subjective”.

One might object, however, that in contexts like (**Experts**) where interlocutors share a standard, they can still have a subjective disagreement. This might be because the standard is not fully determined, in the sense that it does not apply to all potential objects of evaluation; or because the objects under consideration are borderline cases for the application of the common standard.

However that is, we can control for these situations simply by assuming that the relevant standards are not underdetermined nor vague in the way just described. That is, we need to assume that for any pair of cars $\langle a, b \rangle$ and evaluative standard e , it is a determinate matter whether the pair $\langle a, b \rangle$ is in the extension of the relation *better than* relative to standard e . If this qualification is built into all contexts, then only disagreements in which the standard is not shared would appear subjective (this is equivalent to assuming that each standard about cars induces a *total order* on its domain, which might be taken to be unrealistic. Nonetheless, for our purposes it is a harmless idealization).

6.2.4 Action-guidance

Lastly, attributions of value alter the *practical commitments* of interlocutors, while statements of fact do not. This occurs in virtue of the fact, discussed at length in §2.3.1, that evaluative adjectives are ACTION-GUIDING. To see this, consider what happens in a context where (6.11a) and (6.11b) are accepted, respectively.

- (6.11) a. It is cruel to hide your office mates' keys. ↪ practical commitment
 b. It is uncommon to hide your office mates' keys. ↪ practical commitment

If interlocutors accept (6.11a), they are adopting a certain practical stance against performing the kind of action under evaluation, namely hiding office mates' keys. The attitude is one of rejection or avoidance, in virtue of the fact that *cruel* is a negative evaluative adjective. One way to see this is that there would be a strong incoherence if interlocutors who agreed on (6.11a) still went on to hide their office mates' keys. In contrast to this, if (6.11b) were accepted in a conversation, there is no sense in which interlocutors would be adopting any practical stance towards that action—thinking that (6.11b) is true is compatible with any practical attitude that one might have towards that action.

Let us call this the ACTION-GUIDANCE of an utterance:

Definition 24 (Action-guidance) *An utterance is action-guiding just in case its uptake changes the practical commitments of participants in the conversation.*

The capacity to alter the practical commitments of interlocutors is a feature that characterizes evaluative language generally, and it has long been observed and discussed by philosophers working in metanormative theory. The production and uptake of value ascriptions involves the adoption of practical commitments on the part of interlocutors, that is, commitments about how interlocutors are disposed to act.¹⁴

With respect to our examples, the observation is the following: in (*Vis-à-vis*), the practical commitments of interlocutors change as a result of Mora's uptake of Amir's utterance; in (**Experts**) however, the practical commitments of the interlocutors remain unaltered after Mora's uptake of Amir's utterance.

¹⁴ See §2.3.1 in Chapter 2.

Let us elaborate. If in (*Vis-à-vis*) Mora is convinced by Amir's statement that the Mercedes is better than the Audi, then she adopts a commitment to orient her action in a way that coheres with her newly acquired standard. This can be spelled out in a variety of ways: Mora might seek to buy a car like the Mercedes, or she will recommend it to other people over cars like the Audi; if given the choice, she will choose to drive or ride on a car like the Mercedes rather than the Audi, etc. These are all commitments that Amir (in (*Vis-à-vis*)) presumably had already, and that cohere with his antecedently held opinion that the Mercedes is better than the Audi. Note that, if Amir had claimed that the Mercedes was bigger than the Audi and Mora accepted that, no practical commitments would be adopted by her as a result of her uptake of Amir's claim.

(6.7b) (*Vis-à-vis*):

Amir: The Mercedes is better than the Audi. ↪ practical commitment

By contrast, if Mora accepts Amir's statement in (**Experts**), she does not thereby acquire any new practical commitment or change the commitments that she already had. We have intentionally set up (**Experts**) in a way such that Amir and Mora share a standard for cars. In virtue of that, Amir and Mora already share certain practical commitments towards cars that meet their shared standard to various degrees. If Mora interprets Amir's utterance as an evaluation that is being made relative to their common standard, she does not change her standard by virtue of accepting Amir's utterance in that context, and therefore, her practical commitments are not altered either.

(6.7c) (**Experts**):

Amir: The Mercedes is better than the Audi. ↪ practical commitment

Before moving on, a potential confusion should be pointed out: in both contexts, the exchange—if Amir's utterance is accepted—results in the adoption, on the part of the interlocutors, of certain practical commitments towards the *very objects* under evaluation. In both contexts, Mora will come to adopt a more positive outlook towards the Mercedes than towards the Audi, if she accepts Amir's utterance. So the contrast that we want to point out here cannot be spelled out in terms of the interlocutors' practical attitudes towards the cars at the end of the exchange, because those attitudes will turn out to be relevantly similar in both contexts.¹⁵ Rather, the contrast between the (*Vis-à-vis*) and the (**Experts**) contexts with respect to those utterances' action-guidance is best fleshed out by considering *the way* in which Mora comes to have that attitude: in (**Experts**), it is the result of antecedently held practical commitments; in (*Vis-à-vis*), that attitude results from accepting Amir's utterance, and in virtue of that coming to acquire new practical commitments.^{16,17}

¹⁵ As Umbach points out (p.c.), another way of seeing this is to consider that the utterances in both contexts have the status of a recommendation (to choose the Mercedes over the Audi). This is the case regardless of whether the recommendation follows from a previously held standard or from a newly acquired one.

¹⁶ As we said with respect to subjectivity, note that there is a contrast between our previous way of talking about action-guidance and the way we are describing the phenomenon now. Before (cf §2.3.1), we described certain adjectives as action-guiding; and now we are claiming that not all uses of those these adjectives are action-guiding—in particular, in contexts like (**Experts**) they are not. It follows, once again, that action-guidance is not a property of certain words, but of specific, standardized uses of those words.

¹⁷ It may seem too strong to say that, when a value judgment remains unchallenged, all participants in the conversation have accepted it. Moreover, as Neftalí Villanueva remarks (p.c.), on occasions it can be rather pedantic to challenge a value judgments (*it is my duty to speak up*). However, that silence amounts to assent or

Let us take stock: we have diagnosed the contrast in conversational profile between the utterances in (*Vis-à-vis*)/(**Experts**) *via* the following four features: EPISTEMIC STATUS, DECEIT POTENTIAL, SUBJECTIVITY and ACTION-GUIDANCE. Just as other sentences that we intuitively take to express judgments of value, the utterance in (*Vis-à-vis*) lacks the first two features but has the last two. By contrast, just as sentences that we commonly take to express factual judgments, Amir’s utterance in (**Experts**) has the first pair of properties and lacks the second. This is shown in the following table:

	Judgment of value	Judgment of fact	(<i>Vis-à-vis</i>)	(Experts)
Epistemic Statuts	×	✓	×	✓
Deceit Potential	×	✓	×	✓
FD-Potential	✓	×	✓	×
Action-guidance	✓	×	✓	×

Table 6.2: The contrast between (*Vis-à-vis*) and (**Experts**).

We take this to show that, while the utterance in (*Vis-à-vis*) expresses a judgment of value that conveys an appreciation on the part of the speaker, the utterance in (**Experts**) expresses a judgment of fact, with the function of conveying purely descriptive information.

What about uses that are neither purely descriptive nor purely evaluative (cf. (**Strangers**))? How do they fare with respect to the dialectical properties just reviewed? It is a little difficult to say, as such mixed uses presuppose that speakers do not share enough information for many of the dialectical moves that we have considered to be acceptable to begin with. For instance: if Amir and Mora do not share any information about the cars in the auto-show nor about their respective standards, epistemic answers such as *I doubt it* will probably be out. This is not, however, necessarily because Amir has expressed an opinion, as in (*Vis-à-vis*). Rather, that reply is out because Mora has no information to challenge Amir’s statement. The same goes for the question of whether Amir’s statement has deceit potential and whether it can lead to a faultless disagreement. However, there is some reason to think that situations like (**Strangers**) will more easily develop into situations like (*Vis-à-vis*), since it’s easier to come to share factual than practical information. To see this, note that it would be natural for Amir to follow up his claim in (**Strangers**) with certain specs of the cars, in order to clarify what he meant by *better*. And then the situation would become much like (*Vis-à-vis*). It would be comparatively more awkward (and arrogant) for him to follow up his claim in (**Strangers**) by saying what he assumes to be his and Mora’s (whom he’s just met) shared standard for goodness in cars. Either way, none of the proposed tests for distinguishing descriptive and evaluative uses can be applied to (**Strangers**) without either of those two clarification moves on the part of Amir.

6.3 Evaluative *versus* metalinguistic use of language

As mentioned in the introduction of this chapter, the contrast between the contextual setup in (*Vis-à-vis*) and (**Experts**) is structurally similar to the contrast observed in C. Barker 2002 between descriptive and metalinguistic uses of dimensional adjectives (repeated here).

uptake is a useful simplification both in the case of statements of fact as well as attributions of value (see Goldberg 2016).

- (6.6) a. (The threshold for tallness is common ground; Feynman’s height is not):
 Feynman is tall. *≈ descriptive*
 b. (Feynman’s height is common ground, the threshold for tallness is not):
 Feynman is tall. *≈ metalinguistic*

It is therefore tempting to assimilate what we have been calling *evaluative* uses of evaluative adjectives to metalinguistic or interpretational uses of language, and this has been the object of recent proposals in the literature about some evaluative expressions, such as PPTs or aesthetic adjectives.¹⁸

As Barker notes, while dimensional adjectives are most commonly used to impart information about the world, they can also be used to impart *metalinguistic* information, that is, information about the discourse. To wit: (6.6a) conveys information about Feynman’s height, namely, that his height is at least equal to the accepted threshold. It is therefore a DESCRIPTIVE use of *tall*; (6.6a) is the kind of answer that we would give to a question about John; for example, *what does John look like?* By contrast, (6.6b) conveys information about the threshold for tallness (namely that it lies somewhere not above Feynman’s height). It is therefore a METALINGUISTIC use of *tall*. This utterance would be appropriate as an answer to a question not about John but rather about word usage, e.g., *what is considered ‘tall’ around here?*

The hypothesis, then, would be to treat the EVALUATIVE use of *better*—cases like (*Vis-à-vis*)—in the same way that Barker treats metalinguistic uses of *tall*: as utterance types that aim at offering information, not about the cars, but about the accepted standard for goodness in cars.

An immediate disparity between Barker’s contrast and ours is that (6.7) involves a comparative. Note that the comparative *taller*, does not admit of metalinguistic uses, as shown by the fact that it cannot be the object of a question about the discourse nor be an answer to such a question:

- (6.12) a. What is considered tall /#taller around here?
 b. # Skłodowska is taller than Feynman.

In contrast with (6.6), this sentence cannot be used to convey information about the accepted threshold for *tall*; it simply compares the heights of the two individuals. Interestingly, adjectives like *good* seem to allow for such uses in the comparative form.

- (6.13) a. What is considered good/better around here?
 b. The Mercedes is better than the Audi.

This might give initial support to the hypothesis that what we have been calling an evaluative use of *better* (cf. contexts like (*Vis-à-vis*)) is a metalinguistic use. An immediate obstacle with this approach, however, is that it is not appropriate to describe Amir’s utterance in (*Vis-à-vis*) as an answer to a question about *actual* word usage, a question like *what is considered ‘good/better’ around here?*. This is simply because, in order to use that sentence (6.7) evaluatively, Amir need not be prompted by any question about language use. He is merely stating a spontaneous opinion about the cars.

Nonetheless, Amir’s utterance could well be taken to be an answer to a different metalinguistic question; not a *descriptive* metalinguistic question about how words are actually used, but a

¹⁸ See C. Barker 2013; Plunkett and Sundell 2013; Sundell 2016; Umbach 2016 *a.o.*; see Karczewska 2016; Marques 2017 for criticism.

normative metalinguistic question about how words *should* be used. In other words, Amir's utterance in (*Vis-à-vis*) could be treated as a move in a kind of METALINGUISTIC NEGOTIATION (Plunkett and Sundell 2013; Sundell 2016).

Let us illustrate what a metalinguistic negotiation is with the following example (Ludlow 2008; Plunkett and Sundell 2013): speakers discussing whether Secretariat, the famous racehorse, is an athlete, would most likely be taken to be negotiating the appropriate concept of athlete rather than the merits of Secretariat. And most importantly, speakers engaging in such discussions could have what looks like a disagreement even when perfectly aware that the sentences that they use express different and therefore compatible propositions, relative to different concepts of athlete. To see this, suppose that Bill and Mary have different concepts of athlete, and they know that they do; and let ATHLETE- stand for Bill's horse-exclusive concept of athlete and ATHLETE+ stand for Mary's horse-inclusive concept. They might have the following exchange:

- (6.14) a. Mary: Secretariat is an ATHLETE+.
b. Bill: What? Of course he is not an ATHLETE-!

The fact that Mary and Bill have different concepts of athlete does not prevent them from having a dispute, even if their sentences express compatible propositions. But it would be a metalinguistic dispute about how to use the word *athlete*, rather than a disagreement about the properties of Secretariat.

In the case of (*Vis-à-vis*), the situation could be similarly characterized. Amir might be using the sentence in (*Vis-à-vis*) not, or not only, with the intention of informing Mora about the cars, but with the intention of getting Mora to adopt his standard for cars. Conversely, if Mora rejects Amir's claim, she need not be denying that Amir's standard is as he describes it; she might be refusing to adopt Amir's standard.

This suggests a different interpretation of the disagreement continuation of (*Vis-à-vis*) (repeated below): in that context, Amir and Mora would not be exchanging information about the cars. Rather, they would be negotiating the meaning of *better*.

- (6.7b) (*Vis-à-vis*): Amir and Mora have never talked to each other and have absolutely no idea about each other's taste in cars; they are standing in front of a set of cars from the auto-show, whose specs they have carefully studied.
Amir: The Mercedes is better than the Audi.
Mora: I disagree.

If we interpret Amir's utterance in (*Vis-à-vis*) in this way, we come a long way to accounting for the contrast between (*Vis-à-vis*) and (*Experts*): first, we have an available account for why a disagreement would seem SUBJECTIVE in (*Vis-à-vis*) but not in (*Experts*) (§6.2.3): in (*Vis-à-vis*), Amir and Mora would be negotiating the standard to be adopted, and there is no *a priori* way of determining whose scale is the right one (as Barker 2013, p. 242 says, 'no individual has privileged access to or authority over linguistic convention'). In (*Experts*) however, a disagreement would not seem subjective because speakers, having settled on a given standard, would be discussing the objective merits of the cars, relative to that common standard.

Secondly, the lack of EPISTEMIC STATUS (§6.2.1) and DECEIT POTENTIAL (§6.2.2) of Amir's utterance in (*Vis-à-vis*) would be similarly accounted for by noting that, if Amir is using that sentence to try to get Mora to adopt his standard, then it does not make sense for Mora to ask

for further evidence about Amir's claim, or to accuse him of lying or being misleading, even if she thinks that the Mercedes is not better than the Audi relative to *her* standard. In (**Experts**) by contrast, it does make sense for Mora to ask for further evidence. And given the assumption that she and Amir share a standard, if she sees that her own assessment of the cars (relative to their common standard) differs from Amir's, she can question Amir in order to get to the bottom of their conflict. If she figures out that Amir was wrong about the properties of the cars or about their presumably common standard, it makes sense for Mora to accuse him of lying or being misleading.

Unfortunately, the metalinguistic approach has difficulties to accommodate the difference in ACTION-GUIDANCE between (*Vis-à-vis*) and (**Experts**) (§6.2.4). To see this, consider the contrast between a regular, descriptive use of Mary's sentence in (6.14) and a metalinguistic use thereof. In its former use, Mary is simply informing her audience that Secretariat falls under her concept of athlete, ATHLETE+. In the latter case, Mary is trying to get her audience to adopt ATHLETE+ as their concept of athlete. But under none of those interpretations does her sentence alter the practical commitments adopted by speakers. Adopting a new concept of athlete does not affect the practical attitudes of speakers towards the individuals that fall under (or outside of) it. At most, adopting such a new concept alters the *linguistic* commitments of speakers—after Mary's utterance, her audience will be committed to calling any horse as good as Secretariat an *athlete* as well. This is true of metalinguistic uses of language in general: any change in the conventions governing language use affect the linguistic commitments adopted by agents. But over and above linguistic commitments, evaluative language—or more specifically, the evaluative *use* of evaluative language—can alter the practical commitments of speakers, that is, the commitments that orient the interlocutors' actions. Metalinguistic uses of terms do not have this feature, which makes it hard for them to account for all the points of contrast between (*Vis-à-vis*) and (**Experts**).

To put it in dynamic terms, where judgments of fact are assertions that aim at reducing our uncertainty about the world, judgments of value would be assertions that aim at reducing our uncertainty about the discourse (C. Barker 2013, p. 243). Now, our observations about the action-guiding properties of value judgments are meant to highlight that the kind of uncertainty and indeterminacy that value judgments target is of a different kind from worldly as well as discursive uncertainty: it is uncertainty and indeterminacy about what to do. That is, it is *practical* uncertainty. Reducing practical uncertainty to discursive uncertainty loses sight of the action-guiding force of evaluative language, and for this reason, we find the metalinguistic view to lack the resources to accommodate our contrast.

Moreover, proponents of the metalinguistic negotiation approach to evaluative and normative disagreement have been aware from the start that there is a contrast between descriptive and normative metalinguistic negotiations, and that standard uses of evaluative/normative vocabulary, if they can be treated as metalinguistic negotiations, belong to the normative kind. We saw this with (*Vis-à-vis*): in that context, it is clearly wrong to conceive of Amir's utterance as an answer to a question about the currently accepted standard for cars; if anything, it is an answer to a question about what standard for cars *ought* to be accepted. But once that such a normative dimension is allowed into the picture, we seem to be going in a circle: in virtue of what might a standard for cars be the standard that *ought* to be accepted as the meaning of *better car* for Amir, other than the fact that it picks out the better cars according to him?

Nonetheless, the metalinguistic approach goes a long way towards an account, as we have seen. Moreover, we have not attempted to distinguish metalinguistic and evaluative uses of evaluative

adjectives, and we shall acknowledge that the task is not easy. Let us be clear about the merits and limits of this proposal, so that our proposal inherits the former and overcomes the latter. Crucially, we want to retain the idea that in (*Vis-à-vis*) Amir is trying to get Mora to adopt his standard. And if Mora accepts Amir's utterance, she will have thereby adopted Amir's standard. In other words, the conversational import of accepting Amir's utterance in (*Vis-à-vis*) is the adoption of a new standard. All this is already captured by the metalinguistic approach. But in order to account for the observation that accepting Amir's utterance in (*Vis-à-vis*) changes the practical commitments of speakers, we must treat the negotiation of evaluative standards as aiming to reduce, not factual nor discursive uncertainty, but the *practical* uncertainty of interlocutors. This task is taken up in the following section.¹⁹

6.4 Dynamic semantics: a précis

The task before us, then, is the following: evaluative sentences need to be assigned semantic values such that they can play the double role of expressing judgments of value and judgments of fact. And we need to make sure that those semantic values interact with context in the appropriate way: in a context where an evaluative standard is shared but worldly information about the object(s) of evaluation is missing, a speaker uttering an evaluative sentence is making an attribution of fact; and in a context where an evaluative standard is not shared but worldly information about the object(s) of evaluation is, uttering the same evaluative sentence expresses a judgment of value. Although we haven't paid much attention to mixed uses in the previous section, the model of conversation to be laid out now offers a clear way to represent their intuitive contribution.

Dynamic semantics holds the key for a proper treatment of this contrast; and in this section, the main ingredients of the dynamic model of communication to be defended are presented. The proposal is an extension of the basic dynamic semantic framework of Stalnaker, Heim and others (see K. Lewis 2017 for an introduction). In traditional, intensional compositional semantics, the semantic value of a declarative sentence at a context are its truth-conditions, that is, the set of actual and counterfactual circumstances under which it is true. In dynamic semantics by contrast, the meaning of a sentence is its *context change potential*, that is, the way any context changes as a result of uttering and accepting that sentence.

Dynamic frameworks in semantics are well-suited to account for phenomena that rely on the interaction of sentences with the context in which they are uttered, understanding the latter in terms of the information it contains. That is how dynamic accounts of, e.g., presupposition or epistemic modals work: sentences carrying presuppositional content require that the context

¹⁹ In support of the similarity between the non-factualist approach defended here and the metalinguistic approach just discussed, note that there have also been attempts in the opposite direction, that is, attempts to subsume phenomena like vagueness under a non-factualist or expressivist model, cf. MacFarlane 2016. MacFarlane's proposal is to use the tools devised by Gibbard to give a semantics, not (only) to normative or evaluative expressions, but also to a large class of vague expressions. In MacFarlane's view, vagueness is treated as a kind of linguistic *indecision*. Linguistically, this means extending the non-factualist or expressivist model that we have been discussing to the positive form of relative adjectives, something which we have not attempted here. Discussing that work would take us to far afield, however, but one *prima facie* reason to resist MacFarlane's proposal lies in our main motivation for a hyperplan semantics for evaluatives, namely the action-guiding properties of these expressions. In our view, action-guidance motivates a hyperplan semantics. But not all vague expressions are action-guiding, which is why we remain skeptical about MacFarlane's proposal of a hyperplan semantics to tackle vagueness.

in which they are uttered entails their presuppositions, otherwise they are infelicitous (Heim 1992, 2002, *a.m.o.*); epistemic modals such as *might* require that the context in which they are uttered is compatible with their prejacent (Veltman 1996; Yalcin 2007, 2011). It is crucial for a dynamic approach to meaning to understand contexts as informational objects, rather than—in the more traditional Kaplanian fashion—as tuples of indices. As such, they represent the state of a conversation at any given point (D. Lewis 1979b).

A dynamic framework appears well-suited to accommodate the contrast between (*Vis-à-vis*) and (*Experts*) as well. We observed that one and the same sentence appears to have a different communicative “profile” depending on the information available in the context in which it is uttered. In (*Vis-à-vis*), interlocutors share the information about the cars (in virtue of the fact that they have studied their specs carefully), but they do not share a standard. As a result of this, Amir’s utterance in that context lacks epistemic status and deceit potential, but it is subjective and action-guiding. In (*Experts*), it is the other way around. The contrast between these utterances that we have attempted to diagnose in §6.2 arises systematically as a result of the kind of information that is available in each occasion of use. For this reason, we take a dynamic framework to offer the best resources to account for this phenomenon.

In this section, we start by reviewing the dynamic account of the context change potential of declarative sentences—in terms of the elimination of epistemic possibilities in the common ground of the conversation—as well as other clause types that have been observed to motivate the introduction of parameters in the common ground beyond epistemic possibilities, such as so-called *To-Do lists*, associated to the context change potential of directives. Subsequently, we will introduce the notion of a hyperplan parameter, and we will see how hyperplans can be connected to the semantics of evaluative adjectives in this dynamic setting.

6.4.1 Assertion, commands and other animals (in dynamic semantics)

Start, as usual, with the notion of COMMON GROUND (G) as representing shared information among the participants in a conversation at any given stage (Stalnaker 1978/2002). First and foremost, the common ground contains a set C of possible worlds, each of them standing for a maximally specific way that the actual world might turn out to be, for all that interlocutors know (and know that they know, and know that they know that they know, etc). This is often called the CONTEXT SET of a conversation.

Definition 25 (Common Ground #1) $G := C = \{w_0, w_1, w_2, \dots\}$

In dynamic semantics, to make an assertion is to eliminate possibilities from the context set. This operation of eliminating possibilities can, in turn, be represented formally in two ways. We can take the semantic values of sentences to just be their context change potentials, and represent their denotations formally as functions from/to common grounds (this is Heim’s approach in her 1983/2002; Heim 1992; also Barker’s in 2002). Alternatively, one can retain a more traditional notion of semantic values as sets of points on a semantic structure, and understand operations on the common ground in terms of interactions between various entities represented in such structure. In order to ensure consistency with previous chapters, we are going to adopt the second route.

Recall what we said at the start of this section: in static semantics, the semantic value of a declarative sentence at a context is standardly taken to be a set of possible worlds. For any declarative sentence φ :

$$(6.15) \quad \llbracket \varphi \rrbracket = \{w : \varphi(w) = 1\}$$

Given that the context set of a common ground is also a set of worlds, it is straightforward to understand the assertion of sentence φ at a common ground G as the result of intersecting the semantic value of φ with the context set C of G :

Definition 26 (Assertion w/o presupposition) $G[\varphi] = C \cap \llbracket \varphi \rrbracket$

If a sentence φ has a presupposition ψ , its assertion is infelicitous unless the context set entails the presupposition; that is, unless updating the context set with the presupposition has no effect:

Definition 27 (Assertion w. presupposition) $G[\varphi]$ is defined only if $C \cap \llbracket \psi \rrbracket = C$; if defined, $G[\varphi] = C \cap \llbracket \varphi \rrbracket$

It is crucial for our purposes to retain the general idea that the common ground of a conversation has a strong impact on the status of the assertion of a sentence. If a sentence carries a presupposition that is not common ground, then its assertion is not in general felicitous (although it will often be accommodated). I plan to exploit this general mechanism to account for the contrast between descriptive and evaluative uses of evaluative sentences: put simply, the properties of the common ground will determine whether an utterance of an evaluative sentence is felicitous, and if it is, whether it turns out to express a judgment of value, a judgment of fact or a mixture of both.²⁰

This model for assertion offers an intuitive grasp on the idea that assertions eliminate possibilities and thereby *augment* the information accepted by interlocutors. But there are more things one can do in a conversation besides eliminating possibilities. For instance, the basic dynamic story for an epistemic modal like *might* treats it as an operation on the context set that, rather than eliminate possibilities, *tests* whether it's prejacent is compatible with it, and if so, leaves the context set untouched (Veltman 1996; Yalcin 2007).

Moreover, when we turn to clause types other than declaratives, the general idea of operating on the context set becomes insufficient. Take the case of imperatives. What is the contextual effect of issuing (and accepting) a command with an imperative clause such as *pass me the salt*? It cannot be a matter of intersecting the context set with the worlds in which *you*, the addressee, *pass me*, the speaker, *the salt*. Because the fact that I issue the command and that you accept it does not guarantee that you obey it. More generally, as Charlow (2014, p. 646) points out, accounts of imperatives that assign to them traditional propositions fail to capture

²⁰ There might be other theoretical perks of retaining a more conservative view about semantic values, where the semantic values of declarative sentences are model-theoretical objects defined against a semantic structure, and context-change potentials represent the pragmatic effect of asserting sentences with those semantic values. As Camp notes (2017b, p. 1621), Stalnaker's (1978/2002, 2014) notion of a common ground has a markedly social and somewhat provisional status: its contents are accepted by participants for the purpose of conversation. The corresponding dynamic view of assertion is one according to which an assertion is an attempt to eliminate possibilities from the common ground. As such, the dynamic notion of assertion inherits the social and provisional properties from the common ground. However, if one thinks of assertion as a markedly *public* act, so that, e.g., asserting is a way of avowing certain public commitments that go beyond a conversation (Brandom 1983), then the social and provisional notion of assertion as common ground update might appear insufficient. This is a reason to retain a traditional static semantics (see also Stalnaker 2014, p. 162, who also supports this type of conservative view about the dynamics of conversation but for different reasons). Nonetheless, in order to keep matters simple, we will continue to speak of assertion in this dynamic sense, but keeping in mind that it could turn out to be an insufficient model to represent the public, committal aspects of assertion and conversation in general.

the action-directed and non-representational usage of directives.²¹

Considerations like these have led some theorists, notably Portner 2004, to populate the common ground with so called *To-Do lists*. To-Do lists are sets of properties that are indexed to participants in the conversation and that list whatever is common ground that they accept as their obligation or plan of action. If one adopts this view, the common ground becomes a tuple containing a context set and a sequence $\langle t^1, \dots, t^n \rangle$ of To-Do lists indexed to each participant:

Definition 28 (Common Ground #2) $G := \langle C, t^1, \dots, t^n \rangle$, where

$$\begin{cases} C = \{w_0, w_1, w_2, \dots\} \\ \text{for every speaker } s, t^s = \{P : s \text{ has to instantiate } P \text{ in all } w \in C\} \end{cases}$$

According to (a possible implementation of) this view, an imperative sentence expresses a singleton set containing a property. For any imperative sentence $X!$,

$$(6.16) \quad \llbracket X! \rrbracket = \{P : P = X\}$$

Uttering an imperative towards an addressee has the effect of adding that property to the addressee's To-Do list. For any addressee a at a conversation with common ground G and imperative sentence $X!$,

Definition 29 (Command) $G[X!] = t^a \cup \llbracket X! \rrbracket$

We don't need to stop in the details of this strategy. The important thing to retain is that various theorists have considered that the common ground needs to be enriched in order to capture the semantic effect that certain fragments of language, in particular certain clause types, conventionally have. Furthermore, these theorists are aware that assigning a different semantic value to a sentence or clause type is often not enough to capture its typical conversational properties. Note that, in order to capture the conversational effect of issuing a command, it does not suffice to assign properties as the semantic value of imperatives. Portner cleverly devises the notion of a To-Do list, which represents obligations. To update one's To-Do list is to acquire a new obligation, an imperatives do just that. To-Do lists have, in this view, a certain normative status.

It is part of this strategy to assume a generalization associating clause types, on the one hand, and update functions on parameters of the common ground, on the other: declaratives update the context set; imperatives update the (addressee's) To-Do list. But there are reasons to think that, at least in English, sentences that fall under a certain clause type have update operations that do not respect that generalization. Consider the case of promises. As Portner (2004, p. 239) notes, languages like Korean have PROMISSIVES, that is, a special clause type for making promises. However, promises in English are expressed *via* declarative sentences headed by the performative verb *promise* in the first person. Yet similar arguments to those that favor introducing the addressee's To-Do list in order to account for the update potential of imperatives apply here with respect to the speaker. To make a promise is, in this view, to update the *speaker's* To-Do list.

Even more interesting for our purpose are EXHORTATIONS. Exhortations are expressed in English *via* sentences like *let's dance*, a grammatically imperative clause that has fossilized into an almost idiomatic form. Portner 2004 notes that Korean also has a dedicated clause type, namely HORTATIVES, and he suggests that their context change potential is to update

²¹ Although see Kaufmann 2011 for a view treating imperatives as performative modal verbs.

both speaker's and addressee's To-Do list. Promises and exhortations in English suggest that, even though the association between clause types and kinds of update on the common ground is a desirable one, it is not written in stone.

I wish to follow this path in order to account for evaluative uses of evaluative sentences. In a slogan, the idea to be defended is that *evaluative uses of evaluative sentences update the practical commitments of interlocutors*, which might be thought of as a common or shared To-Do list of the common ground. Similarly to promises in English, but contrary to commands, to this particular update on the common ground does not correspond a dedicated clause type—evaluative sentences are just unembedded or atomic declarative sentences that contain evaluative expressions. And unlike commands and promises, but similarly to exhortations, evaluative sentences (in their evaluative use) have the conversational effect of updating a set of shared, rather than individual, commitments.

6.4.2 Hyperplans and common ground

What might this “common” To-Do list be? Put simply, it ought to be a list of things for *any* participant in the conversation to do. So, as a first stab, we could represent it as yet another set of properties. If we did that, we would simply be extending Portner's idea of a To-Do list in a conservative way. On top of each participant's personal To-Do list, there would be something like a “global” To-Do list that represents shared practical commitments.²²

This is going to be roughly our approach, but rather than represent the content of such global To-Do list as a set of properties, we will represent it as a set of hyperplans, or as we called them in §3.3.1, as a PLAN.²³ This way, we can deploy the semantic structure that was proposed in Chapter 3-4 (see especially §4.2.2).

Remember that we defined a semantic structure as a set W of possible worlds, a set of hyperplans H and a set of alternative-generating functions A :

Definition 16 (Alternative-enriched semantic structure) $S = \langle W, H, A \rangle$

And recall as well our definitions of hyperplans and alternatives:

Definition 10 (Hyperplans) A hyperplan is a total function h from non-empty sets of propositions to non-empty sets of propositions such that, for every non-empty set of propositions U, V such that $V \subseteq U$, $h(U) = V$ just in case V are the chosen alternatives at U .

Definition 15 (Alternatives) An alternative set is a function a from possible worlds to sets of sets of possible worlds such that, for every world w and for every set of sets of worlds U , $a(w) = U$ just in case every set of worlds in U is a potential outcome at w .

What we propose to do, put simply, is to interpret this semantic structure *dynamically*. To do that, we take the semantic structure defined above, and we define the different parameters of the common ground as subsets of points from each domain of the structure: the context set C

²² The type of commitments and obligations that we engage with when we use evaluative language presumably concern more than just the participants in a conversation: for example, if asserting that killing is wrong involves adopting a practical attitude against killing, this attitude is expected to be adopted by everyone, not just the people that one is talking with. It is not clear how to incorporate this in our account, although nothing I say will hinge on it.

²³ See Charlow 2015; Starr 2016; Willer 2017; Yalcin 2012, 2018, 2019, for similar moves.

of the common ground is represented as a subset of the set W of possible worlds (as one would expect). The set of alternative-generating functions of the common ground will be a subset O (for “options”) of the set of alternative-generating functions A . And finally, the common ground will determine a plan P , which is a subset of the set of hyperplans H , namely those hyperplans that are *compatible with the shared practical commitments of all participants in the conversation*. Such commitments may very well be implicit at the beginning of a conversation, but they are mobilised by evaluative (and surely other types of expressive) language.

For instance, a common ground might be such that it is accepted by all participants in a conversation that it is raining (hence that its context set is a subset of worlds in which it is raining). The common ground can also determine a set of alternative courses of action, that is, a set of options about what to do. For instance, interlocutors might be considering whether to go continue walking in the rain, or go indoors. Or speakers might be considering as options buying an umbrella of different possible sizes and prices to continue their walk in the rain, etc. In addition to all this, the common ground also contains a hyperplan parameter, which contains information about the shared plans of speaker relative to different available information and alternatives. For instance, the common ground might be such that speakers are against continuing their walk if their clothes get soaked, but are undecided about what to do if it continues raining lightly as it has until now.

The common ground now looks like this (now excluding personal To-Do lists):

Definition 30 (Common Ground #3) $G := \langle C, P, O \rangle$, where
$$\begin{cases} C = \{w_0, w_1, w_2, \dots\} \\ P = \{h_0, h_1, h_2, \dots\} \\ O = \{a_0, a_1, a_2, \dots\} \end{cases}$$

The next step is to associate different types of sentences with different updates on the common ground. In order to keep matters relatively manageable, we are going to ignore the alternative parameter of the common ground from now on. We will assume that the common ground supplies a set of alternatives whenever these are needed, but we will not say anything else about how this mechanism works (beyond what was said in §4.2.1).

6.5 Evaluative & descriptive update (1st attempt)

Here is an initial view: *declarative* sentences can update both the context set and the hyperplan parameter of the common ground. But while *descriptive* sentences restrict the context set and leave the hyperplan parameter untouched, *evaluative* sentences restrict the hyperplan parameter, while leaving the context set untouched. In Chapter 3-4, our proposal was to assign all declarative sentences a uniform semantic type that combined a descriptive and evaluative component, and then let descriptive sentences have an idle evaluative component while evaluatives have an idle worldly component. We will do the same here.

In order to interpret declarative sentences dynamically, we need to enrich our view about assertion, so that the assertion of a declarative sentence can operate an update on the context set as well as on the hyperplan parameter of the common ground. We can do this by assuming that assertion is set intersection with those two parameters (we do not assume any update operation on alternatives). Where φ is any declarative sentence, our new rule of assertion is:

Definition 31 (Generalized Assertion) $G[\varphi] = \begin{cases} C \cap \llbracket \varphi \rrbracket \\ P \cap \llbracket \varphi \rrbracket \end{cases}$

The difference between the conversational impact of evaluative and descriptive sentences is due to their semantic values, which we already defined to be respectively insensitive to either worlds or hyperplans and alternatives. Let us go through each case.

Recall, first, our truth conditions for atomic, descriptive sentences (§4.2.2). Where M is a model consisting of a semantic structure S (as characterized in Definition 16) and a valuation function V ; and where φ is any descriptive and atomic (i.e. non-Boolean) sentence of English,

$$(4.16) \quad \llbracket \varphi \rrbracket_{\langle w_i, h_i, a_i \rangle}^M = 1 \text{ iff } \langle w_i, h_i, a_i \rangle \in V(\varphi) \text{ (we henceforth skip } M)$$

The semantic value that V assigns to this sentence is a set of world-hyperplan-alternative triplets.

$$(6.17) \quad V(\varphi) = \{ \langle w, h, a \rangle : \varphi(w)(h)(a) = 1 \}$$

Since φ is descriptive, $V(\varphi)$ contains all hyperplans and all alternative-generating functions, but not all worlds. In other words, (6.17) is hyperplan- and alternative-insensitive, but world-sensitive.

To see how an assertion of this type of sentence impacts the common ground, we need to apply this semantic object to a common ground G , following our rule of Generalized Assertion. Assuming that φ is compatible with a context set C (and assuming that φ is not entailed by G and there is no presupposition failure),

$$(6.18) \quad G[\varphi] = \begin{cases} C \cap \llbracket \varphi \rrbracket \neq C \\ P \cap \llbracket \varphi \rrbracket = P \end{cases}$$

Since φ is a descriptive sentence, it contains all hyperplans but not all worlds. Therefore, the update on the context set C will result in some epistemic possibilities being ruled out. By contrast, the update on P will leave P untouched.

By contrast, asserting an evaluative sentence performs an update on the hyperplan parameter of the common ground, but leaves the context set untouched. Let us see how this work for the most basic type of evaluative sentence that we discussed, namely a sentence containing a thin evaluative adjective taking a full proposition as argument. These are the truth-conditions that we gave in §4.2.3 for a sentence like *it is good that* φ :

$$(4.17) \quad \llbracket \text{it is good that } \varphi \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff } \{ w' : \llbracket \varphi \rrbracket_{\langle w', h_i, a_i \rangle} = 1 \} \in h_i(a_i[\varphi](w_i))$$

That sentence is true just in case the worldly component of φ is preferred by the hyperplan of the index, relative to the set of alternatives of the index (which includes φ).

The semantic value assigned to the sentence in (4.17) is therefore a set of world-hyperplan-alternative triplets.

$$(6.19) \quad \llbracket \text{it is good that } \varphi \rrbracket = \{ \langle w, h, a \rangle : \{ w' : \llbracket \varphi \rrbracket_{\langle w', h, a \rangle} = 1 \} \in h(a[\varphi](w)) \}$$

But because this sentence is evaluative, its semantic value contains all worlds, but not all hyperplans and not all alternatives. Thus, the result of adding this sentence to a common ground (again, ignoring the alternative parameter) is an update to the hyperplan parameter of the common ground.

$$(6.20) \quad G[(6.19)] = \begin{cases} C \cap \llbracket (6.19) \rrbracket = C \\ P \cap \llbracket (6.19) \rrbracket \neq P \end{cases}$$

In particular, adding (6.19) to a common ground G will leave the context set C untouched, but will rule out hyperplans from P that are not such that the worldly component of φ is among the preferable alternatives.

For instance, when a speaker uses an evaluative sentence such as

(6.21) It's good that we didn't call off our walk.

our proposal is that this sentence should be understood as an update of the hyperplan parameter of the common ground, relative to whatever choice of alternatives is determined by the common ground at that point of the conversation. It is a proposal to rule out hyperplans that do not count the proposition *that we didn't call off our walk* among the preferable alternatives.

The question for us is whether we can take on board this proposal and apply it to our cases, that is, to evaluative sentences such as the one uttered by Amir in the different scenarios of (6.7):

(6.7) The Mercedes is better than the Audi.

Bracketing for now the semantic features of this particular type of sentence (we come back to them in the following section), all we would need to do is associate (6.7) (whose semantic value would also be a set of world-hyperplan-alternative triplets) with an update function. And since (6.7) is evaluative, its semantic value would be hyperplan- and alternative-sensitive, but world-insensitive. Therefore, its update function would target the hyperplan parameter of the common ground, while leaving the context set untouched.

$$(6.22) \quad G[(6.7)] = \begin{cases} C \cap \llbracket (6.7) \rrbracket = C \\ P \cap \llbracket (6.7) \rrbracket \neq P \end{cases}$$

This proposal would be immediately well-suited to account for *evaluative* uses of evaluative sentences (*cf.* (**Vis-à-vis**)): since the function of evaluative sentences (in their evaluative use) is to update the hyperplan parameter (rather than add information to the context set), it is to be expected that utterances of those sentences lack EPISTEMIC STATUS and DECEIT POTENTIAL, as we saw in §6.2.1-6.2.2. Moreover, disagreeing utterances might be taken as proposals to adopt conflicting plans (that is, different sets of hyperplans), which coheres with the idea that no plan of action is *a priori* superior to any other, thereby accounting for the intuition of SUBJECTIVENESS (§6.2.3). Finally, taking evaluative sentences to update the hyperplan parameter of the common ground has the advantage of offering an intuitive grasp of the ACTION-GUIDANCE of evaluative uses of those sentences (§6.2.4). What we would be doing, essentially, is hardwiring the ability to update the practical commitments of interlocutors into the semantic value of those sentences.²⁴

²⁴ See the discussion on action-guidance on §2.3.1-2.3.2.

The obvious shortcoming of this view is that it does not account for descriptive (nor mixed) uses of evaluative sentences, since it predicts that evaluative sentences invariably update the hyperplan parameter. We have seen, however, that there are reasons to think that this is not always the case: in contexts like **(Experts)**, where a standard for good cars is assumed to be shared by speakers, evaluative sentences are used to communicate factual information. Consider again **(Experts)**, repeated here:

- (6.7c) **(Experts)**: Amir and Mora are car experts, they have a very similar standard for cars and they know that they do; Amir is telling Mora on the phone about the cars he is seeing at the auto show.
 Amir: The Mercedes is better than the Audi.

We argued in §6.2 that Amir’s utterance in **(Experts)** has the status of a judgment of fact, offering information about the properties of the two cars. But given that (6.7) contains an evaluative adjective, according to this initial proposal its utterance would effect an update on the hyperplan parameter of the common ground. Recall that this update consists on, roughly, a proposal to adopt a common plan of action. This type of update is hardly reconcilable with the utterance in **(Experts)**’s attested EPISTEMIC STATUS and DECEIT POTENTIAL—as we saw, **(Experts)** appears to provide information about the world that can be supported or falsified by evidence. For that same reason, such an update would also be at odds with the impression that a disagreement in the context of **(Experts)** would not be SUBJECTIVE. Finally, if the hyperplan parameter tracks shared practical commitments and, if what we said in §6.2.4 is right, it follows that the practical commitments of the interlocutors should be updated as a result of accepting this utterance. But that is not the case in **(Experts)**, as we saw; no new practical commitments are adopted there as a result of accepting this utterance. That is, in that context, (6.7) is not ACTION-GUIDING. We conclude that, as it stands, this simple story is insufficient to account for the contrast between *(Vis-à-vis)* and **(Experts)**.

To sum up: although we have yet to give the full picture, a straightforward dynamic version of our proposal for evaluative sentences in Chapter 4 correctly captures the discursive features that characterize evaluative sentences in their evaluative uses. The problem is that such sentences can also have descriptive and mixed uses, and these are unaccounted for. As far as we have said, the prediction of our view is that all evaluative sentences are used evaluatively (in accordance with the piece of received wisdom #1). What we need however, is a view that can accommodate the possibility of using evaluative sentences evaluatively, descriptively and in a mixture of both; moreover, we need a proposal according to which whether an evaluative sentence is used in any of these ways depends on the state of the conversation.

6.6 Evaluative & descriptive update (2nd attempt)

The key to representing the attested double duty of a sentence like (6.7) lies in what we said in Chapter 4 about the evaluation of individuals (§4.3.1). Note that, in our first attempt to assign a dynamic meaning that can capture the contrast between evaluative and descriptive uses of evaluative sentences, we opted for giving a dynamic interpretation to evaluative sentences along the model of thin evaluatives taking full propositions as arguments. But according to the semantics we gave for those type of sentences (§4.2.3), these only have an evaluative meaning. Therefore, this type of sentence is ill-suited to represent the descriptive use of a sentence like (6.7),

because when we interpret its semantic value dynamically, we obtain an update instruction on the hyperplan parameter of the common ground, and never on its context set.

When we turn to the evaluation of *individuals*, things are different. Sentences like *x is good* (where *x* is an individual), we said in §4.3.1, are best analyzed as having a double semantic value. On the one hand, that sentence predicates a descriptive property of *x*, and on the other hand, it expresses support for that property. That property, moreover, is massively underspecified by the lexical meaning of *good* (we represented such underspecification with [...]).

$$(4.42) \quad \text{That white Mercedes is good} = \begin{cases} \text{That white Mercedes is [...]} \\ \text{[...] is good} \end{cases}$$

Relative to the semantics given in §4.3.1, (4.42) has both a descriptive component whereby a non-evaluative property is ascribed to the white Mercedes, and it has an evaluative component whereby such property is supported. As we saw, the semantics for the evaluative component makes it equivalent to sentences of the form *F-ing is good*, which in turn are analyzed as *it is good that PRO is F*.

Letting *F* be a placeholder for the descriptive property, we have the following schema for (4.42):

$$(6.23) \quad \text{That white Mercedes is good} = \begin{cases} \text{That white Mercedes is } F \\ \text{it is good that } PRO \text{ is } F \end{cases}$$

This schema holds the key for capturing the contrast between evaluative and descriptive uses of evaluative sentences. Because according to our previous semantics, the first component in (6.23) is a descriptive proposition, that is, a set of world-hyperplan-alternative triplets which is hyperplan- and alternative-insensitive, but world-sensitive. This semantic object, understood in terms of its context change potential, would effect an update on the context set of the common ground—just as we saw for (6.18). And the second component in (6.23) is an evaluative proposition, that is, a world-insensitive but hyperplan- and alternative-sensitive semantic value. If interpreted dynamically, this semantic value would result in an update on the hyperplan parameter of the common ground—as shown in (6.20).

Before getting down to the details, however, let us note that things are slightly more complicated with (6.7), due to the comparative. The actual schema for a sentence like (6.7), then, would be something like this:

$$(6.24) \quad \text{The Mercedes is better than the Audi} = \begin{cases} \text{The Mercedes is } F\text{-er than the Audi} \\ \text{it is good that } PRO \text{ is } F \end{cases}$$

In the story about the evaluation of action-types/properties that we explored in §4.3.1, a sentence like the second component of (6.24) had the following truth conditions:

$$(4.43) \quad \begin{aligned} & \llbracket \text{it is good that } PRO \text{ is } F \rrbracket_{\langle w_i, h_i, a_i \rangle} = 1 \text{ iff} \\ & \{w' : \llbracket PRO \text{ is } F \rrbracket_{\langle w', h_i, a_i \rangle} = 1\} \in h_i(a_i[PRO \text{ is } F](w_i)) \end{aligned}$$

(4.43) is true just in case *F-ing* is among the preferred alternatives. Correspondingly, the evaluative component of (6.24) would have the following set of world-hyperplan-alternative triplet as its semantic value:

$$(6.25) \quad \llbracket \text{it is good that } PRO \text{ is } F \rrbracket = \{ \langle w, h, a \rangle : \{ w' : \llbracket PRO \text{ is } F \rrbracket_{\langle w', h, a \rangle} = 1 \} \in h(a[PRO \text{ is } F](w_i)) \}$$

Does this capture what we mean when we say that a car is better than another? I think it does. Suppose that we are very environmentally concerned, so that our evaluative standard for cars is tied exclusively to having low CO_2 emissions. The lower the better. Intuitively, when we say that the Mercedes is better than the Audi, we mean (i) that the Mercedes has lower emissions than the Audi, and (ii) that having low emissions is good. This seems correct, and it is predicted by our proposal for (6.7).

However, two difficulties should be pointed out. First, strictly speaking, the evaluative component of (6.24) captures that *having low-emissions is good*, but it does not represent that *the lower the better*. It is not difficult to fix this, though. The idea would be to interpret sentences like *F-ing is good* to mean something like *the F-er the better* in cases where F is a gradable property. Formally, this can be captured by letting the alternatives of the index be composed of propositions ascribing varying quantities of F (rather than different properties) to whoever PRO stands for.

$$\text{Alternatives: } \begin{cases} x \text{ is } n\text{-much } F \\ x \text{ is } m\text{-much } F \\ x \text{ is } o\text{-much } F \\ \dots \end{cases}$$

So the alternatives against which we compare the Mercedes and the Audi would be different degrees of F , rather than different properties altogether. In this way, we can correlate degrees of a non-evaluative property with degrees of support and rejection, which we can model with hyperplans.

Secondly, this view predicts that sentences like (6.7) do after all express outright support. But this goes against **POS-ATT** (see §2.5.1), because the comparative *better* does not invite an inference to the positive form *good*. However, what this proposal predicts is not that uttering (6.7) expresses support for either of the two cars; rather, the prediction is that it expresses support for the descriptive *property* in virtue of which the evaluation is made. So it is not clear that this really counts as a counterexample to **POS-ATT**. Moreover, it seems empirically correct. Suppose that Amir's standard for cars amounts to having low emissions. When he (in any context) utters (6.7), it seems harmless to suppose that his utterance expresses outright support for the property of having low emissions. Moreover, this is independent of how well each car fares relative to that standard. It could be that the Mercedes is a little better than the Audi although both are massive polluters.

The final piece of the puzzle concerns the underspecificity of *good/better*. In the meaning that we are assigning to (6.7), there is a blank. An interlocutor will not know what (6.7) means unless they can figure out what F stands for. And the place to look is the common ground of the conversation. Furthermore, depending on “where” in the common ground we find a value for F , we will determine the type of update that an utterance of (6.7) has: if the information about F is obtained from the context set, then the update will be an update on the hyperplan parameter; if that information is obtained from the hyperplan parameter, an utterance of (6.7) is going to update the context set. And if there is a stand-off, so that both context set and hyperplan parameter provide partial information (as occurs in mixed uses), an utterance of (6.7) will update both context set and hyperplan parameter.

That is what the contrast between (*Vis-à-vis*) and (**Experts**) essentially consists in. The idea is that if you can figure out what the blank represents in terms of the descriptive properties of those cars, then you can plug that information into the evaluative component of (6.24) and obtain an update instruction for the hyperplan parameter. That is what happens in (*Vis-à-vis*): since Mora knows studied the cars' specs, she knows how to fill in the factual blank, and she understands that Amir is trying to get her to adopt those factual properties as the good properties in cars. By contrast, in (**Experts**) Mora does not have the necessary information to fill in the factual blank (because she is not in front of the cars); but she can fill in the evaluative blank, in virtue of the fact that Amir and her share a standard for good cars. And she can use that information to learn about the properties of the Mercedes and the Audi. Finally, in (**Strangers**) Mora lacks both evaluative and factual information. But she can reduce her uncertainty by associating possible properties of the cars with possible standards for Amir. However, that will not reduce her ignorance *about* the cars nor the content of Amir's standard. Let us see how each mechanism works more carefully.

6.6.1 Basic dynamic entry

The picture that emerges is one where evaluative sentences are doing a sort of double duty: they can give factual and practical information (or a mixture of both). So evaluative sentences can perform two semantic operations: (i) an operation of attributing certain factual properties to some object(s) and (ii) an operation of avowing the adoption of a plan. The first operation can be modelled in terms of operations over possible worlds, and the second can be modelled in terms of operations over hyperplans. More precisely, in the dynamic framework that we advocate, the first operation is understood as an update function on the context set parameter of the common ground, while the second can be understood as an update function on the hyperplan parameter.

But each of those operations relies on certain information to be available in the common ground: the context-set update requires antecedently shared information about the shared car preferences of speakers, which is represented in the hyperplan parameter. And the hyperplan update requires factual information about the cars, which is represented in the context set. This describes the contrast between (*Vis-à-vis*) and (**Experts**): in the first situation, the fact that factual information about the cars is available causes the relevant update to be evaluative (an update of the hyperplan parameter); while in the second situation, having information about preferences (i.e. having a common standard) causes the relevant update to be descriptive (an update of the context set). If both context set and hyperplan parameter provide only partial information, the update will coordinate both parameters appropriately. This will be made clearer in the following subsections.

As we just saw, we start off from the assumption that a sentence such as (6.7) has a double semantic value:

$$(6.24) \quad \text{The Mercedes is better than the Audi} = \begin{cases} \text{The Mercedes is } F\text{-er than the Audi} \\ \text{it is good that } PRO \text{ is } F \end{cases}$$

On the one hand, (6.7) predicates that the Mercedes has more of a certain property F than the Audi. On the other hand, (6.7) expresses support for individuals that have property F .

How does (6.7) impact the common ground? It depends on the information that is previously available when (6.7) is uttered. When uttered, (6.7) will rule out certain possible worlds from the context set of the common ground if a certain condition is met by the hyperplan parameter; and it will rule out certain hyperplans from the hyperplan parameter if that same condition is met by the context set. And if the condition is met equally well by the context set and the hyperplan parameter, it will rule out certain combinations of possible worlds and hyperplans.

We can write this as a disjunctive definedness condition to be satisfied either by the hyperplan parameter, the context set or both. Relative to a common ground G with context set C and hyperplan parameter P , the update instruction for (6.7) looks as follows:

- (6.26) $G[(6.7)]$ is defined only if C or P provides the stronger F -value; if defined, then:
- a. if C provides a stronger F , then $G[(6.7)] = G[\text{it is good that } PRO \text{ is } F]$
 - b. if P provides a stronger F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi}]$
 - c. if C & P provide equally strong F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi and it is good that } PRO \text{ is } F]$ for any stronger value of F .

In words: in order to be defined, an utterance of (6.7) at a common ground G requires that either the context set or the hyperplan parameter of G provide the stronger value or specification for the property F that (6.7) makes reference to, or that both parameters provide equally strong values for F . *Stronger* here means logically stronger: for example, if the context set provides information to the effect that F stands for *sporty or reliable* and the hyperplan parameter is such that F stands for *sporty*, then the value of F in (6.7) will be *sporty*, because being sporty logically entails being sporty or reliable.

If the context set provides such specification, then we use that information, which is factual information, to update the common ground with the evaluative component of the meaning of (6.7), namely *that it is good that PRO is F* (where F is the property supplied by the context set). If the hyperplan parameter offers a stronger value for F , then we use that information to update the common ground with the descriptive component of (6.7), namely *that the Mercedes is F-er than the Audi* (where the value of F is provided by the hyperplan parameter).

And finally, if both the context set and the hyperplan parameter provide equally strong specifications for F , then the update consists in ruling out from the common ground the appropriate discrepancies between the context set and the hyperplan parameter *vis-à-vis* (6.7), namely situations where the Mercedes has property F but *being F* is not considered good (for whatever maximally specific value F could take). Suppose, again, that the common ground is such that F can only stand for *sporty or reliable*. This means two things: first, the context is such that the Mercedes is sportier or more reliable than the Audi. And second, the hyperplan parameter is such that sportiness or reliability are supported. In this common ground, if someone utters (6.7), the resulting update will be to rule out worlds w and hyperplans h (ignoring alternatives) such that *either* the Mercedes is sportier than the Audi at w and reliability is supported at h *or* the Mercedes is more reliable than the Audi at w and sportiness is supported at h . Note, however, that this update will not reduce our factual uncertainty about the cars nor our practical uncertainty about car standards: after the update, it's not known yet whether the Mercedes is sportier or more reliable, nor whether reliability or sportiness are supported. What are ruled out in this case are certain combinations of factual and practical states of affairs.

According to the semantics we proposed in §4.2-4.3, each of the semantic components of (6.24) will have a different kind of semantic value: the descriptive component has a descriptive proposition as its semantic value, while the evaluative component has an evaluative proposition. These are both sets of world-hyperplan-alternative triplets, but while a descriptive proposition is world-sensitive but hyperplan- and alternative-insensitive, an evaluative proposition is the other way around: it is hyperplan- and alternative-sensitive, but world-insensitive. This is the semantics that we assigned to each type of sentence in §4.2-4.3:

- (6.27) a. $\llbracket \text{The Mercedes is } F\text{-er than the Audi} \rrbracket =$
 $\{\langle w, h, a \rangle : \text{The Mercedes is } F\text{-er than the Audi at } (w)(h)(a) = 1\}$
 b. $\llbracket \text{it is good that } PRO \text{ is } F \rrbracket =$
 $\{\langle w, h, a \rangle : \{w' : \llbracket PRO \text{ is } F \rrbracket_{\langle w', h, a \rangle} = 1\} \in h(a[PRO \text{ is } F](w_i))\}$

(6.27a) is a set of world-hyperplan-alternative triplets containing all hyperplans and alternatives, but only worlds in which the Mercedes is F -er than the Audi. And (6.27b) is a set of world-hyperplan-alternative triplets containing all worlds, but only hyperplans and alternatives such that being F is supported.

As we saw in §6.5, if we interpret these propositions dynamically, we will obtain an update instruction that will either rule out some hyperplans from the hyperplan parameter while leaving the context set untouched, or will rule out worlds from the context set while leaving the hyperplan parameter untouched:

- (6.28) $G[\text{The Mercedes is } F\text{-er than the Audi}] =$

$$\begin{cases} C \cap \llbracket \text{The Mercedes is } F\text{-er than the Audi} \rrbracket \neq C \\ P \cap \llbracket \text{The Mercedes is } F\text{-er than the Audi} \rrbracket = P \end{cases}$$

 (6.29) $G[\text{it is good that } PRO \text{ is } F] =$

$$\begin{cases} C \cap \llbracket \text{it is good that } PRO \text{ is } F \rrbracket = C \\ P \cap \llbracket \text{it is good that } PRO \text{ is } F \rrbracket \neq P \end{cases}$$

The third type of update, in which neither the context set nor the hyperplan parameters provides a stronger value for property F , results in a partial update of both C and P . The resulting update instructions can be paraphrased like this:

- (6.30) a. Instruction on C : rule out all but worlds relative to which the Mercedes is F -er than the Audi.
 b. Instruction on P : rule out all but hyperplans relative to which being F is supported.
 c. Instruction on C, P : rule out all but worlds and hyperplans relative to which the Mercedes is F -er than the Audi and being F is supported (for all maximally stronger values of F).

The update on the hyperplan parameter is tantamount to adopting a plan to support cars with the features of which the Mercedes has more than the Audi, and what those features are is determined by the context set. In order to know what hyperplans to rule out, we need to know what the cars are like.

Things work exactly the other way around with the worldly update: it is an instruction to rule out worlds where the Mercedes does not have a higher degree than the Audi of the features that are supported. And what features of cars are supported is determined by the hyperplan parameter. In order to know what worlds to rule out, we need to know what the hyperplan parameter looks like.

And finally, in case that there is an “informational stand-off” between context set and hyperplan parameter, so that neither parameter supplies a stronger value to substitute for F , the update consists in coordinating both parameters: for any possible strengthening F' of the value of F , it will be the case that all worlds in the context set are worlds where the Mercedes is F' -er than the Audi and all hyperplans in the hyperplan parameter are hyperplans that support being F' . Let us expand on each of these options.

6.6.2 Evaluative use

Recall that, in (*Vis-à-vis*), we have interlocutors who do not share a standard but share the relevant information about the Mercedes and the Audi (because they are both in front of the cars and have studied their specs). In this context, if Mora accepts Amir’s claim then she does not learn any new information about the cars, but rather she adopts Amir’s standard for cars (or more precisely, she would be adopting a standard that is just like Amir’s at least with respect to cars like the Mercedes and the Audi). How does this come about?

Consider again the double meaning that we assigned to (6.7), which makes reference to a property, F , whose value is underspecified:

$$(6.24) \quad \text{The Mercedes is better than the Audi} = \begin{cases} \text{The Mercedes is } F\text{-er than the Audi} \\ \text{it is good that } PRO \text{ is } F \end{cases}$$

Consider next what the common ground looks like in (*Vis-à-vis*). Mora and Amir have seen and studied cars, but do not share a standard. This has two consequences: first, restricting our attention to information *about the cars*, the context-set of the common ground is rather reduced—in other words, it contains a lot of information. Secondly, given that Amir and Mora do not share a standard for cars, the hyperplan parameter is rather open.

If we consider the update function corresponding to (6.7) (repeated here), it seems that this common ground satisfies the first horn of its definedness condition: the context set provides a stronger F -value.

- (6.26) $G[(6.7)]$ is defined only if C or P provides the stronger F -value; if defined, then:
- a. if C provides stronger F , then $G[(6.7)] = G[\text{it is good that } PRO \text{ is } F]$
 - b. if P provides stronger F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi}]$
 - c. if C & P provide equally strong F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi and it is good that } PRO \text{ is } F]$ for any stronger value of F .

This is the case because Amir and Mora have sufficient information in the context set to define a logically stronger value for F than they could by looking at the hyperplan parameter; that is, they can let F stand for a property (or a combination thereof) such that the Mercedes and

the Audi stand in the appropriate comparative relation. And then, Amir and Mora can use that information to determine a plan, which will result in an update on the hyperplan parameter of the common ground.

To see this, suppose that the only thing that Amir and Mora knew about the cars was that the Mercedes is *sportier* than the Audi. If that was all the information in the context set about the Mercedes and the Audi, then the F in (6.24) could stand for *sporty*.

$$(6.31) \quad \text{The Mercedes is better than the Audi} = \begin{cases} \text{The Mercedes is } \underline{\text{sportier}} \text{ than the Audi} \\ \text{it is good that } PRO \text{ is } \underline{\text{sporty}} \end{cases}$$

In order to obtain an update instruction on the hyperplan parameter, all we need to do is retrieve our semantics for the sentence *it is good that PRO is sporty* and interpret it dynamically; that is, as an instruction to rule out hyperplans from the hyperplan parameter. As we showed in (6.27b), the semantic value of that sentence would be a set of world-hyperplan-alternative triplets. Substituting F for *sporty*, we obtain the following:

$$(6.32) \quad \llbracket \text{it is good that } PRO \text{ is } \underline{\text{sporty}} \rrbracket = \{ \langle w, h, a \rangle : \{ w' : \llbracket PRO \text{ is } \underline{\text{sporty}} \rrbracket_{\langle w', h, a \rangle} = 1 \} \in h(a[PRO \text{ is } \underline{\text{sporty}}](w_i)) \}$$

Since that sentence is evaluative, it contains all worlds but not all hyperplans/alternatives. The resulting instruction on the common ground would be an instruction to update the hyperplan parameter as follows:

$$(6.33) \quad G[\text{it is good that } PRO \text{ is sporty}] = \begin{cases} C \cap \llbracket \text{it is good that } PRO \text{ is sporty} \rrbracket = C \\ P \cap \llbracket \text{it is good that } PRO \text{ is sporty} \rrbracket \neq P \end{cases}$$

If this update is accepted, this results in Mora accepting Amir's standard. The hyperplan parameter of the common ground now supports sporty cars.

This type of update on the common ground, an update on the hyperplan parameter, satisfies the set of features that characterize evaluative uses of adjectives discussed in §6.2.1-6.2.4. Since it is a proposal to adopt a certain plan of action, it is to be expected that such utterance lacks EPISTEMIC STATUS—there is no *evidence* to offer in favor or against a proposal to adopt a certain course of action (§6.2.1). And given that this utterance does not offer information, it lacks the potential to become a lie in light of additional worldly information, that is, it lacks DECEIT POTENTIAL (§6.2.2). Furthermore, since this is a proposal to adopt a plan of action, and no plan seems *a priori* more justified or appropriate than any other, if a speaker rejects the update by signalling disagreement, there is no clear sense in which any of the interlocutors is *at fault*—they have a SUBJECTIVE DISAGREEMENT, in virtue of the fact that they are vouching for different plans (§6.2.3). And finally, a proposal to adopt a plan of action is of course a proposal to alter the practical commitments of interlocutors, so that the perceived ACTION-GUIDANCE of the utterance in (*Vis-à-vis*) is secured (§6.2.4). We conclude that this proposal is capable of accounting for evaluative uses of evaluative sentences.

6.6.3 Descriptive use

Descriptive uses of evaluative adjectives, by contrast, occur in contexts in which interlocutors share a standard but fail to share the relevant information about the objects under evaluation,

that is, in contexts like (**Experts**). In that type of context, the result of uttering an evaluative sentence has the effect of updating the context set, rather than the hyperplan parameter.

Consider again (**Experts**). In that context, Amir and Mora share a standard for good cars, but they do not share the relevant factual information about the Mercedes and the Audi: Amir has seen them, Mora has not. The update function for (6.7) is the same:

- (6.26) $G[(6.7)]$ is defined only if C or P provides the stronger F -value; if defined, then:
- a. if C provides stronger F , then $G[(6.7)] = G[\text{it is good that } PRO \text{ is } F]$
 - b. if P provides stronger F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi}]$
 - c. if C & P provide equally strong F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi and it is good that } PRO \text{ is } F]$ for any stronger value of F .

This time however, the second horn of the definedness condition for (6.7) is satisfied, in virtue of the fact that Amir and Mora share a standard. Their sharing a standard means that they share a plan to support or reject cars with different properties. If this is so, then when (6.7) is added to the common ground in (**Experts**), the update to the context set can succeed, and it results in an instruction to *rule out from the context set all but worlds such that the Mercedes is F -er than the Audi*, according to whatever value for F is supplied by the hyperplan parameter.

To see how this works, suppose that the only content of Mora's and Amir shared standard for cars is *that sporty cars are good*. If that is the case, then we have an F -value that we can plug in the descriptive component of (6.7):

$$(6.31) \quad \text{The Mercedes is better than the Audi} = \begin{cases} \text{The Mercedes is } \underline{\text{sportier}} \text{ than the Audi} \\ \text{it is good that } PRO \text{ is } \underline{\text{sporty}} \end{cases}$$

This time, however, the value of F comes from the hyperplan parameter, not from the context set. And we can use that practical information to obtain a descriptive sentence, *The Mercedes is sportier than the Audi*, that we can update the common ground with. The semantics of such a sentence, as we saw in (6.27a), is a set of world-hyperplan-alternative triplets. Substituting F for *fast*:

$$(6.34) \quad \llbracket \text{The Mercedes is } F\text{-er than the Audi} \rrbracket = \{ \langle w, h, a \rangle : \text{The Mercedes is } F\text{-er than the Audi at } (w)(h)(a) = 1 \}$$

This is a set of world-hyperplan-alternative triplets containing all hyperplans and alternatives, but only worlds in which the Mercedes is F -er than the Audi. Therefore, if we update the common ground with it, we will obtain an update on the context set but not on the hyperplan parameter of the common ground:

$$(6.35) \quad G[\text{The Mercedes is sportier than the Audi}] = \begin{cases} C \cap \llbracket \text{The Mercedes is sportier than the Audi} \rrbracket \neq C \\ P \cap \llbracket \text{The Mercedes is sportier than the Audi} \rrbracket = P \end{cases}$$

If Mora accepts Amir's utterance in (**Experts**), their common ground is now such that the Mercedes is sportier than the Audi.

Treating Amir’s utterance in (**Experts**) as an update instruction on the context set of the common ground means treating it as any factual, descriptive assertion. And this is why that utterance shares with factual, descriptive assertions what we called its conversational “profile”. Briefly: this is a proposal to update the epistemic parameter of the common ground, that is the context set, so it is no wonder that interlocutors can ask for more evidence in support of that type of update; that is, it is not surprising that the utterance in (**Experts**) has EPISTEMIC STATUS (§6.2.1). Secondly, since it is an update that offers certain information about how the world is, it can be falsified by worldly information, and so it has DECEIT POTENTIAL (§6.2.2). Thirdly, the fact that it offers information about the world can also account for the apparent lack of SUBJECTIVENESS if a disagreement arises (§6.2.3). And finally, since this is not an update to the hyperplan parameter of the common ground, it makes sense for the utterance in (**Experts**) to lack ACTION-GUIDANCE (§6.2.4). We conclude that this proposal can also account for descriptive uses of evaluative sentences.

6.6.4 Mixed use

Let us turn now to our last possibility, namely a mixed use of *better*. Recall, once again, the update instruction for (6.7):

- (6.26) $G[(6.7)]$ is defined only if C or P provides the stronger F -value; if defined, then:
- a. if C provides stronger F , then $G[(6.7)] = G[\text{it is good that } PRO \text{ is } F]$
 - b. if P provides stronger F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi}]$
 - c. if C & P provide equally strong F , then $G[(6.7)] = G[\text{The Mercedes is } F\text{-er than the Audi and it is good that } PRO \text{ is } F]$ for any stronger value of F .

For a mixed use of *better* to occur, the common ground has to be such that neither context set nor hyperplan parameter provide stronger information than the other parameter. For if that was the case, there would either be a descriptive or an evaluative update (as the case might be). Suppose that the common ground was such that the Mercedes is either sportier or more reliable than the Audi (but not both) and that reliability (but not sportiness) is preferred. Relative to this common ground, an utterance of (6.7) will have the effect of reducing our factual uncertainty by ruling out worlds where the Mercedes is sportier. Conversely, if the common ground were such that the Mercedes is sportier and either sportiness or reliability is preferred, an utterance of (6.7) will have the effect of reducing our practical uncertainty by ruling out hyperplans where reliability is preferred.

For a mixed use to occur, the context set and the hyperplan parameter have to stand in a sort of informational stand-off, as we said before. Furthermore, there has to remain some uncertainty in both parameters of the common ground—if it’s already known what the cars are like as well as what properties of cars are supported, no further update would be possible (the situation would be similar to (**Sharing-is-caring**)). Now, if the common ground meets the conditions for a mixed update, then such an update will work as follows: for every possible strengthening F' of F , we will rule out combinations of world w and hyperplan h such that the Mercedes is F' -er than the Audi at w and being F' is supported relative to h .

This matches the situation described in (**Strangers**): Amir and Mora do not know each other and do not share information about the cars. But we can assume that background or world

knowledge provides some information both about what type of cars are shown in auto-shows as well as what type of cars people prefer. Moreover, it is natural to think that there is some overlap between those properties, so that we can arrive at a set of properties that cars in auto-shows tend to have *and* that people tend to prefer. Suppose—once again!—that those properties are just *being sporty or reliable*, and assume, further, that the state of the world is such that either the Mercedes is sportier or more reliable than the Audi (but not both); and that the shared preferences are such that either sportiness or reliability are preferred.

Now, when Amir utters (6.7) in such a context, Mora learns that the Mercedes has more of either of those two properties, reliability or sportiness, but she doesn't know which one. And she also understands Amir's utterance as a proposal to update their preferences to either sportiness or reliability, but she doesn't know which one either. Nonetheless, she can rule out combinations of worlds and hyperplans such that the Mercedes is sportier but reliability is preferred, as well as combinations such that the Mercedes is more reliable but sportiness is preferred.

What is interesting (and puzzling) about mixed uses is that their outcome is not, in a sense, more knowledge about the cars nor about preferences: after updating the common ground in the way just described, there are still possible worlds in the context set such that the Mercedes is sportier and worlds in which the Mercedes is more reliable. Similarly, there remain hyperplans in the hyperplan parameter relative to which reliability is preferred and hyperplans relative to which sportiness is preferred. It looks as though mixed uses do not further our knowledge about the world nor about accepted preferences. But what do they do, then? What is their conversational contribution? It is hard to say exactly, but one possible way to look at it is the following:²⁵ when Mora accepts Amir's utterance in (**Strangers**), she does not learn anything about the cars and she does not update her preferences, but she learns something *about Amir*: she learns that Amir values the Mercedes over the Audi. She does not know in virtue of which properties (and therefore she is not in a position to agree or disagree with him) but this is knowledge that she did not have before Amir's utterance. It seems to be factual information about Amir's preferences. As such, it is information that can be used by Mora to predict Amir's behavior.

However, the type of information that Mora can obtain from Amir's utterance in (**Strangers**) should be carefully distinguished from the type of knowledge that results from an evaluative update. What Mora learns in (**Strangers**) is information about Amir: she can rule out (or consider less likely) possible worlds where Amir goes on to choose the Audi over the Mercedes. But that is not information that guides her own behavior. For that to be the case, as we have hypothesized, Amir's utterance would have to have an impact on the hyperplan parameter of the common ground, that is, on Mora's own preferences. But without information about the properties in virtue of which Amir makes his evaluation, that cannot happen.

Finally, as we said at the end of §6.2, it is hard to say what the situation is in (**Strangers**) *vis-à-vis* the criteria that we used to distinguish evaluative and descriptive uses of evaluative adjectives. The reason for this is that all of these criteria require that Mora possess more information about the situation than she actually possesses in (**Strangers**). In the model of conversation that we are proposing, these observations are captured by the fact that Amir's contribution in (**Strangers**) neither augments shared factual knowledge nor restricts the shared standards. Given that, it would be odd for Mora to challenge Amir's evidence in absence of further information. For the same reason, Mora could not accuse Amir of lying. And similarly, she could neither agree nor disagree with him. Finally, with respect to whether Amir's utterance in (**Strangers**) can guide action, the intuitive answer is negative as well. Since Amir's utterance

²⁵This was suggested by Paul Égré (p.c.).

cannot update Mora's preferences in absence of further information about the cars, shared practical commitments of speakers remain as they were before the utterance. Nonetheless, note that this is compatible with the idea that Amir's utterance in (**Strangers**) can be used by Mora to predict *his* behavior. In this sense, Amir's utterance in (**Strangers**) is no more action-guiding than a third-personal report would be (*Amir thinks that the Mercedes is better than the Audi*).

6.7 Conclusion

In this final chapter, we have presented a set of cases in order to show that evaluative sentences can be used to either (i) make judgments of value, (ii) make judgments of fact or (iii) a mixture of both. Crucially, this depends on the state of the conversation in which an evaluative sentence is uttered—following an old suggestion of Hare's against certain received pieces of common wisdom, and prompted by more recent observations stemming from C. Barker 2002 and others about descriptive and metalinguistic uses of dimensional adjectives. However, against the possibility of subsuming what we have described as evaluative uses under metalinguistic uses of language, we have argued that evaluative uses have the crucial effect of updating our practical commitments over and above commitments about language use (which is what metalinguistic moves characteristically do).

We have presented and defended a simple dynamic model of communication in which evaluative sentences, in their evaluative use, can alter the practical commitments of speakers, which are represented *via* a specialized parameter in the common ground of a conversation that we called the *hyperplan parameter*. We have modelled the content of that parameter using Gibbardian hyperplans, following the semantics laid out in Chapter 3 and Chapter 4. But evaluative sentences have descriptive and mixed uses as well, which means that those sentences can also update the factual information of interlocutors, just like any other descriptive claim. We submit that this dynamic model captures satisfactorily the contrast between evaluative and descriptive uses of evaluative sentences, that is, between contexts like (*Vis-à-vis*) and (**Experts**). In the former case, descriptive information is antecedently shared and uttering an evaluative sentence has the conversational effect of updating the practical commitments of speakers; it is therefore an evaluative use. In the latter case, evaluative information is antecedently shared, and uttering that very same sentence has the effect of updating the factual information shared by speakers. It is thereby a descriptive use of an evaluative sentence. Finally, in a context like (**Strangers**), where not enough factual information is shared to perform an evaluative update, and similarly, not enough evaluative information is shared to perform a descriptive update, uttering an evaluative sentence has the effect of informing about speakers' preferences without these preferences being adopted by the hearers. A further assessment of this view, as well as an exhaustive comparison to possible alternative accounts of the same phenomenon, remains the subject of future work.

Chapter 7

Concluding remarks

In this dissertation, I have presented a proposal about the meaning of evaluative adjectives like *good*, *bad*, *beautiful*, *loyal* or *tasty*. The view defended is that the contribution of these words to the unembedded sentences in which they appear—*evaluative sentences*—makes those sentences semantically sensitive to the practical commitments of participants in a conversation. In a static semantics such as the one proposed in Chapter 3-4, such sensitivity is cashed out by proposing that the truth of evaluative sentences depends on hyperplans, which form a distinguished parameter in a semantic model of interpretation. In the dynamic semantics proposed in Chapter 6, such sensitivity is accounted for by assigning to evaluative sentences the dynamic effect of updating the practical commitments of speakers, which are also represented as a specialized parameter of the common ground of a conversation. In this concluding chapter, we review the contents of the dissertation with an eye on possible objections and avenues for future research.

This view is a form of NON-FACTUALISM, because it defends that evaluative sentences do not serve the standard function assigned to declarative sentences, which is to describe the world, or put differently, to offer factual information. To the contrary, evaluative sentences by and large do not offer factual information that is not already available to speakers. Rather, they have the communicative effect of changing our commitments about how to act. In this respect, even though evaluative sentences are declarative sentences, they semantically resemble imperatives, exhortatives and other types of action-oriented language. Contemporary expressivists about evaluative language share the negative thesis about what evaluative sentences do not do, but their positive proposals are sometimes perceived to be problematic. In particular, expressivists associate evaluative sentences with characteristic mental states. They will say, for instance, that a sentence like *stealing is wrong* expresses a certain non-cognitive or non-representational mental state (perhaps a desire to not steal, or a state of moral blame towards those who steal, etc). This way of positively characterizing what evaluative sentences do, insofar as it relies on a strong theoretical link between the meaning of a sentence and the mental state it expresses, gets expressivists into all sorts of troubles. I have talked about evaluatives expressing *practical attitudes*, but I explicitly opted to remain neutral about whether such attitudes are best thought of as types of mental states or as certain types of public commitments (which is, actually, the type of view I favor). It is for this reason that the view defended here is not called ‘expressivist’. Rather, the more neutral label ‘non-factualist’ is preferred and used throughout. This “anti-mentalistic” stance is by no means original, but it is relatively recent (other theorists adopting this type of approach include Charlow 2015; Yalcin 2011, 2012 or Willer 2017).

Moreover, it is assumed in this work that the characteristic communicative function of evaluative sentences should be represented semantically. In other words, that evaluative sentences operate on our practical commitments (cf. that they guide action) is an empirical observation, but the choice to represent that function semantically, that is, as part of the lexical meaning of these words, is a theoretical choice among others. In particular, the possibility that the action-guiding inferences associated with evaluatives arises as a result of some type of pragmatic mechanism was left unexplored. This view was defended by authors like Finlay (2005) or Copp (2001), but recent developments in formal pragmatics suggest that those proposals could be updated and explored more systematically with contemporary linguistic tools. This remains the subject of future work.

The second chapter of this dissertation offers an overview of the linguistic criteria that distinguish evaluative expressions, and in particular evaluative *adjectives* (evaluatives for short), from other expressions in natural language. The main line of argument there is that the linguistics literature only reaches a necessary property of evaluatives (their ‘lexical’ or ‘ordering subjectivity’); while the philosophical literature offers a sufficient property (the capacity to express outright positive or negative practical attitudes), but neither reaches a stable criterion on its own.

From a linguistic point of view, the features that most clearly distinguish evaluatives have to do with their gradability and their subjectivity: as we saw, evaluatives are a type of gradable predicates that are liable to appear in so-called faultless or subjective disagreements both in positive and in comparative form (as well as to embed under subjective attitude verbs such as *find* in the same forms). Consider the following set of things we could say to describe or characterize a certain chord progression.¹

- (7.1) This is an arpeggiated chord progression.
- (7.2) This is a long chord progression.
- (7.3) This is a surprising chord progression.
- (7.4) This is a good chord progression.

This list progresses from objective (*arpeggiated*) to subjective (*surprising*, *good*) predicates. Our question in Chapter 2 was how to distinguish—linguistically—the last of these adjectives, *good*, from the rest.

A first, obvious criterion is that *good* is gradable, which is attested by the admissibility of adjectival modifiers. That serves to rule out *arpeggiated*:

- (7.1) This is a {# very} arpeggiated chord progression.
- (7.2) This is an extremely long chord progression.
- (7.3) This is a less surprising chord progression.
- (7.4) This is a really good chord progression.

This criterion, however, does not yet single out *good*. A further criterion is whether these expressions can figure in so-called faultless or subjective disagreements. There are two crucial ways of applying this criterion to gradable adjectives: one is to ask whether adjectives in the positive form can give rise to a subjective disagreement. Many gradable predicates do

¹I am thinking of the beginning of the [second movement](#) of Falla’s Harpsichord concerto.

(i.e., all of the above), but that has to do with the particular semantics of the positive form of these adjectives, and not with their lexical meaning. The way to test whether adjectives are *lexically* subjective, then, is to ask whether they can figure in subjective disagreements in the comparative form as well. And then, the list reduces. Because many adjectives that give rise to subjective disagreement in the positive form do not do so in the comparative. This is the case of so-called ‘dimensional’ adjectives like *long*:

- (7.5) This is a long-er chord progression than that one. ≈ not subjective
 (7.6) This is a less surprising chord progression than that one. ≈ subjective
 (7.7) This is a better chord progression than that one. ≈ subjective

Now, as we argued in Chapter 2, even though there is an array of proposals that aim at capturing the contrast between adjectives that are subjective in the comparative from adjectives that are not, the linguistics literature on adjectives does not, for the most part, offer any further way of distinguishing *surprising* from *good*. In this sense, linguistics arrives at a necessary property of evaluatives, but not yet a sufficient one.

To reach a more precise criterion to distinguish evaluatives from other subjective adjectives, we turned to philosophy, in particular to metanormative theory. We observed that the feature that has been pointed out as most characteristic of evaluatives is their ACTION-GUIDANCE. This is the property in virtue of which certain expressions can guide action, decision and behavior.

In language, this feature can be observed by considering the fact that some adjectives very clearly invite an inference that the speaker either supports or rejects whatever falls under the predicate that she is using, and some others do not. This sets apart adjectives like *good* from adjectives like *surprising*:

- (7.3) This is a surprising chord progression. ↷ support / rejection ?
 (7.4) This is a good chord progression. ↷ support

However, we noted that this way of conceiving action-guidance is *too precise*: evaluatives in comparative form, for instance, do not express outright support in this way:

- (7.7) This is a better chord progression than that one ↷ support / rejection ?

This suggests that the proposed criterion for action-guidance as the defining property of evaluatives is at best a property of *some forms* of evaluatives.

Summing up, it seems that linguists don’t get close enough to evaluatives—as being lexically subjective is a *necessary* feature of evaluatives; while philosophers get too close—as the expression of outright support and rejection is only a *sufficient* property of evaluatives.

The observation that the expression of outright support and rejection is a property of the positive form of evaluatives has further empirical and theoretical consequences. Empirically, a generalization can be made correlating this property of evaluatives with the presence of an inference to the positive form (we called this generalization **POS-ATT**). Theoretically, insofar as action-guidance informs traditional formulations of non-factualism or expressivism about evaluatives, such theories turn out to be best suited to capture the meaning of evaluatives in their positive form, but they are not straightforwardly applicable to other adjectival forms, cf.

comparatives. To wit: some expressivists say that evaluative sentences express non-cognitive attitudes of (dis)approval. This clearly applies to (7.4), but not to (7.7).

More importantly, given that comparatives and the positive form are compositionally related, if non-factualists are right that practical attitudes of support/rejection figure in the meaning of evaluatives, they owe an account of the way in which those attitudes contribute to, or are the product of, a process of semantic composition. This is what we dubbed the ‘sub-sentential’ Frege-Geach problem. It is the problem of reaching a formulation of non-factualism for evaluatives that can (i) account for the range of possible adjectival form and their attested action-guiding features, and (ii) predict those features compositionally.

A possible criticism of our approach is that perhaps some expressivists or non-factualists might have settled on a vaguer or weaker characterization of what evaluatives do, so as not to incur in the problems just described. Saying that evaluatives express practical attitudes, for instance, would be sufficiently neutral. But it is not enough, as accounting for the compositional semantics of certain basic adjectival forms is an inevitable chapter in building a full theory of the meaning of these expressions. Arguably, the sub-sentential Frege Geach problem is as important as the traditional “supra-sentential” version. Hence, in Chapter 2-4, I lay out a version of non-factualism designed to capture the difference between the action-guiding properties of *good* and *better* and derive that difference compositionally *via* a DELINEATION approach.

The choice of a delineation semantics (in the vein of C. Barker 2002; Benthem 1982; Klein 1980 or Burnett 2017 *a.m.o.*), was justified by arguing that degree semantics (e.g., Cresswell 1976; Kennedy 2007; Kennedy and McNally 2005, *a.m.o.*) faces a problem when accounting for the presence of such outright practical attitudes. Degree semantics, on its most usual construal, holds that the meaning of the positive form of any gradable adjective results from two semantic components, its lexical meaning—i.e., a measure function—and a threshold-contributing silent morpheme *POS*. The problem was that neither of those components seems to be the right source of the outright positive/negative attitudes. If we assign practical support/rejection as the meaning, or part of the meaning, of *POS*, then one is left wondering why are other gradable adjectives not equally action-guiding in their positive form. And if we associate practical attitudes with the lexical meaning of evaluatives, then the correlation **POS-ATT** becomes mysterious (cf. §2.5.2). This argument might be challenged, and it bears mentioning that an alternative formulation of non-factualism using the more standard degree semantics is perhaps an alternative option to the one pursued here. At the very least however, the idea of using simple attitudes of practical support and rejection to construct scales of value could have a metasemantic, if not a semantic, implementation.

In delineation semantics, the positive form of a gradable adjective is treated as a simple predicate (*albeit* with a partial denotation), and nothing prevents assigning outright attitudes of support and rejection as the meaning of these predicates. The basic idea is that a positive evaluative such as *good* expresses, in its positive form, outright support; and a negative adjective such as *bad* expresses, in its positive form, outright rejection. The challenge, then, is to derive all the range of possible adjectival forms starting (compositionally) from the simplest of them (the positive form). This is what Chapter 3-4 are devoted to. Note that, even though a sentence like (7.7) does not express outright support, it does express a sort of conditional support: it expresses support for the first chord progression *over the second*, or *given a choice between the first and the second*.

(7.4) This is a good chord progression.

↷ support

(7.7) This is a better chord progression than that one

↷ support for the first given a choice between the two

So (7.7) does, after all, express an action-guiding attitude. The difference with (7.4) is that the attitude expressed by (7.7) is more sophisticated, as it subsumes support under a sort of proviso. Delineation semantics offers an intuitive way of cashing out this contrast. To see this, take a dimensional adjective like *tall*: in a delineation framework, the positive form *tall* has a denotation relative to a contextually determined comparison class. The sentence *Feynman is tall* is true at a context c , in this view, just in case Feynman falls in the positive extension of the predicate *tall* relative to the comparison class determined at c . At one context, the comparison class might be American physicists from the XXth century and that sentence is true; at a different context the comparison class can be people born in Queens, and the sentence can be false.

The comparative *taller*, by contrast, is insensitive to comparison classes. A sentence like *Skłodowska is taller than Feynman* has the same truth value at any context, and thus relative to any comparison class. In delineation semantics, this is guaranteed by introducing quantification over the comparison class parameter. Whereas an individual x falls under *tall* just in case x is in the positive extension of *tall* relative to a contextually determined comparison class, a pair of individuals $\langle x, y \rangle$ falls under *taller* just in case there exists some comparison class relative to which x is tall but y isn't.

Our proposal for the relation between *good* and *better* is structurally identical. The basic thought, roughly, is that an object of evaluation φ is *good* just in case it is chosen relative to a set of alternatives determined by the context; and a pair of objects $\langle \varphi, \psi \rangle$ stands in the *better* relation just in case there exists some set of alternatives relative to which φ is good but ψ isn't. This proposal is carefully spelled out in Chapters 3-4.

Chapter 3 begins with a proposal to formalise practical commitments as plans of action, which can be accepted and modified by agents through their use of evaluative language. To this effect, in Chapter 3 we recruited Gibbard's notion of a hyperplan, that is, a maximally specific plan of action. Just like possible worlds are maximally specific states of affairs, hyperplans are maximally specific plans of action. For any situation that an agent could find herself in, a hyperplan tells the agent what to do. In the semantic structure that we proposed, declarative sentences are evaluated for truth relative to a semantic structure that introduces two new semantic parameters, in addition to the usual set of possible worlds: a set of alternative-generating functions and a set of hyperplans. A set of alternatives offers something like an occasion for choice, and hyperplans are functions from such sets of alternatives to subsets thereof, where only the preferable options remain. Declarative sentences are evaluated for truth against indices of evaluation, which are triplets of distinguished points of those three parameters.

Theoretically, to adopt this type of view is to adopt a kind of *semantic relativism*. Is this justified? We noted that alternative contextualist proposals are prey to a host of traditional objections. However, many contemporary proposals have tried rescuing contextualism, hence some of those traditional objections might have lost force. And closer to our interests, even the traditional expressivist move of pressing a contrast between expressing a mental state vs. saying that one is in a mental state seems to be accountable within a traditional, descriptivist semantics. But treating hyperplans as an independent semantic parameter of interpretation has other theoretical advantages. Importantly, we can treat the evaluative realm as a quasi-objective

parameter in the semantics and at the same time allow the evaluative and the factual realms to vary independently from each other.

Truth conditions work, informally, as follows: (7.4) is true relative to an index of evaluation formed by a world, hyperplan and alternative-set triplet $\langle w, h, a \rangle$ just in case this chord progression is among the alternatives that h selects relative to a ; that is, the first chord progression is chosen, or preferred, among a set of alternatives determined by the index of evaluation. And (7.7) is true relative to an index $\langle w, h, a \rangle$ just in case there is a subset X of the set of alternatives a such that (i) the first chord progression is among the alternatives that h selects relative to X , but (ii) the second chord progression is not among the alternatives that h selects relative to X .

The reason why only (7.4) expresses outright support is because it is made true by hyperplans that actually prefer this chord progression over the available alternatives at the index. By contrast, (7.7) does not express outright support because the preferences of the hyperplan relative to the available alternatives remain open. All that the truth of (7.7) requires is that we can find a subset of the alternatives determined by the index relative to which the first chord progression would be preferred to the second. Hence, the speaker of (7.7) does not express a preference for either chord progression, but would prefer the first over the second if she had to choose between them.

This proposal leaves a number of open questions. Most importantly, we hypothesized that the sets of alternatives at an index of evaluation coincides with the set of alternatives relevant for computing a focus alternative set, and can be affected in a similar way. This might come into question, and certainly requires more investigation than was offered here. Furthermore, our proposal was to offer an lexical entry for the positive form and then derive a few relational predicates, such as equatives and comparatives. But many more adjectival forms remain to be accounted for. What is the meaning of, e.g., *very*, *extremely*, *a little... good*? More importantly, does our plan semantics in any way prevent or hinder standard accounts of such modifiers available in the formal semantics literature? We have not attempted to answer the former question, but we remain hopeful that the answer to the latter is negative.

In Chapter 4 we explored how our basic proposal could be extended in two interrelated directions: first (*pace* the current running example) we began by giving a semantics for evaluative adjectives in their use as sentential operators (*it is good/bad that*), and then we considered other objects of evaluation—namely, action-type or properties (*stealing is bad*) and individuals ((7.4)/(7.7)). The main complication introduced by these extensions is due to the evaluation of individuals. As argued extensively in Chapter 6, when an evaluative adjective is applied to an individual, it both attributes a set of under-specified descriptive properties and contributes an evaluation of those properties. Thus, strictly speaking, evaluative adjectives do not take individuals as arguments; at most they take properties, which in turn are analyzed as propositional functions, and therefore follow the model of evaluative sentential operators.

Secondly, we considered evaluative adjectives of increasing degrees of descriptive *thickness*. The distinction between different degrees of thick adjectives is best seen by considering evaluatives taking individuals as their arguments. Once established that evaluatives predicate certain descriptive properties of the individuals that they apply to, increasing degrees of thickness correlate with increasing specificity in the descriptive component of these adjectives. The thinnest evaluatives, namely all-purpose evaluatives such as *good* and *bad*, attribute a completely under-specified set of descriptive properties. “Flavor”-specific evaluatives, such as moral and aesthetic adjectives (*ethical*, *beautiful*, *tasty*), introduce restrictions on the underspecification of such de-

scriptive properties. And finally, thick adjectives (*cruel, elegant, loyal*) predicate very specific descriptive properties. In our view (and restricting our attention to individuals), the difference between thin and thick adjectives is not that the latter carry an additional semantic component that the former lack; rather, thick adjectives simply introduce stronger restrictions on the descriptive component that all evaluatives have (when they take individuals as arguments, that is).

Future work should come to terms with the ramifications of this generalization about evaluatives. In general, as was mentioned in §4.3, there are reasons to think that a proposal that relies so heavily on a clear distinction of evaluative and descriptive components will have trouble dealing with certain thick adjectives like *dumpy* or *delicate*, where it just feels really difficult to separate a descriptive and an evaluative component. More specifically, a prediction that falls out from my proposal is that thick adjectives cannot enter into the type of context-dependent evaluative/descriptive usage described in Chapter 6, because they always carry both a descriptive and an evaluative component. That seems to be *prima facie* correct, but a full account of the whole range of evaluatives along the thickness spectrum should explain more data. For instance, although there has been quite some debate (Cepollaro and Stojanovic 2016; Väyrynen 2013, 2017), it seems that the evaluative component of thick adjectives projects much in the way that some presuppositions do. Is this compatible with our view about thickness being a *continuum*, and more importantly, with the way I claim that thin adjectives—which are on the other side of that *continuum*—behave?

In the second part of Chapter 4 we looked at further properties of evaluative adjectives, especially in relation to their scale structure. We considered previous attempts in the literature on subjective predicates to argue that evaluative and akin adjectives challenge the traditional distinction between relative and absolute adjectives (Kennedy and McNally 2005), and we saw those arguments as wanting. For the most part, evaluatives appear to be relative adjectives with some special features, such as the fact that they are not especially sensitive to comparison classes.

In relation to their scale structure, we considered Lassiter’s (2017) discussion about the scale structure of sentential *good* and we concluded that there are definitely reasons to rule out an ordinal scale, yet a decision between an interval and a ratio scale could not be reached. Data about the inferential behavior of sentential *good* points to an interval scale, but the admissibility of loose ratio modifiers suggests that perhaps a case can be made against interval scales, and possibly in favor of a very simple ratio scale. The argument in that section was involved and I will not attempt to summarize it here, but since we could not reach a final decision, that discussion awaits future research.

Another important thing to note, however, is that considerations about scale structure no longer apply only to evaluatives, but to other subjective predicates as well. An exciting project beyond this dissertation is to explore further scalar properties of subjective predicates in general, that is, including experiential and emotional predicates, and possibly others too. An attractive hypothesis in this respect is something that was hinted at in Chapter 5, namely that all predicates whose semantics involves an ‘inherent human element’ (McNally and Stojanovic 2017; Solt 2018) might share certain scalar properties.

Finally, Chapter 6 explored a systematic contrast that arises when speakers apply thin evaluatives to individuals: most often, speakers will use these adjectives assuming that their audience has access to the descriptive properties in virtue of which they make their evaluations (this is

the case, for example, of the usual disagreement dialogues between people discussing whether their food is tasty). And then, an evaluative expresses a positive attitude towards those properties. But sometimes, a speaker will use those same adjectives in a context where their audience lacks access to the relevant descriptive properties but knows and shares a certain evaluative standard. Following Hare (1952), I have argued that in those cases people make a completely descriptive use of those adjectives. And sometimes too, evaluatives appear to be used both descriptively and evaluatively at once.

Coming back to our earlier set of examples, the descriptive contexts are speech situations in which *good* turns out to behave conversationally just like *arpeggiated*.

(7.1) This is an arpeggiated chord progression.

(7.4) This is a good chord progression.

Suppose that you and I agree that arpeggiated chord progressions are the pinnacle of musical quality. And suppose that I have heard the beginning of Falla's second movement of his Harpsichord concerto, but you haven't. For me to tell you that it starts with a really good chord progression is a way for you to learn that it begins with an arpeggiated chord progression. This is, I argued, a completely descriptive use of *good*.

Chapter 6 develops a set of criteria to set apart evaluative and descriptive uses of adjectives like *good*. These criteria revolved around the idea that, when a speaker makes a descriptive use of *good*, she is conveying factual information. Therefore, her audience can request evidence in favor of her statement and accuse her of lying if she thinks that her statement is untrue. Moreover, the factual character of these uses makes them unsuitable to give rise to subjective or faultless disagreement, at least in the way that these are usually characterized. Finally, descriptive uses seem not to guide action in the way that evaluative adjectives normally do. In light of this diagnostic, we concluded that a semantics for evaluatives is needed where these adjectives can perform a double duty, namely informing the audience of certain worldly facts and changing the practical commitments of speakers (or both).

This contrast, however, is largely compatible with an alternative hypothesis: perhaps what we have been calling evaluative uses are just *metalinguistic* uses of those adjectives, that is, uses where speakers negotiate among a set of possible interpretations of the relevant adjective. So, in the situation where we are both listening to Falla's concerto and you tell me that that's a good chord progression, you'd be proposing to use *good* with a denotation such that this chord progression falls under it; and if I resist that, I would be rejecting your metalinguistic proposal about the denotation of *good*. This is a suggestive view, as it appears to capture most features of the contrast between (what we called) evaluative and descriptive uses of evaluative adjectives. The main issue with this proposal however, is that evaluative uses have an additional property, which is their capacity to manipulate practical commitments (over and above commitments about discourse or language use). In this respect, evaluatives appear to do something that neither metalinguistic nor descriptive uses of language can do. Nonetheless, let us acknowledge that distinguishing metalinguistic and evaluative uses of language is notably tricky, and requires more attention than we have accorded it here.

The last sections of Chapter 6 were devoted to presenting a dynamic model of conversation, built with the tools that were developed in Chapters 3-4. We mapped our previous semantic structure onto the common ground of a conversation. We distinguished two different parameters in the common ground: a context set, containing factual information, and a hyperplan

parameter, representing shared practical commitments of speakers.

In their evaluative uses, evaluative adjectives update the content of the hyperplan parameter. If we're both listening to Falla's concerto but we know that we have different standards and I utter (7.4), then my utterance should be interpreted as a proposal to update our shared practical commitments in a way that we both accept being arpeggiated as a supported, or preferable, property of chord progressions. On the other hand, descriptive uses update the context set: if only I have listened to the concerto but we both share a love for arpeggiated chord progressions, then my utterance of (7.4) should be interpreted as a proposal to update the context set, that is our shared factual knowledge, with the information that the chord progression is arpeggiated.

And finally, if both the context set and the hyperplan parameter contain only partial information both about the chord progression (maybe that the progression is either arpeggiated or in a major key) as well as about our shared preferences (similarly, that we support either arpeggiated or major chord progressions), then an utterance of (7.4) has the effect of ruling out possible discrepancies between facts and preferences. For example, we can now rule out from consideration a common ground relative to which the chord progression is arpeggiated but we prefer major chord progressions, and *vice versa*.

Surprisingly, this type of update has the effect that all live options about our preferences as well as about the properties of the chord progression remain in the common ground. What we rule out, in this type of situation, are certain combinations of facts and plans. In a slogan: mixed uses appear to reduce *global* uncertainty in the common ground without reducing *local* uncertainty, that is, uncertainty within each parameter of the common ground. Moreover, we have only hinted at the way in which mixed uses behave in communication. It has been suggested to me that perhaps such mixed uses are in fact the norm (Paul Égré, p.c.), but I remain doubtful. The reason for this is that we often assume that either we share evaluative or factual knowledge with our interlocutor, and thus proceed to interpret her utterance as either an evaluative or descriptive statement. For instance, if I do not know you and I haven't heard Falla's concerto, I will most likely assume that the chord progression that you are talking about has certain factual properties, and interpret your claim as a judgment of positive evaluation of those properties. Alternatively (but less probably I think), I may assume that you are supposing that everyone shares a certain standard of quality in chord progressions, and that you are informing me of those properties, thus making a descriptive claim about the chord progression.

My proposal leaves a few loose ends, but I hope to have shed some light in the study of such a fascinating parcel of natural language as are adjectives, and especially, evaluatives. Philosophically, I have presented (and tackled) some of the challenges that a compositional implementation of non-factualism about evaluatives faces. Linguistically, I hope to have made a convincing case that extant theories of subjectivity leave something out when it comes to evaluatives—and that the place to look for it is metanormative theory. Future research should aim to account for the full syntactic behavior of evaluatives *vis-à-vis* their communicative function, as well as investigate further connections between evaluatives and other forms of non-factual language, (such as slurs and expressives).

References

- Aloni, Maria (2007). “Free choice, modals, and imperatives”. *Natural Language Semantics* 15.1, pp. 65–94.
- (2018). “FC disjunction in state-based semantics”.
- Alonso-Ovalle, Luis (2006). *Disjunction in alternative semantics*. University of Massachusetts Amherst Amherst.
- Alston, William P (1978). “Meta-ethics and meta-epistemology”. In: *Values and morals*. Springer, pp. 275–297.
- Ayer, A J (1946). *Language, Truth and Logic*. London: Penguin.
- Baker, Derek and Jack Woods (2015). “How Expressivists Can and Should Explain Inconsistency”. *Ethics* 125.2, pp. 391–424.
- Bar-On, Dorit and Matthew Chrisman (2009). “Ethical neo-expressivism”. *Oxford studies in metaethics* 4, pp. 133–165.
- Barker, Chris (2002). “The dynamics of vagueness”. *Linguistics and Philosophy* 25.1, pp. 1–36.
- (June 2009). “Clarity and the grammar of skepticism”. *Mind and Language* 24.3, pp. 253–273.
- (2013). “Negotiating Taste”. *Inquiry* 56.2-3, pp. 240–257.
- Barker, Stephen J (2000). “Is value content a component of conventional implicature?” *Analysis* 60.3, pp. 268–279.
- Bear, Adam and Joshua Knobe (2017). “Normality: Part descriptive, part prescriptive”. *cognition* 167, pp. 25–37.
- Bedke, Matthew S. (2017). “Cognitivism and Non-Cognitivism”. In: *The Routledge Handbook of Metaethics*. Ed. by Tristram MacPherson and David Plunkett. Routledge, pp. 292–307.
- Benthem, J Van (1982). “Later than late: On the logical origin of the temporal order”. *Pacific Philosophical Quarterly* 63.2, pp. 193–203.
- Bierwisch, M (1989). “The semantics of gradation”. In: *Dimensional adjectives*. Ed. by Manfred Bierwisch and Ewals Lang. Berlin: Springer-Verlag, pp. 71–261.
- Björklund, Fredrik et al. (Jan. 2012). “Recent work on motivational internalism”. *Analysis* 72.1, pp. 124–137.
- Björnsson, G. et al., eds. (2015). *Motivational Internalism*. Oxford Scholarship Online.
- Blackburn, Simon (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Clarendon Press.
- (1993). *Essays in quasi-realism*. Oxford University Press.
- (1998). *Ruling passions: a theory of practical reasoning*. Oxford University Press.
- (2006). “The Semantics of Non-factualism, Non-cognitivism, and Quasi-realism”. *The Blackwell Guide to the Philosophy of Language*. Ed. by Michael Devitt and Richard Hanley, pp. 244–252.

- Blome-Tillmann, Michael (2009). "Knowledge and presuppositions". *Mind* 118.470, pp. 241–294.
- Boghossian, Paul A (1990). "The Status of Content". *The Philosophical Review* 99.2, pp. 157–184.
- Boisvert, Daniel R (2008). "Expressive-assertivism". *Pacific Philosophical Quarterly* 89.2, pp. 169–203.
- Bordonaba Plou, David (2017). "Operadores de orden superior y predicados de gusto: Una aproximación expresivista." PhD thesis. Universidad de Granada.
- Brandom, Robert (1983). "Asserting". *Noûs*, pp. 637–650.
- Burnett, Heather (2017). *Gradability in Natural Language : Logical and Grammatical Foundations*. Oxford University Press.
- Bylinina, Lisa (2017). "Judge-dependence in degree constructions". *Journal of Semantics* 34.2, pp. 291–331.
- Camp, Elisabeth (2017a). "Metaethical Expressivism". In: *The Routledge Handbook of Metaethics*. Ed. by Tristram MacPherson and David Plunkett. Routledge, pp. 87–101.
- (2017b). "Pragmatic force in semantic context". *Philosophical Studies* 174.6, pp. 1617–1627.
- Cariani, Fabrizio (2013). "'Ought' and resolution semantics". *Noûs* 47.3, pp. 534–558.
- Carr, Jennifer (2017). "Deontic Modals". In: *Routledge Handbook of Metaethics*. Routledge.
- Cepollaro, Bianca and Isidora Stojanovic (2016). "Hybrid Evaluatives: In Defense of a Presuppositional Account". *Grazer Philosophische Studien* 93.3, pp. 458–488.
- Charlow, Nate (2014). "The problem with the Frege-Geach problem". *Philosophical Studies* 167.3, pp. 635–665.
- (2015). "Prospects for an Expressivist Theory of Meaning". *Philosophers' Imprint* 15.23, pp. 1–43.
- Chrisman, Matthew (2007). "From Epistemic Contextualism to Epistemic Expressivism". *Philosophical Studies* 135, pp. 225–254.
- (2008). "Expressivism, inferentialism, and saving the debate". *Philosophy and Phenomenological Research* 77.2, pp. 334–358.
- (2012). "Epistemic Expressivism". *Philosophy Compass* 7.2, pp. 118–126.
- (2016). "Metanormative Theory and the Meaning of Deontic Modals". *Deontic Modality*, pp. 395–430.
- (2018). "Two nondescriptivist views of normative and evaluative statements". *Canadian Journal of Philosophy* 48.3-4, pp. 405–424.
- Copp, David (2001). "Realist-expressivism: A neglected option for moral realism". *Social Philosophy and Policy* 18.2, pp. 1–43.
- (2009). "Realist-expressivism and conventional implicature". *Oxford studies in metaethics* 4, pp. 167–202.
- Coppock, Elizabeth (2018). "Outlook-based semantics". *Linguistics and Philosophy* 41.2, pp. 125–164.
- Cresswell, Maxwell J. (1976). "The Semantics of Degree". In: *Montague Grammar*. Ed. by Barbara Partee. New York: Academic Press, pp. 261–292.
- Darwall, Stephen, Allan Gibbard, and Peter Railton (1992). "Toward Fin de Siècle Ethics: Some Trends". *The Philosophical Review* 101.1, pp. 115–189.
- Davidson, Donald (1969). "How Is Weakness of the Will Possible?" In: *Moral Concepts*. Ed. by Joel Feinberg. Oxford University Press.

- Del Pinal, Guillermo and Kevin Reuter (2017). "Dual character concepts in social cognition: Commitments and the normative dimension of conceptual representation". *Cognitive science* 41, pp. 477–501.
- Deroose, K (1992). "Contextualism and Knowledge Attributions". English. *Philosophy and Phenomenological Research* 52.4, pp. 913–929.
- DeRose, Keith (Oct. 1991). "Epistemic possibilities". *The Philosophical Review* 100.4, pp. 581–605.
- Dreier, James (1990). "Internalism and speaker relativism". *Ethics* 101.1, pp. 6–26.
- (2009). *Relativism (and expressivism) and the problem of disagreement*.
- Dretske, Fred (1972). "Contrastive statements". *The Philosophical Review*, pp. 411–437.
- Egan, Andy (2012). "Relativist dispositional theories of value". *The Southern Journal of Philosophy* 50.4, pp. 557–582.
- Egré, Paul and Florian Cova (2015). "Moral asymmetries and the semantics of "many"". *Semantics and Pragmatics* 8.13, pp. 1–45.
- Égré, Paul and Jérémy Zehr (2018). "Are gaps preferred to gluts? A closer look at borderline contradictions". In: *The Semantics of Gradability, Vagueness, and Scale Structure*. Springer, pp. 25–58.
- Faroldi, Federico L G and Andrés Soria Ruiz (2017). "The Scale Structure of Moral Adjectives". *Studia Semiotyczne* 31.2, pp. 161–178.
- Field, Hartry (2000). "Apriority as an evaluative notion". *New essays on the a priori*. Ed. by Christopher Peacocke and Paul A Boghossian, pp. 117–149.
- (2009). "Epistemology without metaphysics". *Philosophical Studies* 143.2, pp. 249–290. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- (2018). "Epistemology from an Evaluativist Perspective". *Philosophers' Imprint* 18.12, pp. 1–23.
- Finlay, Stephen (Sept. 2004). "The Conversational Practicality of Value Judgement". *The Journal of Ethics* 8.3, pp. 205–223.
- (2005). "Value and implicature". *Philosophers' Imprint* 5.4, pp. 1–20.
- Fintel, Kai von and Anthony S. Gillies (2011). "'Might' Made Right". In: *Epistemic Modality*. Ed. by Andy Egan and Brian Weatherson. Oxford University Press, pp. 108–130.
- Fintel, Kai von and Sabine Iatridou (2005). "What to do if you want to go to Harlem".
- Fleisher, Nicholas (2013). "The dynamics of subjectivity". In: *Semantics and Linguistic Theory*. Vol. 23, pp. 276–294.
- Fletcher, Guy and Michael Ridge, eds. (2014). *Having it Both Ways: Hybrid Theories and Modern Metaethics*. Oxford University Press.
- Foot, Philippa (1958). "Moral Arguments". *Mind* 67.268, pp. 502–513. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- Frank, Annette (1997). "Context Dependence in Modal Constructions". PhD thesis. Universität Stuttgart.
- Franzén, Nils (2018a). "Aesthetic Evaluation and First-Hand Experience". *Australasian Journal of Philosophy*, pp. 1–14.
- (2018b). "Sense and Sensibility: Four Essays on Evaluative Discourse". PhD thesis. University of Uppsala.
- Frápolli, María José and Neftalí Villanueva Fernandez (2013). "Frege, Sellars, Brandom: expresivismo e inferencialismo semánticos". *Perspectivas en la filosofía del lenguaje* 14, pp. 583–617.
- Frápolli, María José and Neftalí Villanueva (2012). "Minimal expressivism". *Dialectica* 66.4, pp. 471–487.

- García-Carpintero, Manuel (2015). “Contexts as shared commitments”. *Frontiers in psychology* 6, p. 1932.
- García-Carpintero, Manuel and Max Kölbel (2008). *Relative Truth*. Oxford University Press, pp. 1–336.
- Geach, P. T. (Oct. 1965). “Assertion”. *The Philosophical Review* 74.4, pp. 449–465.
- Gibbard, Allan (1981). “Two Recent Theories of Conditionals”. In: *IFS: Conditionals, Belief, Decision, Chance and Time*. Ed. by William L Harper, Robert Stalnaker, and Glenn Pearce. Dordrecht: Springer Netherlands, pp. 211–247.
- (1986). “An Expressivistic Theory of Normative Discourse”. *Ethics* 96.3, pp. 472–485.
- (1990). *Wise Choices, Apt Feelings*. Clarendon Press ; Oxford University Press.
- (2003). *Thinking how to live*. Harvard University Press.
- Glanzberg, M (2007). “Context, content, and relativism”. *Philosophical Studies* 136.1, pp. 1–29.
- Goldberg, Sanford C (2016). “II—Arrogance, Silence, and Silencing”. In: *Aristotelian Society Supplementary Volume*. Vol. 90. 1. Oxford University Press, pp. 93–112.
- Graff, Delia (2000). “Shifting sands: An interest-relative theory of vagueness”. *Philosophical topics* 28.1, pp. 45–81.
- Gutzmann, Daniel (2016). “If expressivism is fun, go for it”. *Subjective Meaning: Alternatives to Relativism*. Ed. by C. Meier and J van Wijnberger-Huitink, pp. 21–46.
- Hamblin, Charles Leonard (1973). “Questions in Montague English”. *Foundations of Language* 10.1, pp. 41–53.
- Hansson, Sven Ove and Till Grüne-Yanoff (2018). “Preferences”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N Zalta. Summer 201. Metaphysics Research Lab, Stanford University.
- Hare, Richard M (1952). *The Language of Morals*. Oxford University Press.
- (1970). “Meaning and speech acts”. *The Philosophical Review* 79.1, pp. 3–24.
- Heim, Irene (1992). “Presupposition projection and the semantics of attitude verbs”. *Journal of Semantics* 9.3, pp. 183–221.
- (2000). “Notes on Interrogative Semantics”, pp. 1–16.
- (2002). “On the Projection Problem for Presuppositions”. In: *Formal Semantics: The Essential Readings*. Ed. by Paul Portner and Barbara Partee. Blackwell Publishing Ltd/Inc., pp. 249–260.
- Heim, Irene and Angelika Kratzer (1998). *Semantics in generative grammar*. Malden, MA: Blackwell.
- Hume, David (1999). *An enquiry concerning human understanding*. Oxford University Press.
- Jackson, Frank and Philip Pettit (1998). “A Problem for Expressivism”. *Analysis* 58.4, pp. 239–251.
- Jaeger, T. Florian (Nov. 2008). “Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models”. *Journal of Memory and Language* 59.4, pp. 434–446.
- Kaplan, D (1989). “An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives”. *Themes from kaplan*.
- Karczewska, Natalia (2016). “Disagreement about Taste as Disagreement about the Discourse: Problems and Limitations”. *Studies in Logic, Grammar and Rhetoric* 46.1, pp. 103–117.
- Karttunen, Lauri (2008). “Syntax and Semantics of Questions”. In: *Formal Semantics: The Essential Readings*, pp. 382–420.
- (2014). “Three ways of not being lucky”. *Slides from SALT* 24.

- Kaufmann, Magdalena (2011). *Interpreting imperatives*. Vol. 88. Springer Science & Business Media.
- Kennedy, Christopher (1999). “Gradable adjectives denote measure functions, not partial functions”. *Studies in the Linguistic Sciences* 29.1, pp. 65–80.
- (2007). “Vagueness and grammar: The semantics of relative and absolute gradable adjectives”. *Linguistics and Philosophy* 30.1, pp. 1–45.
 - (2013a). *Projecting the adjective: The syntax and semantics of gradability and comparison*. Routledge.
 - (2013b). “Two Sources of Subjectivity: Qualitative Assessment and Dimensional Uncertainty”. *Inquiry* 56.2-3, pp. 258–277.
- Kennedy, Christopher and Louise McNally (2005). “Scale Structure, Degree Modification, and the Semantics of Gradable Predicates”. *Language* 81.2, pp. 345–381. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- Kennedy, Christopher and Malte Willer (2016). “Subjective attitudes and counterstance contingency”. *Proceedings of SALT* 26, pp. 0–0.
- Klein, Ewan (1980). “A semantics for positive and comparative adjectives”. *Linguistics and Philosophy* 4.1, pp. 1–45.
- (1991). “Comparatives”. *Semantik: Ein internationales handbuch der zeitgenössischen forschung*. Ed. by A. Stechow and D. Wunderlich, pp. 673–691.
- Kölbel, M (2003). “Faultless disagreement”. *Proceedings of the Aristotelian society* 104.1.
- Kolbel, Max (2002). *Truth without Objectivity*. Vol. 71(2). Routledge.
- Kolodny, Niko and John Macfarlane (2010). “Ifs and Oughts”. *Journal of Philosophy* 107, pp. 1–31. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- Krantz, David H, Patrick Suppes, and Robert Duncan Luce (1971). *Foundations of measurement*. New York academic press.
- Kratzer, Angelika (2012). *Modals and conditionals: New and revised perspectives*. Oxford University Press.
- Krifka, Manfred (1989). “Nominal reference, temporal constitution and quantification in event semantics”. *Semantics and contextual expression*. Ed. by R Bartsch, J Van Benthem, and P Van Ende Boas, pp. 75–115.
- (2014). “Embedding Illocutionary Acts”. *Recursion: Complexity in Cognition* 1995, pp. 59–87.
- Kyle, Brent G. (2013). “How Are Thick Terms Evaluative?” *Philosophers’ Imprint* 13.1.
- Lasersohn, Peter (1999). “Pragmatic halos”. *Language*, pp. 522–551.
- (2005). “Context dependence, disagreement, and predicates of personal taste”. *Linguistics and Philosophy* 28.6, pp. 643–686.
- Lassiter, Daniel (2017). *Graded Modality : Qualitative and Quantitative Perspectives*. OUP Oxford.
- Leslie, Sarah-Jane (2015). ““Hillary Clinton is the only man in the Obama administration”: Dual Character Concepts, Generics, and Gender”. *Analytic Philosophy* 56.2, pp. 111–141.
- Lewis, David (1973). *Counterfactuals*. Blackwell.
- (1979a). “Attitudes de dicto and de se”. *The Philosophical Review* 88.4, pp. 513–543.
 - (1979b). “Scorekeeping in a language game”. In: *Semantics from different points of view*. Springer, pp. 172–187.
 - (1996). “Elusive knowledge”. *Australasian Journal of Philosophy* 74.4, pp. 549–567.
- Lewis, Karen (2017). “Dynamic Semantics”.
- Liao, Shen-yi, Louise McNally, and Aaron Meskin (2016). “Aesthetic adjectives lack uniform behavior”. *Inquiry* 59.6, pp. 618–631.

- Liao, Shen-yi and Aaron Meskin (2015). “Aesthetic Adjectives: Experimental Semantics and Context-Sensitivity”. *Philosophy and Phenomenological Research*, n/a–n/a.
- López de Sa, Dan (2015). “Expressing Disagreement: A Presuppositional Indexical Contextualist Relativist Account”. *Erkenntnis* 80.1, pp. 153–165.
- Ludlow, Peter (2008). “Cheap Contextualism”. *Philosophical Issues* 18, pp. 104–129.
- MacFarlane, John (2014). *Assessment sensitivity : relative truth and its applications*. Oxford University Press.
- (June 2016). “Vagueness as Indecision”. *Aristotelian Society Supplementary Volume* 90.1, pp. 255–283.
- Maldonado, Mora, Emmanuel Chemla, and Benjamin Spector (2019). “Revealing abstract semantic mechanisms through priming: The distributive/collective contrast”. *Cognition* 182, pp. 171–176.
- Marques, Teresa (2017). “What Metalinguistic Negotiations Can’t Do”. *Phenomenology and Mind* 0.12, pp. 40–48.
- McDowell, John (2002). “Non-cognitivism and rule-following”. In: *The New Wittgenstein*. Routledge, pp. 48–62.
- McNally, Louise and Isidora Stojanovic (2017). “Aesthetic Adjectives”. In: *The Semantics of Aesthetic Judgment*. Ed. by J. Young. Oxford University Press.
- Miller, A (2003). *An introduction to contemporary metaethics*. Cambridge: Polity Press.
- Moltmann, Friederike (2006). “Generic one, arbitrary PRO, and the first person”. *Natural language semantics* 14.3, pp. 257–281.
- (Sept. 2010). “Relative truth and the first person”. *Philosophical Studies* 150.2, pp. 187–220. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- Moore, G. E. (1993). *Principia Ethica*. Ed. by Thomas Baldwin. Cambridge University Press.
- Partee, Barbara and Mats Rooth (2002). “Generalized conjunction and type ambiguity”. *Formal semantics: the essential readings*. Ed. by Paul Portner and Barbara Partee, pp. 334–356.
- Pérez Carballo, Alejandro (2011). “Rationality without representation”. PhD thesis. MIT.
- (2014). “Semantic hermeneutics”. In: *Metasemantics: New essays on the foundations of meaning*. Ed. by Alexis Burgess and Brett Sherman. Oxford University Press, USA, pp. 119–146.
- Pérez Carballo, Alejandro and Paolo Santorio (2016). “Communication for expressivists”. *Ethics* 126.3, pp. 607–635.
- Plunkett, David and Tim Sundell (2013). “Disagreement and the semantics of normative and evaluative terms”. *Philosophers’ Imprint* 13.23, pp. 1–37.
- Portner, Paul (2004). “The semantics of imperatives within a theory of clause types”. In: *Semantics and linguistic theory*. Vol. 14, pp. 235–252.
- Portner, Paul and Aynat Rubinstein (2016). “Extreme and non-extreme deontic modals”. *Deontic modals*. Ed. by Charlow Nate and Matthew Chrisman, pp. 256–282.
- Recanati, François (1987). *Meaning and force: The pragmatics of performative utterances*. CUP Archive.
- Rett, Jessica (2007). “Antonymy and evaluativity”. In: *Proceedings of SALT XVII*. Ed. by T. Friedman and M. Gibson. Ithaca, NY: Cornell University, pp. 210–227.
- Ridge, Michael (2006). “Ecumenical expressivism: finessing Frege”. *Ethics* 116.2, pp. 302–336.
- (2013). “Disagreement”. *Philosophy and Phenomenological Research* 86.1, pp. 41–63.
- (2014). *Impassioned belief*. Oxford University Press.
- Roberts, Debbie (2013). “Thick concepts”. *Philosophy Compass* 8.8, pp. 677–688.

- Rooth, Mats (1992). “A theory of focus interpretation”. *Natural language semantics* 1.1, pp. 75–116.
- Rosen, Gideon (1998). “Critical study of Blackburn’s Essays in Quasi-Realism”. *Noûs* 32.3, pp. 386–405.
- Rothschild, Daniel (2012). “V—Expressing Credences”. In: *Proceedings of the Aristotelian Society*. Vol. 112. Wiley Online Library, pp. 99–114.
- Ruiz, Andrés Soria and Isidora Stojanovic (2019). “On Linguistic Evidence for Expressivism”. *Royal Institute of Philosophy Supplement* 86, pp. 155–180.
- Sæbø, Kjell Johan (2009). “Judgment ascriptions”. *Linguistics and Philosophy* 32.4, pp. 327–352.
- Sassoon, Galit W. (2013a). “A typology of multidimensional adjectives”. *Journal of Semantics* 30.3, pp. 335–380.
- (2013b). *Vagueness, gradability and typicality: The interpretation of adjectives and nouns*. Brill.
- (2016). “Multidimensionality in the grammar of gradability”. arXiv: [arXiv : 1011 . 1669v3](https://arxiv.org/abs/1011.1669v3).
- Schaffer, J (2011). “Perspective in taste predicates and epistemic modals”. In: *Epistemic Modality*. Ed. by Andy Egan and Brian Weatherson. Oxford University Press, pp. 179–227.
- Schellekens, Elisabeth (2008). “Three Debates in Meta-Aesthetics”. In: *New Waves in Aesthetics*. Springer, pp. 170–187.
- Schroeder, Mark (2008a). *Being for: Evaluating the semantic program of expressivism*. Oxford University Press.
- (2008b). “Expression for Expressivists”. *Philosophy and Phenomenological Research* 76.1, pp. 86–116.
- (2008c). “How expressivists can and should solve their problem with negation”. *Nous* 42.4, pp. 573–599.
- (2008d). “What Is the Frege-Geach Problem?” *Philosophy Compass* 3.4, pp. 703–720.
- (2009). “Hybrid Expressivism: Virtues and Vices”. *Ethics* 119.2, pp. 257–309.
- Silk, Alex (2015). “How to be an ethical expressivist”. *Philosophy and Phenomenological Research* 91.1, pp. 47–81.
- (2016). *Discourse Contextualism: A framework for contextualist semantics and pragmatics*. Oxford University Press.
- (2017). “Evaluational Adjectives”. *Unpublished manuscript*.
- Smith, E. Allyn, Elena Castroviejo, and Laia Mayol (2015). “Cross-Linguistic Experimental Evidence Distinguishing the Role of Context in Disputes over Taste and Possibility”. *International and Interdisciplinary*.
- Smith, Michael (1994). *The Moral Problem*. Blackwell.
- Sode, Frank (2018). “good as a predicate of worlds”. In: *Proceedings of Sinn und Bedeutung* 22. Ed. by Uli Sauerland and Stephanie Solt. Berlin: ZAS, pp. 407–424.
- Solt, Stephanie (2015). “Scales in natural language”. *Language and Linguistics Compass* 9.1, pp. 14–32.
- (2018). “Multidimensionality, subjectivity and scales: experimental evidence”. In: *The Semantics of Gradability, Vagueness, and Scale Structure*. Ed. by Elena Castroviejo, Louise McNally, and Galit Weidman Sassoon. Springer, pp. 59–91.
- Stalnaker, Robert (2002). “Assertion”. In: *Formal Semantics: The Essential Readings*. Ed. by Paul Portner and Barbara Partee. Blackwell Publishing Ltd/Inc., pp. 147–161.
- (2014). *Context*. Oxford University Press.

- Starr, William B (2016). “Dynamic Expressivism about Deontic Modality”. In: *Deontic Modality*. Ed. by Nate Charlow and Matthew Chrisman. New York: Oxford University Press, pp. 355–394.
- (2018). “A Preference Semantics for Imperatives”.
- Stephenson, Tamina C (2007a). “Judge dependence, epistemic modals, and predicates of personal taste”. *Linguistics and Philosophy* 30.4, pp. 487–525.
- (2007b). “Towards a Theory of Subjective Meaning”. PhD thesis. MIT.
- Stevenson, Charles Leslie (1937). “The Emotive Meaning of Ethical Terms”. *Mind* 46.181, pp. 14–31.
- (1963). *Facts and Values: Studies in Ethical Analysis*, p. 244.
- Stine, G C (1976). “Skepticism, Relevant Alternatives, and Deductive Closure”. English. *Philosophical Studies* 29.4, pp. 249–261.
- Stojanovic, Isidora (2007). “Talking about taste: Disagreement, implicit arguments, and relative truth”. *Linguistics and Philosophy* 30.6, pp. 691–706.
- (2016). “Expressing aesthetic judgments in context”. *Inquiry*, pp. 1–23.
- (2017a). “Context and disagreement”. *Cadernos de Estudos Lingüísticos* 59.1, pp. 9–22.
- (2017b). “Metaethical Relativism”. In: *The Routledge Handbook of Metaethics*. Ed. by Tristram McPherson and David Plunkett. Routledge, pp. 119–132.
- (2018). “An Empirical Approach to Aesthetic Adjectives”. In: *Advances in Experimental Philosophy of Aesthetics*. Ed. by Florian Cova and Sébastien Réhault. Bloomsbury Publishing.
- (2019). “Disagreements about Taste vs. Disagreements about Moral Issues”. *American Philosophical Quarterly* 56.1, pp. 29–42.
- Strandberg, Caj (2011). “The pragmatics of moral motivation”. *The Journal of ethics* 15.4, pp. 341–369.
- Sundell, Tim (2016). “The tasty, the bold, and the beautiful”. *Inquiry* 59.6, pp. 793–818.
- Szabolcsi, Anna (1982). “Model Theoretic Semantics of Performatives”. In: *Hungarian General Linguistics*, pp. 515–535.
- Tappolet, Christine (2013). “Evaluative vs. Deontic Concepts”. In: *International Encyclopedia of Ethics*.
- Umbach, Carla (2016). “Evaluative propositions and subjective judgments”. In: *Subjective Meaning: Alternatives to Relativism*. Ed. by Cécile Meier and Janneke van Wijnbergen-Huitink. De Gruyter, pp. 127–168.
- (2019). “Evaluative Predicates beyond fun and tasty”.
- Väyrynen, Pekka (2013). *The Lewd, the Rude, and the Nasty*. Oxford University Press. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- (2017). “Thick Ethical Concepts”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N Zalta. Fall 2017. Metaphysics Research Lab, Stanford University.
- Veltman, Frank (1996). “Defaults in update semantics”. *Journal of philosophical logic* 25.3, pp. 221–261.
- Von Fintel, Kai and Sabine Iatridou (2008). “How to say ought in foreign: The composition of weak necessity modals”. In: *Time and modality*. Springer, pp. 115–141.
- Willer, Malte (2017). “Advice for Noncognitivists”. *Pacific Philosophical Quarterly* 98, pp. 174–207.
- Williams, Bernard (1987). *Ethics and the limits of Philosophy*. Vol. 8. 6. Routledge.
- Wittgenstein, Ludwig (1965). “A lecture on ethics”. *The Philosophical Review* 74.1, pp. 3–12.
- Wright, G.H. von (1972). “On So-Called Practical Inference”. *Acta Sociologica* 15.1, pp. 39–53.

- Yalcin, Seth (2007). “Epistemic modals”. English. *Mind* 116.464, pp. 983–1026.
- (2011). “Nonfactualism about epistemic modality”. In: *Epistemic Modality*. Ed. by Andy Egan and B Weatherson. Oxford University Press.
 - (2012). “Bayesian Expressivism”. *Proceedings of the Aristotelian Society (Hardback)* 112.2pt2, pp. 123–160.
 - (2018). “Expressivism by force”. In: *New Work on Speech Acts*. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford University Press.
 - (2019). “Modeling with Hyperplans”.
- Zvolenszky, Zsófia (2002). “Is a possible-worlds semantics of modality possible? A problem for Kratzer’s semantics”. In: *Semantics and Linguistic Theory*. Vol. 12, pp. 339–358.

RÉSUMÉ

Cette thèse porte sur les adjectifs évaluatifs: bon, mieux, mauvais, beau, beau, amusant ou cruel. Je défends une forme de non-factualisme, à savoir l'idée que les expressions évaluatives n'offrent pas d'informations sur le monde. Au lieu de cela, ils expriment et coordonnent les engagements des interlocuteurs en matière d'action. Mais les adjectifs évaluatifs sont graduels, autrement dit ils acceptent différents degrés. Par conséquent, la théorie non-factualiste doit être formulée de manière à ce que les expressions évaluatives puissent exprimer des engagements pratiques avec une structure scalaire appropriée.

MOTS CLÉS

Évaluation • Jugements évaluatifs • Expressivisme • Adjectifs • Gradation • Subjectivité

ABSTRACT

This is a dissertation about words like good, better, bad, beautiful, fun or cruel. I defend a form of non-factualism, which is the view that evaluative sentences do not offer information about the world. Instead, they express and coordinate commitments about action. But evaluative adjectives are gradable, which means that they accept different degrees. Therefore, non-factualism has to be formulated in a way that evaluative sentences can express practical commitments that have the appropriate scalar structure.

KEYWORDS

Evaluation • Evaluative judgment • Expressivism • Adjectives • Gradability • Subjectivity

