THE UNIVERSITY OF CHICAGO

SOURCES OF SUBJECTIVITY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE HUMANITIES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF LINGUISTICS

BY
TAMARA NIKOLAEVNA VARDOMSKAYA

CHICAGO, ILLINOIS
JUNE 2018

To the protagonist of the novel I didn't write because I was writing this.

There's a possible world where you're real.

It precludes the possible world where you read this.

I agreed that what really matters is what you *like*, not what you *are like*... Books, records, films – these things matter. Call me shallow but it's the fuckin' truth.

— High Fidelity (2000)

TABLE OF CONTENTS

ACKNOWLEDGMENTS

First of all, I would like to thank my dissertation advisor, Chris Kennedy. Ever since, on my acceptance to graduate school at the University of Chicago, he became a person rather than half of Kennedy & McNally 2005, he has been the finest advisor, instructor, supervisor, coach who pushes me to write and reason better than I thought I could, supporter, promoter and fellow fan that can be imagined. Sometimes he is too good a professor to be real. And there were times when I stuck the course simply because Chris needs someone around who would appreciate his science fiction references.

Itamar Francez has been a source of advice, argument, and support both on pragmatics and on morale. I also thank his wonderful family.

Malte Willer was a splendid mentor for avoiding linguists reinventing the wheel that had rolled and broken a hundred times in the philosophy world. I think the voice of sanity in my head is Malte's voice.

Anastasia Giannakidou was my cheerleader from the start, and her enthusiasm for my work and for the semantics of modality never wavered. Though I am sorry for not actually using any examples from Greek.

Thank you to the other professors in the linguistics department who have helped me in innumerable ways: Lenore Grenoble, Amy Dahlstrom, Salikoko Mufwene, Greg Kobele, Alan Yu, Jason Riggle, Jason Merchant, Karlos Arregi, Jerry Sadock, John Goldsmith.

To my cohort, my fellow team of CLS 49, we always have been and always will be: Helena Aparicio Terrasa, Gallagher Flynn, Katie Franich, Asia Pietraszko, Mike Pham and Diane Rak. You have made my time in the linguistics department unique and wonderful. Also to the other linguistics and allied fields students of Chicago for years of advice, friendship, and adventure: Carissa Abrego-Collier, Anais Ahmed, Rebekah Baglini, Andrea Beltrama, Ryan Bochnak, Tasos Chatzikonstantinou, Kate Christensen, Carlos Cisneros-Juarez, Emily Coppess, Marina Ermolaeva, Ksenia Ershova, Josh Falk, Jeff Geiger, Tim Grinsell, Julian Grove, Emily Hanink, Jonathan Keane, Peter Klecha, Alex Kramer, Jackson Lee, Robert

Lewis, Yaron McNabb, Cherry Meyer, Patrick Munoz, Chieu Nguyen, Jacob Phillips, Adam Singerman, Matt Teichman, Julia Thomas Swan, Woody Wu, Orest Xherija, and Anqi Zhang.

To other linguists for swapping ideas that added to this research and to my graduate success: Chris Barker, Eva Csipak Bochnak, Heather Burnett, Elizabeth Coppock, Gwendolyn Gillingham, Thomas Graf, Hadas Kotek, Peter Lasersohn, John MacFarlane, Eric McCready, Louise McNally, Walter Pedersen, Rachel Rudolph, E. Allyn Smith, Kjell Johann Saebø, Galit Sassoon, and Sarah Zobel.

I couldn't have lived through those years in Chicago without music, so I would like to thank the friends I made through vocal studies — MurphyKate Montee, Grace Tsiang, and especially Patrice Michaels — and the people in University Ballet for the three wonderful years.

For the last two and a half years of dissertating, I especially thank the University of Chicago quiz bowl team for reminding me that besides subjectivity, there is also Hurrian mythology and Chilean literature and Baroque music and postmodern art installations: Alston Boyd, Halle Friedman, Athena Kern, Samir Khan, Olivia Kiser, Teddy Knox, Abby Kuchnir, James Lasker, John Y. Lawrence, Matthew Lehmann, Chris Ray, Federico Scivittaro, Francesco Scivittaro, Max Schindler, Kai Smith, Sam Winikow, Marianna Zhang, and Jason Zhou. Also the friends I have made through quiz bowl outside Chicago: Rob Carson, Naveed Chowdhury, Auroni Gupta, Andrew Hart, Will Nediger, Tejas Raje, Jacob Reed, Ryan Rosenberg, Andrew Wang, and many others present and future. Thank you for this game of forty conversations about things someone may like, and for all the conversations afterwards. When all my life in the late stages of a Ph.D seemed like a conversation about one thing, this mattered.

For dear friendship in Chicago, I also must thank the Palmer-Schiller bash' and their allies: Michael Lueckheide, Michael Mellas, Ada Palmer, Lauren Schiller, and Jonathan Sneed.

Thank you to Juana M. Liceras, my first mentor and patron in the linguistics world, and to her son Marco Llamazares, who remains my oldest friend. And to Jonathan Crowe and Jennifer Seely, for all the weekends in Shawville, sanity checks and food and books and movies and photography and snakes and Goober, Doofus and Scourge. May Goober showing up in my examples in this dissertation be a fitting legacy for that charismatic megafauna.

Thank you to my parents, Nikolai Vardomskiy and Valentina Vardomskaya, and my brothers Stanislav Vardomskiy and Boris Vardomskiy, for support. And to Ambrose, Tania and Benjamin Tam, for taking me into their home and family.

And last and most of all, I thank Nicholas Tam, for bringing beauty I could not have imagined into my life and for urging me to bring it back into my writing, for care and compassion, for adventure and daring and repose and laughter. (And for doing line edits on this thing; all remaining typos, stylistic and grammatical errors, misquotations, and flat-out lies are mine, of course.) For all the faultless disagreements we had and have, and the love anyway.

ABSTRACT

Subjectivity is the phenomenon of the apparent truth of a predicate depending on a perspective of evaluation, such that one person may sincerely assert a proposition p while another may sincerely assert *not-p*. Among the numerous analyses of the semantics of subjective predicates (Lasersohn 2005, Stojanovic 2007, Stephenson 2007, MacFarlane 2014, Barker 2002, a.o.), few consider why they differ from objective ones: what makes *delicious* allow faultless disagreement while *wooden* or *red* do not? Assumptions that subjective and objective predicates differ In their semantics (do not have truth conditions as per expressivism, have another index or argument as per relativism or contextualism) ignore the fact that the same predicate may be subjective in a context where it is loosely defined and objective in a context where it is stringently defined. E.g. the truth of *good figure skater* is objective to trained figure skating judges but subjective to casual TV watchers.

I provide a relatively theory-neutral analysis of what makes subjective predicates what they are. I argue that objective predicates are precisely those for which there is a reliable consensus of what evidence matters (to distinguish from a reliable consensus as to whether propositions containing them are true: we do not know whether there is life on other planets, but we know what it would take to prove it). For subjective predicates, and propositions containing them, there is no reliable and socially enforced consensus as to what evidence matters, and how much, and what does not, and for some predicates, there cannot be. Thus, speakers are allowed, in a pragmatic context, to perceive the evidence differently (to have different taste perceptions due to genetic differences in smell receptors) or to classify it differently (looking at a painting, to judge whether it is excellent or poor based on differing prior expertise in painting). If we allow differing perceptions or different categorizations to be valid, we have a subjective predicate.

As a follow-up, I explain the selection criteria of *find (NP) (Predicate)* - "I find the soup disgusting/wonderful" - which is known (Saebø2009 a.o.) to select for subjective constructions. I argue that *find* actually selects for direct experience of its object, as was proposed

by Stephenson (2007), and I address subsequent criticisms of that analysis and extend it to modal expressions such as "I find the Cubs winning unlikely," which had not been previously considered in the literature. I conclude by showing how my analysis fits into different theories (expressivism, relativism, contextualism, metalinguistic negotiation) by providing them with clearer selection criteria for not only what a subjective predicate is, but why it is so.

CHAPTER 1

THE SEMANTICS OF SUBJECTIVITY

ROSALIND: Where learned you that oath, fool?

TOUCHSTONE: Of a certain knight that swore by his honor they were good pancakes, and swore by his honor the mustard was naught. Now, I'll stand to it, the pancakes were naught and the mustard was good, and yet was not the knight forsworn.

— William Shakespeare, *As You Like It*, Act I, Scene 2, lines 51-52.

1.1 Two kinds of predicates

Semanticists and philosophers, and indeed ordinary users of language, generally assume that words like *tasty, fun, interesting* are different from *wooden, extinct,* and others. The former are words that do not depend on 'matters of fact' but simply on 'matters of opinion.' These words are known as *subjective predicates* and propositions that contain them are known as *subjective propositions*; predicates and propositions that do not meet such criteria — that is, disagreements about them are factual disagreements where at least one party must be wrong — are known as *objective*.

These words pose challenges to truth-conditional semantics because of the phenomenon of faultless disagreement: A can assert P, B can assert not-P, and they would be perceived as disagreeing but neither would be saying something false by outside observers — yet people within the dialogue would treat P and not-P as being true or false. This is shown in the following examples, variants on which are a staple of the subjectivity literature:

(1) Sue: Avocados are delicious.
Zoe: No, they're not! – SUBJECTIVE

Compare with the following example:

(2) Sue: Avocados are vegan.

Zoe: No, they're not vegan! — OBJECTIVE (Zoe is wrong; avocados contain no animal products)

Readers may naturally be skeptical that such examples sound very artificial: across paper after paper, John and Mary, Kim and Lee, Tarek and Irma, Yuri and Zelda are endlessly arguing about various foods. So I provide an example of such a disagreement occurring in natural conversation, where a statement that is seemingly about a personal opinion is treated as true or false, with different parties disagreeing:

(3) (Example in an online chat between University of Chicago students, used with permission, October 2016):

A: *Diners, Drive-Ins and Dives* is a fantastic show.

B: Literally false.

C: Literally true.

In this example, *fantastic* is such a subjective predicate. B is asserting that A has said something false ("literally" so). Intuitively, B is disagreeing with A, and C is disagreeing with B (and agreeing with A). However, despite B's assertion, we as external observers do not feel that A, B or C are somehow in the wrong or mistaken. Even as B takes A to be saying something false, an outside observer would be reasonable to claim that no, A is not, but neither is B.

A similar example in classical literature is the quotation from Shakespeare that serves as this chapter's epigraph. Touchstone's knight is treating *They were good pancakes and the mustard was naught [not good at all]* as a statement on the truth of which one can stake one's honor (let's say, one's reliability in telling the truth). Touchstone claims that it was a false statement (*the pancakes were naught and the mustard was good*) but the knight did not have to give up what he'd sworn by, his honor, because he didn't have any by that point since he had already made many similar false assertions. This assumption relies on statements

about the goodness of pancakes and mustard having truth conditions and being true or false, Touchstone and the knight disagreeing on it — yet even Touchstone qualifies this by *I'll stand to it*, putting his own judgement of the pancakes and mustard as the correct one for evaluating truth or falsehood. An outside observer can claim that neither Touchstone nor the knight are actually saying something false, because pancakes and mustard being good/delicious is a matter of taste.

But note that if *good* and *naught* describing pancakes and mustard is taken to mean 'fresh' and 'spoiled' rather than 'delicious' and 'not delicious,' Touchstone has a much firmer claim that the knight is actually saying something false, and thus losing his honor if he has any. How these two interpretations of *good* differ is the problem at the heart of this dissertation. In Chapter 5, I propose an answer.

How to model this problem — that propositions containing subjective predicates can seem true, false or "it depends" depending on whose perspective is considered — has been the subject of many papers and debates. In this chapter, I summarize the prevailing models that seek to find space within truth-conditional semantics to model predicates that inspire such disagreements, and in the following chapter, I add some observations and attempts to model the pragmatics of these predicates.

Things that fall into the category of subjective predicates, in the literature, include epistemic modals (Stephenson, 2007a); predicates of personal taste (Lasersohn, 2005); evaluative and aesthetic predicates and those of moral judgement (many going back to Kant (1790) and noncognitivist philosophers in the 1930s through the 1960s such as Ayer (1936) and Stevenson (1937)); vague predicates and multidimensional predicates (Kamp, 1975; Klein, 1980; Bierwisch, 1989; Bylinina, 2013).

However, most discussions of subjective predicates try to model subjectivity, but don't explain *why* a predicate falls into one class rather than another. As Lasersohn (2005) remarks, "Exactly which predicates qualify as predicates of personal taste is an interesting question...I do not offer here any firm diagnostic criteria for identifying predicates of personal taste,

though I will return to this question briefly in Section 7.2, below. Despite the absence of such criteria, I think the intuitive idea should be reasonably clear, and that at least some examples, including *fun* and *tasty*, may be identified on that basis." (p. 2-3). In Section 7.2, he concludes, "I see little reason to expect that subjective predicates may be identified by any straightforward linguistic test; it may be that the status of predicates must be argued on more on philosophical than linguistic grounds, on a case-by-case basis." (p. 35).

An issue with many models, as I show in the following sections, is that they take the difference between subjective and objective predicates as something "built-in" to the semantics (and argue *how* it is built-in) rather than something that may shift with the context. Whereas in natural speech, we see that there may be propositions which shift in whether they are subjective or not depending on the context. While the example of Touchstone's *good* meaning either 'delicious' or 'not rotten' can be ascribed to polysemy of a single lexical item, Pedersen (2012) gives an example of faultless disagreement with no clear lexical item triggering the subjectivity:

(4) (At the 2012 U.S. presidential election):
A: Obama won the debate last night.
B: No, he didn't. Obama didn't win the debate. Romney won the debate.

This disagreement is faultless, whereas an almost identical case is not:

(5) (At the World Universities Debating Championships):
A: Sharon won the debate last night.
B: No, she did not. Sharon didn't win the debate. Wallis won the debate.

In that case, the context makes the disagreement objective: the World Debating Championships, unlike presidential debates, have judges that officially decide winners and losers to every debate round. Yet we cannot pin down that the difference between (4) and (5) is

due to a particular word that has particular features (an indexical, a hidden argument, a first-person generic, etc.) associated with it.

I argue in the following chapters that subjectivity can be predicted in the pragmatic context and should be treated separately from truth conditions, but to set the stage, this chapter will provide background for how semanticists and philosophers have dealt with the issue.

First, in Section 1.2, I present the four accepted indicators of subjective propositions in English: faultless disagreement, *find*-embedding, and 'judge' prepositional phrase modifiers. (Although subjective propositions seem to be common and perhaps universal across languages, the majority of research using philosophy of language and the generativist framework has been done on English; I will incorporate examples from German, Japanese and Russian where needed in this dissertation.) Then in Section 1.3 I present some further puzzles that have influenced the development of the theories I describe: the problem of exocentricity, the definition of disagreement, the problem of retraction, and the puzzle of whether there is one kind of subjectivity or multiple kinds.

Section 1.4 is the meat of the chapter; in it, I give a tour of the major theories on the market based on whether they have addressed subjectivity as residing in the speech act, the Kaplanian context, the Kaplanian content, the metalinguistic pragmatics, or elsewhere. To keep this chapter from becoming the length of a dissertation by itself, my overviews will be generally focused on brief introductions to each theory and the justifications behind it, criticisms from other theories, and what this theory says about the source of subjectivity, if anything. Section 1.5 is a much shorter overview of the general criticisms and a look forward to the pragmatics issues that we will discuss in the following chapter. I conclude in Section 1.6.

1.2 How to tell a subjective predicate

1.2.1 Faultless disagreement

First of all, let us define what we intuitively mean when we say about a predicate, "That's subjective." *Subjectivity* is the phenomenon of the apparent truth of a predicate depending on a perspective of evaluation, such that one person may sincerely assert a proposition p while another may sincerely assert $\neg p$. Thus, the phenomenon of faultless disagreement arises. Sue may assert *Mint is tasty* while Bob asserts *Mint is not tasty*, without either being judged wrong by an external observer.

The opposite of *subjectivity* is *objectivity*. If a predicate p' is objective, then if one person asserts p' while another asserts $\neg p'$, one of them is right (is saying something true) while the other is saying something false:

(6) Sue: 27 is a prime number.
Bob: No, 27 isn't a prime number.

Bob is correct and Sue is wrong.

Let's keep in mind that predicates may be objective even if no external observer knows which speaker is right:

(7) Sue: There is life on other planets, somewhere.
Bob: No, there's no life on other planets, we are alone in the universe.

Even though (barring alien abduction reports) as I write this sentence we do not know which of Bob and Sue are correct, we know that there is a fact of the matter, and we have some idea of the evidence required to prove it: say, a probe finding bacteria on Mars or Europa would immediately make it that Sue is saying the truth and Bob's utterance is false. However, in the case of whether mint is tasty or not, both Bob and Sue seem to be telling the truth.

(If they have both tasted mint, that is; if Sue says that mint is tasty without ever having actually tasted it, we get the sense that she is prevaricating or being insincere; see Chapter 2 for much more discussion of this).

1.2.2 Faultless disagreement in the comparative

Kennedy (2013) points out a stronger diagnostic for predicates of personal taste: only typically-subjective predicates allow *faultless disagreement in the comparative form.*

(8) a. Zelda: Skiing is the most fun.
Yuri: No, knitting is!
b. Zelda: Mary is smarter/taller than John.
Yuri: No, John is smarter/#taller than Mary.

Although one may have subjective disagreement over the positive form of *tall*, *taller* allows only objective disagreement — John's and Mary's heights may be easily measured and compared, even by such a simple test as standing the two of them next to each other.

1.2.3 Find and subjective attitude verbs

An indicator that distinguishes the two is that many languages have an embedding verb that will only embed subjective predicates. In English, this is *find (X) (Pred)*:

(9) a. Sue finds mint tastes good.
b. #Sue finds the Amur leopard is extinct.
c. #Sue finds it's Tuesday.

Although this observation has been known since (Borkin, 1973), Saebø (2009) sparked a new wave of discussion by placing its analysis within modern semantics models. He argues that this restriction on *find* is evidence for a contextualist view (of which more below and

much more in Chapter 4). Saebø adds that similar verbs in other languages, such as French *trouver*, Norwegian *synes*, German *finden*, Swedish *tycker*, and Mandarin Chinese *jué dé* share the same property, and he dubs the class *subjective attitude verbs*. Even though otherwise subjective predicates seem to syntactically behave just like ordinary predicates, those two factors — the apparent dependence of their truth on the judge, and the ability to be embedded under subjective attitude verbs — present a puzzle that models need to account for. I offer my own account in Chapter 4.

1.2.4 Judge prepositional phrases

This leads us to the fourth (and least consistent) diagnostic in English, that of the *'judge PP'* (Lasersohn, 2005; Stephenson, 2007a). Certain subjective predicates allow a prepositional phrase (PP) to specify who the truth of the predicate should be 'relativized' to:

(10) a. Wine is tasty to Yuri, while roller coasters are fun for Zelda.

Once that PP is overt, faultless disagreement is no longer felicitous:

(11) Yuri: The wine is tasty to me.
Zelda: #No, it's not!
Zelda: Well, it's not tasty to *me*.

However, as Stephenson (2007a) and Bylinina (2013) point out, this 'judge PP' is not available to all predicates that pass the other tests for subjectivity, and which preposition is used seems to be idiosyncratic.

(12) a. fun for Sue / *fun to Sue
b. *tastes good for Sue / tastes good to Sue
c. *smart for Sue / *smart to Sue
d. ?beautiful for Sue / ?beautiful to Sue

Thus the question of what role this prepositional phrase occupies is one that a model of subjectivity in English needs to account for. Certain languages such as Russian and Hungarian (Bylinina, 2013) seem to have an equivalent in a dative construction showing the 'judge', which seems to work for subjective but not objective predicates (and in case of Japanese, has an additional first-person restriction), as shown here for Russian, Japanese and Hungarian:

(13) Russian:

a. Mne jeta kniga interesnaya/#russkaya/#chornaya.
me-DAT this book interesting/Russian/black
"This book is interesting/#Russian/#black to me."

(14) Japanese (Bylinina, 2013, 2016):

a. watasi-ni-(totte)-wa / *John-ni-(totte)-wa kono keikii-wa oisii
I-DAT-TOTTE-TOP / John-DAT-TOTTE-TOP this cake-TOP tasty
"This cake is tasty to me / *to John."

b. John-ni-(totte)-wa kono keikii-wa oisii noda / ni
John-DAT-TOTTE-TOP this cake-TOP tasty EVID / there's.no.mistake
tiganai

"The cake must be tasty to John."

(15) Hungarian (Bylinina, 2013, 2016):

a. János-nak fontos-ak vagyunk
John-DAT important-PL be.1PL
'John finds us important.'

b. János-nak minden felńott magas
John-DAT every adult tall
'For John, all grown-ups are tall.'

However, the dative arguments do not coherently and reliably map to English prepositional phrases either, as for example, the Russian equivalent of *this is tasty to me* can only take a dative in the adverbial but not the adjectival form, as I illustrate below:

(16) a. mne vkusno
me-DAT delicious-ADV
'(It/Some salient object I am consuming) is delicious to me.'

b. (po-moemu), avokado vkusnoje
in.my.opinion avocado-NOM delicious-NEUT
'Avocado is delicious (in my opinion).'

c. *avokado mne vkusnoje
avocado-NOM me-DAT delicious-NEUT
Intended: 'Avocado is delicious to me.'

d. *avokado mne vkusno
avocado-NOM me-DAT delicious-ADV
Intended: 'Avocado is delicious to me.'

So although these tests are not necessarily all reliable, they have all led to much conversation concerning subjectivity. In the next section, I give some of the main challenges faced by theories of subjectivity armed with these tests.

1.3 An introduction to the puzzles

Before I plunge into the details of what various theories have proposed for the difference between subjective expressions (including moral terms such as *wrong*, evaluative adjectives such as *beautiful*, predicates of personal taste such as *fun* or *delicious*, vague predicates such as *tall* or *rich* and epistemic modals such as *likely* or *might*) and objective expressions, let me introduce a few more puzzles that analyses of this class need to account for.

1.3.1 Exocentricity

First is the problem of *exocentricity*: for predicates of personal taste, one can (debatably) use them where the judge/assessor/taste holder is not the speaker. Stephenson has a famous example of *tasty* as uttered by B reflecting the taste of B's cat (while B herself would probably find the cat food disgusting):

(17) A: How's that new cat food you bought?

B: It must be tasty, because the cat ate almost all of it.

Similarly, there is Lasersohn's (2009) example of the adult John, when telling Mary about his trip to the amusement park with his four-year-old son Bill, using *fun* and *scary* as Bill would perceive it:

(18) Mary: How did Bill like the rides?

John: Well, the merry-go-round was fun, but the water slide was kind of scary.

Here, John as an adult would consider the merry-go-round boring, and the water slide not at all frightening, but he is using these terms to apply to Bill's perceptions.

This is true of predicates of personal taste, but becomes more difficult with epistemic modals. Although the use of an epistemic modal seems to be anchored to a particular knowledge state, it is difficult to make that knowledge state someone's other than the speaker's. Stephenson (2007a) contrasts the cat food example above with the following dog food example:

(19) Mary: Wow, the dog really likes the dog food you're feeding him.

Sam: (#)Yeah, I think it might be table scraps.

It is very difficult to coerce the judge of *might be* to be the dog in this example (that is, the dog believes that there is an accessible possible world where the dog food is table scraps), whereas it had been fairly easy to coerce the judge of *tasty* to be the cat.

However, Egan *et al.* (2005) claim there is a use of an exocentric modal in the following case:

(20) [Context: Ann is planning a surprise party for Bill. Unfortunately, Chris has discovered the surprise and told Bill all about it. Now Bill and Chris are having fun

watching Ann try to set up the party without being discovered. Currently Ann is walking past Chris's apartment carrying a large supply of party hats. She sees a bus on which Bill frequently rides home, so she jumps into some nearby bushes to avoid being spotted. Bill, watching from Chris's window, is quite amused, but Chris is puzzled and asks Bill why Ann is hiding in the bushes. Bill says:]

I might be on that bus.

Here the judge of *might be* is expected to be Ann, not Bill or Chris.

An explanation of how subjectivity works should be able to explain why predicates of personal taste and evaluative predicates could be exocentric, while epistemic modals could not. In Chapter 5, I return to the 'table scraps' example and conclude that exocentricity is a universal feature, though it is rarer among modals than it is among predicates of personal taste.

1.3.2 Disagreement

A second problem is the issue of how to define disagreement. How do we understand when two parties disagree? MacFarlane (2007) considers the following, but argues that it cannot be right; however, many linguists do implicitly or explicitly take it as their model of disagreement:

(21) Two parties disagree iff there is a proposition that one party accepts and the other rejects.

This poses a problem if we do not consider worlds and times in this proposition: John may assert *Tom is sitting* at 2 pm and Mary may assert *Tom is not sitting* at 3 pm, and although one utterance is the negation of the other, we do not get the sense that John and Mary disagree, as MacFarlane argues. MacFarlane accepts (with reservations, as pointed out by Lasersohn (2017)) the following modalization of this:

(22) Two parties disagree iff (i) There is a proposition that one party accepts and the other rejects, and (ii) the acceptance and the rejection cannot both be accurate.

A crucial point is whether two parties need to be aware of each other to disagree. Lasersohn (2017) believes not:

(23) As we have already seen, two people may disagree even if they are geographically separated and unaware of each other's existence: If John, in Urbana, Illinois, asserts that licorice is tasty, and Mary, a stranger he has never met, in Santa Cruz, California, asserts that licorice is not tasty, their assertions contradict each other, even though they are not engaged in conversation with each other and do not presume each other to share a conversational context. (Lasersohn (2017), p. 33)

Not all linguists accept this; to me, *I disagree with Aristotle on the nature of blood circulation* seems distinctly odd since I am separated from Aristotle in both space and time and he cannot be aware of my existence.[1] But my point is that defining disagreement is an issue that any theory of subjectivity has to address. (I provide my definition and criteria in the following chapter.)

1.3.3 Retraction

A third problem is the issue of Retraction, which particularly concerns MacFarlane (2014). He presents the following case, known as the Coffee Shop example, about epistemic modals:

(24) Suppose you are standing in a coffee line, and you overhear Sally and George discussing a mutual acquaintance, Joe:
SALLY: Joe might be in China. I didnt see him today.
GEORGE: No, he can't be in China. He doesn't have his visa yet.

1. Some readers accept this, but hesitate at *Aristotle disagrees with me on the nature of blood circulation.*

SALLY: Oh, really? Then I guess I was wrong.

SALLY: Oh, really? #Still, I was right when I said "Joe might be in China," and I stand by my claim. (MacFarlane (2014), p.240)

MacFarlanes's concern is why we consider it reasonable for Sally to retract her claim rather than stand by it in the case of epistemic modals, since at the time she made the claim, she was speaking to the best of her knowledge and she didn't assert that Joe *was* in China, only that he might be. It isn't clear, either, why George felt he could reject her claim that Joe might be in China, if she was talking about her own mental state, which George has no basis to reject.

MacFarlane notes that retraction is also a valid concern for predicates of personal taste: when a speaker's tastes change, he or she feels justified in retracting their earlier statement, as MacFarlane illustrates in his 'Fish Sticks example':

(25) When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I've changed my mind about the tastiness of fish sticks. So, if someone said, "But you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I didn't mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane (2014) pp. 13-14)

Why does the adult MacFarlane feel the need to retract a statement made by the child MacFarlane, when the child MacFarlane was judging his tastes based on an entirely different context and set of life experiences? Not many of the competing theories on the market do address the issue of Retraction, but it crucially influences MacFarlane to distinguish the

context of use from the context of assessment, as discussed in Section 1.4.

1.3.4 Kinds of subjectivity

A fourth problem is whether subjectivity is a homogeneous phenomenon: are all subjective predicates alike or is each class of subjective predicates subjective in its own way?

Kennedy (2013, 2016) argues that there must be two kinds of subjectivity: vague predicates such as *tall* allow faultless disagreement only in the positive form, while predicates of personal taste and evaluative predicates allow it in the comparative form as well, as shown in the examples below.

(26) a. Sue: John is tall/rich.
Zoe: No, he is not! (Faultless disagreement)

b. Sue: John is taller/richer than Bill.
Zoe: No, he is not! (Objective disagreement)

(27) a. Sue: Avocados are tasty.
Zoe: No, they're not! (Faultless disagreement)

b. Sue: Avocados are tastier than chocolate.
Zoe: No, chocolate is tastier than avocados. (Faultless disagreement)

(28) a. Sue: The Chicago Picasso is ugly/beautiful.
Zoe: No, it's not! (Faultless disagreement)

b. Sue: The Chicago Picasso is uglier/more beautiful than Calder's Flamingo.
Zoe: No, Calder's Flamingo is uglier/more beautiful than the Chicago Picasso. (Faultless disagreement)

Stephenson (2007a,b) does not say explicitly that there are different kinds of subjectivity, but she is forced to give a different analysis for epistemic modals than she does for predicates of personal taste, due to the aforementioned problem of exocentricity: the latter allow

exocentric readings, while the former, pace Egan *et al.* (2005), generally do not.

Bylinina (2013, 2016) considers the problem and concludes that although vague predicates like *tall* and *rich*, evaluative predicates like *beautiful* or *intelligent*, and predicates of personal taste like *tasty* and *fun* have different *sources* of their subjectivity (as I discuss in Section 1.4), she claims that there is no good argument that there are different *kinds* of subjectivity per se.

Whereas in the counterstance contingency model espoused by Kennedy & Willer (2016), as I discuss in Section 1.4, predicates that can be embedded under *find*, such as *fascinating* (*Kim finds Lee fascinating because he's an expert on oysters*) must have further properties than predicates that can just be embedded under *consider*, such as *vegetarian* (*Kim considers Lee vegetarian because the only animals he eats are oysters*). Thus, Kennedy and Willer do argue that there are two types of subjectivity, although the distinction falls along slightly different lines than the distinction of whether faultless disagreement in the comparative is available.

Having presented the problems that any definition of subjective predicates should account for, I now present some possible solutions. I classify them by in what aspect of the semantics (or the pragmatics) they consider the difference between subjective and objective predicates to lie.

1.4 Previous theories

Intuitively, subjective expressions are context-sensitive. Whether they are perceived as "true" varies depending on when, where and by whom they are used (and by the audience). Sue would judge *Avocados are delicious* as true, while Zoe would judge it false. Five-year-old Sue would claim *Playing on the see-saw is fun* while *Discussing philosophy with Ph.D students is not fun at all*; twenty years later, though Sue is the same person, she would assert the opposite — in a world where she becomes a semanticist. (She would still agree that *Discussing philosophy with Ph.D students is not fun at all* in a world where she

becomes a professional alligator wrestler.)

Semanticists' most useful tools for dealing with context-sensitive expressions were invented by David Kaplan (1989) to account for indexical expressions like 'I.' Let us use the following example:

(29) John: I am a doctor.
Mary: I am a doctor.

John and Mary's assertions "mean the same thing" in that they both claim to be doctors, and in fact use the same words to say so. However they "mean different things" in that John's statement would be true iff John was a doctor at the relevant world and time, while Mary's statement would be true iff Mary was a doctor at the relevant world and time. Kaplan explained that such pairs of expressions have the same *character* but different *content*.

The Kaplanian understanding of *context* is the set of indices — time, world, agent, location — that need to be resolved in order to give a proposition an appropriate truth value. Content is the function from context to truth value for a given proposition. Context-sensitive propositions vary in truth value when the context varies; e.g. *I am a doctor* is true when uttered by Dr. Atul Gawande (the index of *I* resolves to [[the individual called Dr. Atul Gawande]]) and false when uttered by Queen Elizabeth II (the index of *I* resolves to [[the individual called Queen Elizabeth Windsor of the United Kingdom]]).

Every indexical expression has a single character that is fixed, but will vary in content between contexts. The varied contents would pick out different extensions. The extension of a term is the set of all objects that meet the criteria for that term; the extension of a proposition is its truth value.

The question then becomes how subjective expressions differ from objective ones — in the content, in the extension, in the context or somewhere else entirely? And as we will see, different models of subjectivity approach this puzzle differently. In the following subsections, I will present the models based on the criteria of whether they consider the difference to be

nonexistent (objectivism), not in the semantics (expressionism), in the content (contextualism), in the context and extension (relativism), or in the metalinguistics (metalinguistic negotiation and counterstances).

1.4.1 Realism/Objectivism: there is no difference

As Kölbel (2002) interprets it, unmitigated realism denies that subjective propositions exist and draws no distinction between them and objective propositions. One of Sue or Zoe is (or both are) simply wrong when they disagree. Kölbel claims, and I agree, that this position is unattractive because intuitively, people *do* have a distinction that some disagreements are faultless and some are not.

Mitigated realists believe that even though one of the disputants in such a disagreement is at fault, and holds a false belief, neither of them have made a mistake that they should have avoided. This poses a cognitive problem: how could flawless reasoning have arrived at a false belief?

This led Wright (1994) to propose *intuitionist realism*, a term coined by Kölbel. Intuitionist realism can acknowledge that neither disputant is making a mistake while remaining silent on one of them having a false belief. This poses even more problems than the position of the unmitigated realist. Kölbel dismisses intuitionist realism with the words, "Faced with the unmitigated realist's affirmation that one of them is mistaken, the best [the intuitionist realist] can do is remain silent. The intuitionist is not able to deny the claim that one of them is mistaken. Worse: intuitionist realism does not even allow us to state that there is some sense in which the disputants in question are blameless. I conclude that intuitionist realism is in a worse position than mitigated realism to account for our intuition that some disagreements are faultless." (Kölbel (2002), p. 62)

So we see that assuming that there *isn't* a difference between subjective and objective predicates leaves us with problems. We can take another assumption, though: that the difference is that subjective expressions are not actually semantic — they do not carry

logical meaning or factual language. This is the position of expressivism.

1.4.2 Expressivism: there is too much of a difference, as subjective expressions are about emotion, not truth

The early noncognitivists, such as Ayer (1936), Stevenson (1937) and Hare (1952) argued that there is a distinction between factual language ("The cat is on the mat"; "Stealing is illegal") and moral language ("Stealing is wrong"). Ayer asserts that there is no further logical distinction between "You were wrong in stealing that money" and "You stole that money"; there is only an additional shade of moral disapproval. Thus expressions like *wrong* are merely emotive expressions that do not actually have propositional content. Kölbel (2002) distinguishes Ayer-style expressivism, which argues that any real disputes about questions of value are actually non-faultless disputes about related factual questions. There is no room for faultless disagreement; "it is impossible to dispute about questions of value," as Ayer asserted in 1946.

To the expressivist, stating a subjective or moral proposition is a statement of attitude. To simplify the position, *Stealing is wrong* means approximately the same as "Boo, stealing!" or "I disapprove of stealing." *Avocados are tasty* means approximately the same thing as an expression of pleasure: "Yum, avocados!"

For some expressivists, a subjective proposition is a non-assertive kind of speech act, in the distinction that Austin (1962) and Searle (1962) propounded. Where imperatives and directives such as "Go to bed!" issue a command and commissives such as "I promise you ten dollars" create promises, subjective propositions express feelings.

This runs quite quickly into the problem: if we are not actually disputing, why do we feel like we are? Why does (a) sound natural and (b) sound wrong?

(30) a. A: Avocados are tasty.

B: No, that's not true. They're not tasty at all.

b. A: Yum, avocados!

B: #No, that's not true. Not yum, avocados, at all.

It does not account for why "Stealing is wrong" and "Stealing is not wrong" are not truth-compatible. "#Not yum, avocados" is a nonsensical statement, while "Avocados are not tasty" is a perfectly comprehensible one that we understand contradicts the statement "Avocados are tasty."

The problem that most expressivist models run into is widely known as the Frege-Geach problem, after Geach (1965) who developed it from the writings of Frege (1956). Essentially, we can have subjective propositions compositionally enter into logical relations like modus ponens, which expressions of pleasure or distress cannot do, as the following examples demonstrate:

(31) Stealing is wrong.

If stealing is wrong, then getting your little brother to steal is wrong.

Therefore, getting your little brother to steal is wrong.

Trying to analogize this to expressions of attitude does not make sense:

(32) Boo, stealing.

??If boo, stealing, then boo, getting your little brother to steal...

If we analyze subjective propositions as a type of speech act, we run into the problem that subjective propositions can also be embedded under negation or attitude verbs and still remain subjective, which conventional speech acts cannot do while remaining the same speech act type. This is illustrated by the following example comparing the commissive speech act *promise* (which is a promise, inducing an obligation on the part of the speaker) with a subjective proposition like *this is tasty*:

(33) a. I promise you ten dollars, — Commissive speech act

b. I don't promise you ten dollars. — Not a commissive speech act

c. I think I promised you ten dollars. — Not a commissive speech act

d. This is tasty. — Subjective proposition/"expressive"

e. This is not tasty — Still a subjective proposition

f. I think this is tasty – Partly a subjective proposition.

To enter into logical relations like modus ponens, a statement should have propositional content and a truth value. It certainly *seems* like we can do that with subjective expressions containing moral terms, evaluative terms, and predicates of personal taste, whereas we cannot do this with canonical emotive statements like cries of pain or pleasure.

Essentially, the emotivist / noncognitivist / expressivist models predict that moral terms like *good* and *wrong* (and other subjective terms like *tasty*, presumably), as statements of the speaker's attitude, must have different functions and force than descriptive terms like *green* or *wooden*. (Schroeder (2008a) has a good explanation) But there is no difference in how moral terms and descriptive terms enter semantic composition: moral and descriptive terms alike can be negated and we do have the distinct sense that *not-p* contradicts p whether p includes a moral or a descriptive statement. Moral and descriptive terms alike can enter into logical relations like modus ponens.

Blackburn-style expressivism does try to bypass the Frege-Geach problem by arguing that even though they are not truth-evaluable, expressive and subjective statements can enter into a kind of truth relation. Gibbard (1992) sets up a system for expressivism related to norms. However, this still does not address the issue completely: how do we tell a descriptive statement from a normative one when they fit the same syntactic patterns and patterns of entailment?

The noncognitivists and the expressivists take the distinction between moral language and descriptive language for granted. One expresses emotions, or norms, while the other

expresses how the world is — but how can we tell the two apart? The Frege-Geach problem, as Schroeder (2008b) explains, addresses the very fact that descriptive and moral terms look pretty much alike under syntactic and semantic operations.

More recent theories of expressivism, such as Gutzmann (2016) and Silk (2016), argue that there can be expressivist theories that account for the Frege-Geach problem. Gutzmann (2016) argues that treating subjective propositions as speech acts is not as big a problem as Lasersohn (2005) considers it when he rejects that idea. Lasersohn rejects the notion that subjective propositions are anything but truth-conditional assertions because, he argues, they can be denied, can be embedded in modus ponens constructions, and give an intuition of contradiction, which requires truth conditions. Gutzmann counters that other kinds of speech acts besides assertions can also be negated and give an intuition of contradiction, and enter modus ponens constructions, as shown in the following examples.

(34) a. Subjective propositions can be denied:

Sue: That's tasty!

Zoe: That's not true! That's not tasty at all!

b. Lasersohn's argument that true expressives cannot be denied:

Sue: Mm-mm!

Zoe: ??No!/That's not true! That's not tasty at all!

c. Gutzmann's counterargument that some expressive and commissive speech acts (congratulations and promises) can be denied:

(i) A: I congratulate you!

B: That's not true! I know you are actually really envious of me.

(ii) A: I promise you I will come to your party.

B: That's not true! I know you are not planning to come as you won't even be in town then.

(35) a. Subjective propositions can be embedded under modus ponens:

Premise: If there is a loop, the roller coaster is fun.

Premise: There is a loop.

Conclusion: The roller coaster is fun.

b. Gutzmann's argument that so can non-assertoric speech acts, such as directives, commissives and expressives:

(i) Premise: If the roller coaster has a loop, go for it.

Premise: The roller coaster has a loop.

Conclusion: Therefore, go for it.

(ii) Premise: If the roller coaster has a loop, I promise to ride it with you.

Premise: The roller coaster has a loop.

Conclusion: Therefore, I promise to ride it with you.

(iii) Premise: If the roller coaster has a loop, I congratulate you for being brave.

Premise: The roller coaster has a loop.

Conclusion: Therefore, I congratulate you for being brave.

(36) a. Subjective propositions can be contradicted:

Sue: That's tasty!

Zoe: That's not tasty at all!

b. Gutzmann's argument that so can non-assertoric speech acts, such as directives and expressives:

(i) Tom to Suzy: Take a ride on this roller coaster!

Hanna to Suzy: Do not take a ride on this roller coaster!

(ii) Tom (eating tofu): Mm-mm!

Hanna (eating from the same tofu): Ugh!

Lasersohn (2005) argues that contradiction means that two sentences cannot be accommodated under a single perspective, for example:

(37) Sue: #That's tasty and it's not tasty.

Gutzmann argues that other speech acts also have contradictory forms that cannot be accommodated under a single perspective, such as expressives, directives and commissives, which all give the same sense of self-contradiction as the example with *tasty*:

(38) a. Tom (eating tofu): #Mm-mm! Ugh!

b. Tom: #Take the roller coaster! Do not take the roller coaster!

c. Tom: #I promise that you'll get a bike for Christmas. I promise that you won't get a bike for Christmas.

We will see that both metalinguistic negotiation (of which more below) and pragmatic theories of subjectivity related to common ground update (of which much more in the following chapter) are quite close to forms of expressivism, as they depend on pragmatics rather than semantics. I devote a fair amount of space to explaining modern expressivism here because in Chapters 2 and 5 I will argue for a model that is very close to it.

Viewed through our lens, expressivists have the opposite problem that objectivists do: instead of taking there to be no difference between subjective and objective expressions, they take there to be too much of a difference, as subjective expressions should not be able to enter into logical compositional constructions at all. Yet there is endless evidence that they do, and modern expressivists like Gutzmann (2016) skirt around this issue.

So we can look for the difference in the relation between referent and indexical, and that is what the camp known in semantics circles as the *relativists* does.

1.4.3 Indexical relativism and assessment sensitivity: the difference is in the indexical

A prominent model in recent years has been the so-called judge-index relativist model, often just called relativism. Although it had been previously discussed in philosophy circles as far

back as Lewis (1970), as well as by Richard (2004) and others (it is also addressed by Kölbel (2002)), it came to the attention of semanticists with Lasersohn (2005) in *Linguistics and Philosophy*, where Lasersohn applied the tools of modern semantics to the problem of how to model subjectivity. Most of the semantics literature on subjectivity since 2005 has been in conversation with Lasersohn, so his model will be the one I examine most closely.

Semantics already uses indices to show pronoun reference, time, and world. Lasersohn's innovation was to suggest that subjective predicates like *tasty* and *fun* have an index for the judge as well. That is,

(39) a. $[[tasty]]^{c;w,t,j} = \lambda x.x$ tastes good at time t, in world w and to judge j

b. $[[\text{Mint is tasty }]]^{c;w,t,j} = 1$ if mint is tasty to j in w at time t.

Thus, the interpretation of *tasty* is relativized to a particular judge. The *content* of *tasty* is the same no matter who says it, similar to how *me* always has the same content referring to the speaker. However, the context and extension resolve the index — we learn from the context which particular judge is in question.

It is perhaps confusing that it is relativism that sees the difference between subjective and objective predicates to be in the context, rather than contextualism (of which more below).

This proposal has been popular with several other semanticists and philosophers. Stephenson (2007a) incorporates it into a model of common ground update based on Stalnaker (1978). MacFarlane (2014) expands on the role of context in resolving a judge index by introducing both a *context of utterance* and a *context of assessment*.

MacFarlane is interested in why we would consider retracting subjective statements like his 'Fish Sticks Example' I presented in Section 1.3 and repeat here:

(40) When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; Ive

> changed my mind about the tastiness of fish sticks. So, if someone said, "But you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I didn't mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane (2014) pp. 13-14)

MacFarlane seeks to explain the feeling we have that we need to retract previous taste assertions if our tastes change. He posits that the truth of a statement can be ascertained both based on indexicals of context of utterance and of context of assessment. At the context of utterance, ten-year-old MacFarlane saturates the judge indexical, and his statement "Fish sticks are tasty" should be true. However, at the context of assessment, the MacFarlane of 2014 is the judge, and according to his tastes, the statement "Fish sticks are tasty" is false, and he should retract it.

In later work, Lasersohn (2017) adopts MacFarlane's view for its ease of handling changing tastes, retractions and withdrawals.

Relativism has the advantage of explaining disagreement better than differences in content, as Lasersohn and MacFarlane explain it: the reason that Sue and Zoe contradict each other about avocados or other foods being tasty is that they have different indexicals, but the content and extension of *tasty* is still the same, as Lasersohn puts it. As MacFarlane puts it, Zoe's context of assessment is different from Sue's context of assessment.

Relativism builds on a semantic framework that has already worked well for indices of world and time, and allows the same machinery to handle subjectivity and modality (which also allows for faultless disagreement in such cases as *If X, we could have won the game. — No, even if X, we couldn't have.*). The flexibility of filling in a judge index also allows for perspectives to not be anchored to the speaker, such as in Stephenson's cat example, already

discussed above:

(41) Sue: How's that new cat food you bought?

Bob: It must be tasty, because the cat has eaten almost all of it.

What do Lasersohn and MacFarlane have to say about how to differentiate expressions that call for a judge from those that do not? Here are Lasersohn's (2017) thoughts on the subject:

> If we assume that for each use of an expression, there is a matter of fact about what the context of that use is, and that each context supplies a judge, then for each use there will be a matter of fact about who the judge is. In such a system, a sentence as used in context c could be regarded as true *simpliciter* iff it is true relative to the judge (and time and world) of c. But the point of introducing the judge parameter was to draw a semantic distinction between sentences whose truth values depend purely on matters of fact and sentences whose truth values depend on matters of taste; only the former should have truth values *simpliciter*, while the latter should have truth values only relative to the judge parameter.
>
> Therefore, I suggested in Lasersohn (2005), we should give up the assumption that there is a matter of fact about what the context is, in which a given use of an expression occurs. Rather, for any concrete situation in which a sentence is used, there are as many different "contexts" as there are individuals who might judge whether or not that use was truthful. The formal notion of context as represented in the grammatical theory was thus distinguished from the intuitive idea of the practical environment in which an expression is used; the practical environment underdetermined the formal context. Contextual parameters other than the judge were assumed to be fixed by matters of fact, of course; so the connection between the context and matters of fact about the practical environment was not entirely severed. Sentences that were purely about matters of fact could be distinguished

> from sentences about matters of taste in that their contents did not vary in truth value among contexts which differed only in the value of the judge parameter. (p.63)

But Lasersohn recognizes that this system would not work, because

> If each context supplies a judge, then each use will have (in any world where it occurs) a particular judge specified as the judge of its context in that world, and we will have returned to a system in which each sentence content can be assigned a definite truth value, regardless of whether it is about matters of fact or matters of taste...I think a better response would be to recognize that "judging" whether a use of a sentence is true is a separate matter from using the sentence; and that such judgment may be based in part on contexts which are different from the contexts in which the sentence is used. That is, we may articulate each formal context in the sense of Lasersohn (2005) into two parts, corresponding to the situation in which an expression is used, and the situation in which a truth value is "judged." (ibid.)

Thus Lasersohn endorses the variant on judge-index relativism proposed by MacFarlane (2007, 2014).

However, the choice of which predicates get the judge index seems quite stipulative in this model, whereas we see in ordinary English that some predicates can change from subjective to objective with context: *a lighter cake* may mean either one that sits more easily in the stomach (subjective) or one that weighs less than another (objective); *good mustard* could mean either mustard that tastes delicious (subjective), or mustard that has not yet spoiled (objective). If we introduce judges, how do we get real objective predicates back? How do we get back to *truth*? The idea of truth being relative is problematic to many, as van Wijnbergen-Huitink (2016) asserts in her introduction to a volume she edited with Cecile Meier (Meier & van Wijnbergen-Huitink, 2016) which is based around alternatives to relativism.

Yet another problem is the issue of disagreement itself. If Sue and Bob's notions of *tasty* are evaluated from different contexts, and Sue and Bob *know* they evaluate them from different contexts, why would they disagree about this? Hirvonen (2014) tries to get around this by positing 'folk objectivism': speakers use different perspectives to evaluate taste (she rejects both relativism and contextualism as formal theories) but do not realize that taste properties are not objective, thus motivating them to disagree as if those were objective properties, like Touchstone treats the knight's assertion about pancakes and mustard as an assertion on whose objective truth value one may place a publicly-announced stake. Although much of Hirvonen's analysis is worthwhile, the idea that speakers are naive about their language use strains our sense of psychological realism. It may have some grounds, and in fact I argue for what properties of taste predicates lead speakers to differ in their judgements about them; however, she does not explain in depth why this should be so.

For another problem, our most common pre-existing model for semantics does not accommodate introducing a new index too well. As Pearson (2012) summarizes the issue, "Such a move complicates the semantics by requiring that each predicate takes an additional individual argument as well as a world (and time) argument, with implications across the inventory of semantic types: the intension of a one-place predicate like *table* or *run* is now of type $< s, < e, < e, t >>>$, that of a two-place predicate like *love* is of type $< s, < e, < e, < e, t >>>>$ and so on." (Pearson (2012), p. 2.)

The assumption that the English phrase *to/for X*, i.e. *tasty to me* or *fun for me* is an explicit pronunciation of the judge index in English poses yet more problems for Lasersohn's account. The patterns of distribution of the *to X* phrase are idiosyncratic. Many speakers, like Stephenson (2007a), do not accept *tasty to me*. Even though *tasty* and *delicious* are otherwise synonyms, *?delicious to me* seems a lot worse than *tasty to me* even to speakers like me who do accept *tasty to me*. Additionally, many predicates that are indisputably subjective, like *smart* or *lazy*, do not accept *to me*: *?smart to me* or *??lazy to me* are out for speakers. Overall, although this phrase seems like an appealing indicator, a closer look

shows that as an indicator of the judge index, it almost creates more problems than it solves.

The problems with relativism as an analysis of predicates of personal taste have led more semanticists and philosophers to argue for its main opposition, contextualism, which I describe in the next subsection.

1.4.4 Contextualism: the difference is in the content

Looking at subjective predicates within a Kaplanian framework, *contextualism* argues in favor of the actual content of the predicate differing depending on who says it. A contextualist semantics for subjective predicates includes an implicit argument showing the judge (similar to *pro*). The content of subjective predicates, rather than the context indexicals that form its functional domain, is what makes their truth values differ. Thus *tasty* when Sue utters it is actually a different *tasty* than when Bob does; Sue is looking for different things when she tastes her food, even if her world, time and location are the same.

(42) a. $[[tasty]]^{c;w,t,j} = \lambda x.\lambda y.y$ tastes good to x at time t, in world w.

b. $[[\text{Mint is tasty } pro]]^{c;w,t,j} = 1$ if mint is tasty to pro_j in w at time t.

Intuitively, this seems reasonable: if we try picking apart our taste in, say, movies, we realize that one person's notion of a *good movie* may have completely different requirements than another's (witty dialogue versus well-staged action scenes, etc.), suggesting that $good\,movie_{Sue}$ should have different semantic content from $good\,movie_{Bob}$. Stojanovic (2007) has even argued that contextualism has all the virtues of relativism, and that anything that relativism can do, contextualism can also do; she argues that the difference between indices and arguments is merely a notational variant, siding in favor of contextualism and implicit arguments for that reason.

Stephenson (2007a,b) suggests a model that is intermediate between contextualism and relativism. She proposes that predicates of personal taste take a silent nominal argument PRO_j that refers to the judge, and assumes that prepositional phrases like *to me* or *for the*

cat occupy the role of PRO_j, thus saturating the judge argument. On the other hand, she has an indexical analysis for epistemic modals, very similar to that of Lasersohn (2005). "The difference between epistemic modals and predicates of personal taste, then," Stephenson says (p. 500), "is that epistemic modals are inherently judge-dependent, whereas predicates of personal taste become judge-dependent only if they take PRO_j as an argument." In essence, Stephenson wants to explain why *tasty to me* is fine, whereas *?Rain is likely to me* is very odd, and also why exocentric use of modals is very odd — although it is easy to say something like the following:

(43) The cat food must be tasty, because Goober ate all of it. (tasty = tasty to Goober)

Whereas it is very hard to find an exocentric use of an epistemic modal, although as we saw in the introduction, Egan *et al.* (2005) provide a possible case for one, which I repeat below:

(44) [Context: Ann is planning a surprise party for Bill. Unfortunately, Chris has discovered the surprise and told Bill all about it. Now Bill and Chris are having fun watching Ann try to set up the party without being discovered. Currently Ann is walking past Chris's apartment carrying a large supply of party hats. She sees a bus on which Bill frequently rides home, so she jumps into some nearby bushes to avoid being spotted. Bill, watching from Chris's window, is quite amused, but Chris is puzzled and asks Bill why Ann is hiding in the bushes. Bill says:]
I might be on that bus.

Stephenson thus argues that predicates of personal taste and epistemic modals need to be analyzed differently, with predicates of personal taste requiring what is basically a contextualist analysis, whereas epistemic modals need a relativist one.

Stephenson's account, though, comes under criticism by Lasersohn and MacFarlane, although Lasersohn (2017) acknowledges the importance of Stalnakerian common-ground up-

date that Stephenson brings forward in her paper and dissertation. They would much prefer a uniform account of subjective predicates across epistemic modals and predicates of personal taste alike.

Contextualism has one big problem to its critics: There is no accounting for the faultlessness of disagreement. If the content of the same sentence is different for different speakers, i.e. if Sue says *Mint is tasty (to Sue)* and Bob says *Mint is not tasty (to Bob)*, they should not be in contradiction. It seems absurd for a contradiction to go as follows:

(45) Sue: Mint tastes good to me.
Bob: #No, it doesn't taste good to me!

This parallels the equally absurd example:

(46) Sue: I'm a doctor.
Bob: #No, I'm not a doctor!

But in the standard scenario, we have the strong sense Sue and Bob are indeed contradicting each other, which is a problem. Stojanovic (2007) deals with this by dispensing with faultlessness; she claims that any scenarios of faultless disagreement are either genuine disagreement, or a misunderstanding: that is, Bob didn't know that when Sue said *mint is tasty* she was referring only to her own taste rather than taste *tout court* (in which case she would be wrong).

However, there still remains the problem that if such conversations are not disagreements, they certainly feel that way. Following Stojanovic's view would make misunderstanding incredibly prevalent. Furthermore, how does Sue feel that she is justified in asserting that mint is tasty if she is aware that some people may disagree?

A second major problem is that if there is a hidden argument that reflects the judge for subjective predicates, there should be syntactic evidence for that argument also (the

way *pro* is revealed for through verb inflection in pro-drop languages and *PRO* is revealed through subject and object control tests). But if so, then Pedersen's example of *win the debate*, presented in the introduction, poses a problem, as we have two syntactically-identical sentences where one is subjective and one is not. There is no test I know of that would reveal a PRO_j in *Obama won the debate* and not in *Sharon won the debate*.

This motivates an alternative approach: that there is indeed something in the content of subjective predicates that causes them to be subjective, but it is not explicitly an argument reflecting the judge.

1.4.5 Judge-free contextualism: the difference is in genericity, multidimensionality, the Experiencer argument, or vagueness

The problem of where to locate the judge in the semantics of subjective predicates — as an index or as an argument? In the content or in the context and extension? — has led some scholars to suggest that the judge can be dispensed with entirely. In this subsection I will look at models that try to differentiate subjective from objective predicates in another part of the semantics.

Genericity and identify-with relations

Moltmann (2010) proposes that there is no concrete judge, as an index or an argument, in an assertion like *Mint is tasty*; instead, it contains a first-person generic that extends from the speaker to generic people. She argues this with data such as the following:

(47) Premise: John believes the wine is tasty (and is a wine connoisseur).
Premise: Mary believes the wine is tasty (and is a wine novice).
Conclusion: John and Mary believe the same thing/have the same belief.

This conclusion is valid and a reasonable statement to make, Moltmann argues, even though

the contexts of Mary's and John's beliefs are very different. On the other hand, beliefs do not have this property:

(48) Premise: John believes he is the winner.
Premise: Bill believes he is the winner.
Conclusion 1: John and Bill believe the same thing.
Conclusion 2: John and Bill believe different things.

Both these conclusions, Moltmann argues, are valid, whereas this cannot be said of the conclusion that John and Mary believe different things about the wine. She uses this data to argue that the difference is due to subjective predicates like *tasty* having a generic component that full truth-conditional beliefs do not have. This generic component is a difference in content; thus, Moltmann argues for a form of contextualism. The genericity is very similar, she argues, to generic *one*. Like subjective propositions, generic *one* used by a speaker implies that the speaker herself has the relevant ability:

(49) Premise: One can see the picture from the entrance.
Conclusion: I can see the picture from the entrance.

(50) Premise: John believes that one can see the picture from the entrance.
Conclusion: John believes that he/John can see the picture from the entrance, based on personal experience.

Compare this example to:

(51) Premise: John: the wine is tasty.
Conclusion: John has tasted the wine and it is tasty to him.

In both cases, the use of *one* or a subjective predicate anchors the perspective to the speaker. Moltmann also notes that such assumptions of joint belief as (47) are also applicable to uses

of *one*, as in the following example:

(52) Premise: John believes one can sleep on this sofa.
Premise: Mary believes one can sleep on this sofa.
Conclusion: John and Mary believe that one can sleep on the sofa.

In both the case of *tasty* and the case of *one*, it is required that *John and Mary believe the same thing* follows from each of them holding the belief. This leads Moltmann to conclude that both *one* and subjective predicates contain first-person-oriented generic content. An indexical cannot capture this, Moltmann argues, because it does not reflect our intuitions on disagreement as well as a first-person generic does. She defends her argument that faultless disagreement is best explained by analogizing to generic-*one* sentences as follows:

(53) The intuition of faultlessness of disagreement about a generic-*one* sentence arises because the sentence requires the predicate to be applied in a first personal way, 'as if to oneself', to every individual in the domain. Whether the content of a generic-*one* sentence is accepted, rejected, just entertained, or merely understood, this requires the same first-personal access for any agent, that is, an application of the predicate to the individuals in the domain as if to oneself. Formally, the condition that the content of a generic-*one* sentence can be accessed only in a first-personal way consists in a self-attribution of the property expressed by the generic-*one* sentence. (Moltmann (2010), p. 208).

This is an intriguing model, but is rightly criticized for not accounting for statements like the cat food example, where the speaker has no first-person experience with the cat food:

(54) The cat food must be tasty, because Goober ate all of it. (tasty = tasty to Goober)

Pearson (2012, 2013) further refines the model of subjective predicates having a first-person

component by analyzing a subjective predicate as referring to the speaker and those the speaker empathizes with (including, as necessary, cats). Pearson posits an *identify with* relation in the first-person generic content of a subjective predicate; this fixes the problems in Moltmann's account to some degree, as it accounts for exocentricity. The strengths of Pearson's proposal are that it requires very little adjustment to pre-existing Heim and Kratzer-style semantic types.

Although this is an interesting premise, it still does not explain where the first-person generic would come from, or why it applies to subjective predicates but not objective ones.

Vagueness, multidimensionality, and Experiencers

Bylinina (2013) also argues that the 'judge' is not a uniform class — differences in perspective lie either in the Experiencer thematic role of a predicate like *tasty*, or varying standards for a predicate like *tall*, or a predicate like *smart* being multidimensional (composed of several different component predicates, like smart at math, smart at languages, smart at social skills, etc.) and different speakers prioritizing those dimensions differently.

Bylinina (2013) points out that there are three classes of adjectives that show different subjectivity patterns:

(55) a. Predicates of Personal Taste (PPTs), e.g. *tasty, fun*, allow faultless disagreement both in the positive and in the comparative form, and take 'judge PPs.'

b. Evaluative Adjectives (Bierwisch, 1989), e.g. *smart, lazy*, allow faultless disagreement both in the positive and in the comparative form, but do not take 'judge PPs.'

c. POS-Dimensional Adjectives, e.g. *tall, deep*, allow faultless disagreement in the positive form only, and do not take 'judge PPs.'

She concludes that the sources of subjectivity in all three of these classes are different.

POS-dimensional adjectives are subjective due to the subjectivity of the silent POS mor-

pheme indicating the degree standard ((Kennedy 1999), (Kennedy & McNally, 2005) a.o.). Thus, the subjectivity lies in two people disagreeing what the standard should be, e.g. how many centimetres of height are required to qualify as *tall*? This is induced by the well-known property of *vagueness*.

PPTs, Bylinina argues, are subjective due to them having Experiencer semantics. As Nouwen (2007) proposed, the reason that we find *tasty* to be subjective is that the speaker has privileged information — only she is the experiencer of pleasure at the taste of the item in question. Using a cross-linguistic analysis with data from Japanese, Russian, and Hungarian, Bylinina shows that for predicates involving an Experiencer argument, the 'judge PP' is the overt manifestation of the Experiencer thematic argument, with a requirement that the 'judge' be the Experiencer. In Chapter 4, I discuss this hypothesis further and point out some issues with the difference between the thematic role of Experiencer and the requirement that the subject had the required experience.

Evaluative adjectives, on the other hand, are subjective for a different reason according to Bylinina. They are multidimensional, in that they aggregate several different predicates in themselves; for example *smart* may encompass skill at mathematics, language, wilderness survival, etc. When they were first described, by Kamp (1975) and Klein (1980), they were referred to as *nonlinear adjectives* — nonlinear for the very reason that they allow faultless disagreement in the comparative form, because different speakers may prioritize different dimensions. More recently, Sassoon (2012) has extensively studied the properties of multidimensional adjectives, while van Rooij (2010), whom Bylinina follows, provides a model for how subjectivity may arise from their multidimensionality.

Subjectivity arises from a specific subclass of multidimensional adjectives. Sassoon (2012) distinguishes the following subclasses:

(56) a. *Conjunctive*: Entities are required to reach the standard in *all* of the component dimensions. Exception phrases are permitted:

(i) John is healthy (in all respects) except for his blood pressure.

These pattern in the same way as upper-closed adjectives such as *closed, straight* (Kennedy & McNally, 2005; Rotstein & Winter, 2004): they allow modification with *completely* but not *slightly*, for example.

b. *Disjunctive*: Entities are required to reach the standard in *any one* of their component dimensions. Exception phrases are not permitted:

(i) John is sick #except for his blood pressure.

These pattern in the same way as lower-closed adjectives, such as *open, bent* allowing modification with *slightly* but not *completely*.

c. *Mixed*: Context determines how many dimension standards the entity must exceed in order to count; thus exception phrases may or may not be acceptable. *Smart* is an example. Neither *completely* not *slightly* are acceptable; thus, they pattern with relative-standard adjectives such as *tall, deep*.

Van Rooij points out that in a ('mixed') multidimensional predicate, some dimensions may 'count' for more than others in certain contexts. For example, taking the classic multidimensional adjective *healthy* (comprising the different predicates of, say, blood pressure, cholesterol level, pain, presence of cancer, etc.), we observe that 'having a mild cough' may still qualify someone as *healthy* for coming in to office work or school — but not for going into space or performing surgery. In certain contexts, 'mathematical proficiency' would count for a great deal towards being counted as *smart*, while 'Classical Greek proficiency' might be mildly impressive, but not counterbalance a shortage of math skills. He thus proposes a 'weigh function' that a speaker may apply to the component dimensions of a multidimensional adjective in a given context, in the form of an ordered vector: if a dimension should not count for that much, its weight would be low, and if it should not count in the context, its weight would be zero.

Thus, Bylinina argues, evaluative adjectives are precisely the 'mixed' multidimensional

adjective class, and subjectivity in evaluative adjectives can be explained merely by different speakers (in different contexts) having different weightings.

(57) Yuri: John is smarter than Mary: he is a mathematical genius.
Zelda: No, Mary is smarter than John: she is an expert wilderness guide, and John won't last a second in the wilderness.

In this case, Yuri is weighting *good at math* heavily in his assessment of *smart*, while Zelda weights *good at wilderness survival skills* much heavier than mathematical proficiency in her own assessment.

Bylinina's analysis is, I think, a great step forward in studies of subjective phenomena. However, the problem is that it overgenerates. If multidimensionality were the sole basis for subjectivity in evaluative adjectives, then we would expect that such well-known multidimensional predicates as *big* (van Rooij, 2010) would also be subjective. And most of the time, they aren't (although perhaps this is because the dimensions for *big* are mostly determinate as width, height, length — but that shows that it is not the multidimensionality per se but its indeterminacy that becomes the key issue). On the other hand, *healthy*, being an example of a conjunctive multidimensional adjective, should never be subjective, yet some of the time, it is. The following example seems like a faultless disagreement:

(58) A: Chantal Petitclerc, Paralympic wheelchair-racing champion, is healthier than I am.
B: No, she isn't. You can walk and she can't.

Since the number of contexts and the number of speakers are endlessly varied, and, as Kamp (1975) observed, multidimensional adjectives form the majority of adjectives, how do we understand each other at all? Why doesn't every dialogue collapse in faultless disagreement?

Grinsell (2012) has a very interesting analysis of why multidimensional adjectives such

as *healthy* lead to vagueness effects; he brings in economics, specifically social choice theory, as the reason. Vagueness in multidimensional adjectives, he argues, is the result of multi-criterial decision making: speakers are forced to evaluate multiple criteria, the same way voters in a political election have to evaluate a number of different candidates who have differing platforms concerning foreign policy, health care, policing, etc. There is an economic result, known as *Arrow's Theorem* after economist Kenneth Arrow (1951):

(59) Arrow's Theorem is easy to state: there is no collective decision procedure that respects certain rational assumptions and avoids circularity, or cycling. (Grinsell, 2012)

That is, assume that there is are candidates Alice, Bob, and Chris. On criterion X, Alice is better than Bob, who is better than Chris. On criterion Y, Bob is better than Chris, who is better than Alice. On criterion Z, Chris is better than Alice, who is better than Bob. Thus, how do we choose the optimal candidate for all three criteria? We end up in a circle, with the logical reasoning to the absurd conclusion that if Alice is better than Bob, and Bob is better than Chris, and Chris is better than Alice, then Alice is better than herself.

Grinsell's proposal is that this is what happens to create vagueness in such multidimensional adjectives as *healthy*. He uses vagueness in the sense of Kennedy (2007), as involving three criteria:

(60) a. Vague predicates are susceptible to the sorites paradox, for example:
   (i) A person 2 m tall is tall.
   (ii) A person 1 mm shorter than a tall person is tall.
   (iii) A person 1 mm tall is therefore tall.

   b. Vague predicates are susceptible to borderline cases: even if we know someone is 180 cm in height, we may be unsure as to whether that counts as *tall*.

   c. Vague predicates have contextual variability: to use Kennedy's (2007) famous

example, the statement *The coffee in Rome is expensive* is true if the comparison class is other Italian cities, but false when comparing Rome to Chicago.

Grinsell argues that these vagueness effects in multidimensional predicates arise because of the issues of cycling: because speakers are trying to aggregate multiple choice functions, they find that their evaluation varies by context (including by speaker), that there are borderline cases (should Candidate Bob above be considered *good* if he is better than Chris, but worse than Alice?), and that we can form a sorites series:

(61) (Grinsell, 2012): A person with low blood pressure, low cholesterol, and a low heart rate is healthy. A person with high blood pressure, high cholesterol, and a high heart rate is not healthy. Is a person with low blood pressure, low cholesterol, and a high heart rate healthy? A person with mid-values for all the measures?

(62) A person whose systolic blood pressure is one unit higher than a healthy person's is healthy.

We recognize that our issue of subjective predicates also has the two features of strong context-sensitivity, and borderline cases as to whether someone should count as *smart* or a meal count as tasty. Grinsell (p.c.) admits that it is difficult to form a sorites series with subjective predicates. However, in cases of evaluative predicates that contexts may assign a numerical value, I claim one may indeed construct a sorites series:

(63) a. Many wine critics rate the wines they review out of 100. So a wine that is rated a 98, this critic definitely found very tasty. A wine that is 97, likewise...A wine that is a 10/100 is definitely undrinkable pig slop. But by the sorites series, where is the line actually drawn between wines fit for the Queen, and pig slop?

b. A figure skater who scores 295 points (the current world record under the modern ISU judging system) is obviously a good figure skater. One who scores 294

points, likewise...so does it follow that a skater who scored zero points is a good figure skater?[2]

The vagueness effects are due, in essence, to the aggregation function being unclear. (Unlike van Rooij (2010) and Bylinina (2013), Grinsell assumes for his analysis that all of the dimensions are equally weighted.) Thus, one way to break the cycling and to restore precision, Arrow's analysis says, is through allowing a dictator, one participant whose preference function overrides every other. Grinsell does not extend his proposal to faultless disagreement, but it is possible to do so by analyzing faultless disagreement as a conflict between equally valid preference functions. Lacking a dictator, we have no way to decide between them.

However, these models still run into the problems with disagreement that the judge-argument contextualists run into. Having *Mint is tasty* mean *Mint is tasty/tastes good to everyone* or *Mint is tasty/tastes good to my kind of people* or *Mint is tasty according to my preference function* does not account for why Sue can still assert it even after she hears Bob assert *Mint is not tasty*, proving that it isn't tasty for everyone. Also, as Pearson (2013) observes, generics are generally a lot more tolerant of exceptions than taste statements are; in the following example (a) knowing that Mary and I are linguists would still make the utterance felicitous, but in example (b), the first person is not allowed to be an exception, although the third person is.

(64) a. Linguists typically write their papers using LaTeX, but {I use, Mary uses} Microsoft Word.

b. The cake that Mary and I ate was tasty, but {#I,she} didn't like it.

So if neither differences in context nor differences in content give us everything that we need to explain where subjectivity comes from, an alternative approach might be to step outside the semantics entirely, and revert to the pragmatics to show the difference between

2. We will ignore the issue that someone who skates well enough to participate in international ISU-sanctioned events already surpasses some minimal standard of good figure skating.

subjective and objective propositions. So say the scholars featured in the next subsection.

1.4.6 Metalinguistic negotiation: the difference is pragmatic and about prevailing norms

A proposal put forth by Barker (2002), later expanded by Sundell (2010) and Plunkett & Sundell (2013) is that faultless disagreement is actually a matter of negotiation. When we argue whether gymnast Svetlana Khorkina is tall, whether the Chicago Picasso is ugly, or whether avocados are tasty, we are actually negotiating how these words should be interpreted in our speech community. If our speech community decides that Khorkina *is* tall, that means that 5'5" is greater than the minimal height to count as *tall* for us. By coming to a consensus that the Chicago Picasso is beautiful, we are establishing a local norm of what *beautiful* means: it includes things that look like the Chicago Picasso. Ludlow (2008) discusses a faultless disagreement as to whether the racehorse Secretariat should be on a list of greatest athletes of the 20th century, which is a disagreement on whether the meaning of *athlete* includes the requirement *human* or may apply to nonhumans.

In a sense, this proposal hearkens back to the concerns of expressivism, as in (Gibbard, 1992), which strove to separate language about norms and morals from language about facts and the world. However, by placing the difference in the pragmatics of language use, Barker, Plunkett and Sundell circumvent the Frege-Geach problem expressivists had to tackle.

This proposal has interesting explanatory advantages. It explains why people would engage in disputes about taste even though they are aware that *de gustibus non est disputandum*: by declaring our opinions and disputing them, we establish social cohesion and shared meanings and standards for language, or at least, awareness of where our meanings and standards may diverge.

However, Fleisher (2013) argues that this cannot be the whole story: we may be in complete agreement as to what *tasty* (or *good*) means, and the standard that a wine needs to exceed in order to be tasty/good, and still disagree on whether a certain wine meets that

standard. One may wonder, given the direct experience requirement, whether we know for sure that any two people actually can agree on experience, but Fleisher's is an objection well worth considering.

Another criticism is that modelling faultless disagreement as negotiation does not always reflect reality. People may assert that something tastes bad even if they know that everyone else disagrees, and would not accept negotiation of this (although even for food tastes, we must accept some negotiation and changing of our tastes based on someone else's persuasion, else wine-tasting courses would not exist). Also, we would expect much more metalinguistic negotiation: why do we negotiate whether *athlete* includes humans, but not whether, say, *bird* includes ducks?

In (Vardomskaya, 2014) I argued for metalinguistic negotiation explaining all subjectivity, claiming that subjectivity is the default case, and that we come to objectivity through consensus, or through invoking an authority. Although I still believe that there is truth in that, this dissertation presents a deeper and more refined view.

1.4.7 Outlook-based Semantics: the difference is in a simplified context

A new subjectivity model on the market is a version of relativism recently proposed in manuscript by Coppock (2015), called *outlook-based semantics*. Instead of a semantics for subjective predicates involving a possible world and judge, as with Lasersohn's (2005) model and subsequent work, she replaces both world and judge with an *outlook*.

Outlooks have most of the same properties of possible worlds: they determine the extensions of predicates and relations, including both objective and subjective (what Coppock calls *discretionary*) predicates. That is, the proposition *Mint is delicious/Mint tastes good* is the set of outlooks in which mint tastes good; the proposition *Mint is not delicious/mint tastes bad* is the set of outlooks in which mint tastes bad. These outlooks do not overlap, so the statements contradict each other. However, unlike a set of possible worlds, neither set of outlooks contains the actual world.

Coppock still includes sets of worlds in her overall semantics, but only for objective predicates (e.g. *red, wooden, prime number*). An objective proposition describes a set of worlds, as well as a set of outlooks, while a subjective proposition describes only a set of outlooks.

(65) Properties of Outlooks (Coppock, 2015):

a. An outlook-based model is a tuple $M = \langle C, D, D_t, \Omega, A, R, I \rangle$, where
   (i) C is the non-empty set of contexts;
   (ii) D is a set of individuals;
   (iii) D_t is a set of truth values T, F and #;
   (iv) Ω is a set of outlooks;
   (v) A is the subset of D that is the set of agents;
   (vi) R is a set of accessibility relations, one for each $a \in A$, each being a binary relation on Ω specifying the doxastic state for each agent at each outlook;
   (vii) I is a valuation function assigning to each non-logical constant in this language an intension, a function from outlooks to extensions.

b. An outlook o *satisfies* proposition p if and only if o is an element of p.

c. Let ϕ be a sentence and let $[[\phi]]^{c;o}$ be the semantic value of ϕ relative to context of utterance c and outlook o. The proposition expressed by ϕ in context c is the set of outlooks o such that $[[\phi]]^{c;o} = T$, where T stands for 'true.'

d. Two propositions are contradictory if there is no outlook satisfying both of them.

The set of outlooks compatible with an agent's beliefs and opinions are *doxastically accessible* to the agent (thus, Coppock allows agents who do not have a single fixed perspective: they may be ambivalent on some matter of opinion, not be sure what they think, or simply not care). If a proposition is in all of the agent's outlooks, she accepts it; if it is in some but not all, she does not accept it; if it is in none of her outlooks, she rejects it.

(66) Properties of Outlooks (Coppock, 2015):

a. Disagreement: If at a given outlook o, one agent accepts p and another rejects p, the two agents disagree about p in o.

b. Think: Relative to a given context of utterance c, 'a thinks ϕ' is true at an outlook o if and only if, in o, a accepts the proposition expressed by ϕ in c.

c. Satisfaction: Let p be the proposition expressed by ϕ with respect to context c; the presuppositions of ϕ are satisfied in context set S (with respect to context c) iff for every outlook o in S, $p(w) = T$ or $p(w) = F$.

d. Sincerity: An agent a asserts ϕ sincerely in c if a accepts p in $o(c)$.

(67) Let us use $O(w)$ to refer to the intersection of all facts in world w. Since facts are propositions, and propositions are sets of outlooks, $O(w)$ is a set of outlooks. Thus, a world corresponds to an equivalence class of outlooks.

(68) a. An *objective proposition* is one that corresponds to a set of possible worlds, as the union of a set of facts $O(w)$ from a set of worlds w. A proposition is discretionary (subjective) if it is not objective.

b. An agent is *at fault* if in world w, or more precisely, in any outlook in $O(w)$, she accepts an objective proposition that is not a fact in w. Otherwise, she is faultless.

Pragmatically, the essential purpose of conversation is to find a common outlook, as Coppock views the common ground of conversation as a set of outlooks rather than a set of possible worlds as in (Stalnaker, 1978).

However, Coppock's model does not provide a clear way to distinguish worlds from outlooks, even as it elegantly models disagreement and "truth" using them. Although she offers evidence that some predicates can be either subjective or objective, testing them embedding under the Swedish verb *tycka*, which acts similarly to English *find*, she does not explain what is going on in these cases.

(69) Jag tror/tycker att soppan är god.
'I believe/find that the soup is good.'

(70) Jag tror/tycker att det är viktigt.
'I believe that it is important.'

Contrast these with examples that show *tycka* is unacceptable with matters of fact, but acceptable with subjective predicates:

(71) a. Jag tror/#tycker att hon är lärkare.
'I think that she is a doctor.'

b. Jag tror/#tycker att det är tisdag idag.
'I think that it is Tuesday today.'

c. Jag tror/#tycker att jag kommer att vinna.
'I think I'm going to win.'

d. Jag tror/#tycker att det kanske borjar kvart över.
'I think that it maybe starts quarter past.'

(72) a. Jag tycker att skolmaten är god.
'I think that the school food is tasty.'

b. Jag tycker att det är fel att inte hela Sverige hjälps åt.
'I think it's wrong that not all of Sweden helps.'

c. Jag tycker att vi ska ta hand om varandra.
'I think that we should take care of each other.'

d. Jag tycker att den ser ut som en champinjon.
'I think it looks like a mushroom.'

Coppock's model is much simpler, in terms of moving parts, than most of the preceding models. However, it is not immune to criticism. As Kennedy & Willer (2016) note, the model

of outlooks presupposes a fixed and radical distinction between subjective and non-subjective predicates. We have evidence that this is not so, and subjectivity can be context-sensitive, as with Pedersen's (2012) *win the debate* example which I repeat below:

(73) (At the 2012 U.S. presidential election):
A: Obama won the debate last night.
B: No, he didn't. Obama didn't win the debate. Romney won the debate. (SUBJECTIVE - FAULTLESS DISAGREEMENT)

(74) (At the World Universities Debating Championships):
A: Sharon won the debate last night.
B: No, she did not. Sharon didn't win the debate. Wallis won the debate. (OBJECTIVE - NON-FAULTLESS DISAGREEMENT)

In that case, the context makes the disagreement objective: the World Debating Championships, unlike presidential debates, have judges that officially decide winners and losers of every debate round. Thus, subjectivity, or triggering faultless disagreement, whether a proposition is in the world or in an outlook, seems to rely at least partly on the pragmatics of the context.

As well, Coppock does not predict a distinction between those predicates that allow faultless disagreement in both the positive form and the comparative form (*tasty, beautiful*) and those that allow it only in the positive form (*tall, rich*). Nor does she predict any difference between those predicates that can be embedded under *find* and those that can be embedded under *consider*.

A proposal that models what is required of the context in order to trigger subjectivity is Kennedy & Willer's (2016) counterstance model, which is described in the next subsection.

1.4.8 Counterstances: features about the pragmatic context allow coordination by stipulation

Kennedy & Willer (2016) have a recent proposal that seeks to account for the subjective attitude verbs *find* and *consider*. I investigate this account in more depth in Chapter 4 on *find*, but for now I give a brief sketch of it. As we have already seen, the subjective attitude verbs *find* embeds subjective predicates, as shown:

(75) a. Kim finds Lee fascinating, because he is an expert on oysters.
b. #Kim finds Lee vegetarian, because the only animals he eats are oysters.

On the other hand, *consider* allows embedding of both subjective and partially-objective predicates as shown:

(76) a. Kim considers Lee fascinating, because he is an expert on oysters.
b. Kim considers Lee vegetarian, because the only animals he eats are oysters.

Both these verbs reject fully objective predicates:

(77) a. #Kim finds the sum of two and two equal to four.
b. #Kim considers the sum of two and two equal to four.

And as the following example shows, whether they apply or not seems determined by context, not by semantic type:

(78) a. #Kim considers Burgundy part of France.
b. Kim considers Crimea part of Russia.

(a) is infelicitous, whereas in the current global situation as of 2017, (b) is felicitous because the status of sovereignty over Crimea is debated.

Kennedy and Willer argue that the felicity of *consider* and *find* depends on the availability of doxastic alternatives in the discourse, which sets of alternatives they call *counterstances*. In a sense, we have the intuition that it is 'up to Kim' to decide whether a certain word or phrase (*fascinating, vegetarian, part of Russia*) applies in the given context, and there is an equally felicitous option where it does not.

(79) [W]e embellish our context model with a function κ that tracks the contingency of the stipulations involved in achieving an information state. κ takes an information carrier s and derives a set $\kappa(s)$ of s's COUNTERSTANCES: each such counterstance agrees with s on its factual information but disagrees on contextually salient decisions about linguistic practice. (Kennedy & Willer (2016) p. 5)

A counterstance of s is a world just like s except for contextually salient decisions about how to resolve indeterminacy of meaning. Kennedy and Willer are clearly inspired by the Metalinguistic Negotiation model, allowing counterstances to differ in those aspects of language use that are stipulative rather than factual. They define the difference between subjective and objective propositions as follows:

(80) Relatedly, we can say that a proposition is subjective just in case we can find some information state s that accepts itthe proposition is true at all possible worlds in $s$but that fails to be accepted by one of s's counterstances. We label this COUNTERSTANCE CONTINGENCY. (Kennedy & Willer (2016) p. 6)

Essentially, they can stipulate that some x is p, but there are equally valid stipulations available where x is not p, given the facts of the matter.

Consider, according to Kennedy and Willer, presupposes mere counterstance contingency, whereas *find* requires a more stringent condition, which they call *radical counterstance contingency*. With radical counterstance contingency, the prejacent is counterstance

contingent relative to every partition introduced by κ_c , i.e. no matter how we resolve uncertainty of meaning based on parameters that support coordination by stipulation, it remains uncertain.

Kennedy and Willer explain that once they assume counterstance contingency, it does promise a way to explain faultless disagreement:

(81) We have set the issue of faultless disagreement aside partly because it is a non-trivial question how exactly this kind of disagreement is to be characterized in theoretical terms (see, for instance, MacFarlane 2014 for discussion), and partly because it is unclear whether the possibility of faultless disagreement has any distinct semantic implications once we allow for a sufficiently rich conception of the dynamics of conversation (see, for instance, Barker's (2013) model for negotiating contextual parameters for such a conception). Here we just point out that our concept of counterstance contingency also promises to be relevant for the issue of faultless disagreement: treating an issue as counterstance contingent is just to say that the objective facts (whatever those are, according to the conversational context) do not select for a unique resolution of that issue, and intuitively it is exactly the absence of a single correct view on an issue that underwrites intuitions of faultless disputes. One important implication of this approach is that since counterstance contingency is a matter of what alternatives a discourse context provides, we predict faultless disagreement to be an essentially context sensitive phenomenon — just like the acceptability of certain predicates in the scope of [subjective attitude verb]s. (Kennedy & Willer (2016) p. 18)

So in Kennedy and Willer's perspective, as I understand it, the difference between subjective predicates (at least, those that can be embedded under subjective attitude verbs like *consider* and *find*) and objective ones is that the applicability of subjective predicates is underdetermined by the facts of the matter at hand (which may or may not be known —

observers not knowing for sure whether a soup contains pork or whether there is life on other planets does not make these propositions subjective). This allows for flexibility and indeed describes the use of *consider* quite well: we can consider X Y if we have a choice, given the facts, whether X could be Y or not (however, see the Appendix to Chapter 4 for a polysemy involving *consider*). As they themselves acknowledge, the situation seems more complicated with *find*. They use Bylinina's (2016) reasoning that since predicates that are acceptable under *find*, such as *charming, lazy, ugly* or *brave*, lack overt Experiencer arguments in languages that show them like Russian or Hungarian, a direct experience presupposition must not be enough. However, as I argue extensively in Chapter 4, the Experiencer argument as a semantic-syntactic theta role must not be confused with a pragmatic presupposition that the agent or subject has direct experience of the matter, and Bylinina's confusion of the two undermines her analysis. Thus, that *find* rests on direct experience cannot be ruled out, and in Chapter 4 I argue that direct experience is sufficient to explain the distribution of *find*.

Their analysis that all radically counterstance contingent predicates (those that can take *find*) are counterstance contingent (those that can take *consider*) poses a challenging prediction: it should be possible to *consider* something a certain way but have it impossible to *find* it so (just *not finding* it so would violate the presupposition of belief that is preserved under negation). That is, the predicate should be radically counterstance contingent, but the direct experience requirement, or some other factor, is not met. An attempt to bear out that challenge is shown in the following sentence:

(82) John considers mint delicious but cannot find it delicious.

(83) Sue: John, do you find mint delicious?
John: Well, I consider it so.

This would seem to imply that John has not tasted mint but for some reason is in a language use context where claiming mint is delicious is the norm. However, what kind of language use

context would that be? (In a dictatorship that forces the populace to affirm that the dictator's favorite foods are delicious, wouldn't people just lie?) Although Kennedy and Willer's model is quite effective at predicting what should happen if a predicate is counterstance contingent or radically counterstance contingent, they fail to explain what characteristics make predicates such in the first place.

The counterstance model is still undergoing development, but as it was presented in 2016, it laid most of the responsibility for subjectivity on the stipulative aspects of language use; thus, it is based on linguistic negotiation. It provides a framework for why some predicates, the radically counterstance contingent ones, cannot undergo metalinguistic negotiation. But it still does not predict which predicates those would be.

1.5 General criticism of previous models

The modern study of subjective predicates has been focused on using semantic tools (the character-content distinction, tools developed for generics, degree scales and multidimensionality) to try to build a semantics for them that would reflect why they seem to be true when uttered by one person while false when uttered by another.

Semantic models for subjectivity that target certain words as *subjective* tend to run into the same problems. For one, subjectivity does not live purely at the lexical level. Some predicates, like *light* when describing food, slip between subjective meanings (how a food feels to the eater) and objective meanings (the food's mass or fat percentage) with protean ease. Others may turn out to unexpectedly trigger faultless disagreement until resolved by an outside authority: recall the seemingly faultless disputes as to whether Pluto should count as a *planet*, which have now died down. As we saw with Pedersen's (2012) example I showed in the introduction, there can be different contexts that determine whether a proposition is subjective or objective, with no actual difference in the lexical items:

(84) (At the 2012 U.S. presidential election):

A: Obama won the debate last night.

B: No, he didn't. Obama didn't win the debate. Romney won the debate. (FAULTLESS)

(85) (At the World Universities Debating Championships):

A: Sharon won the debate last night.

B: No, she did not. Sharon didn't win the debate. Wallis won the debate. (NON-FAULTLESS)

Thus, subjectivity, or triggering faultless disagreement, whether a proposition is in the world or in an outlook, seems to rely at least partly on the pragmatics of the context. Other than the counterstances model, very few models try to account for this. I discuss the pragmatic effects of subjectivity much more extensively in Chapter 2.

Relativist models account for faultlessness but don't seem to have a very good account of disagreement; contextualist models are the other way around, accounting for disagreement but not faultlessness.

Many models try to account for exocentricity, but other than MacFarlane's, few models try to account for retraction. Thus it seems the only models that can take care of retraction are relativist models.

Since Kennedy (2013)'s observation, several models have tried to address whether there are two kinds of subjectivity (or more); however, the two kinds that they come up with differ (positive vs. comparative faultless disagreement, modals vs. predicates of personal taste, consider vs. find, vagueness vs. predicates of personal taste vs. multidimensionality). If there are two kinds of subjectivity, should we at least be consistent on which two kinds they are?

This dissertation is not going to solve all of these issues. But over the next two chapters, I provide an evidential answer for where the source of subjectivity lies. In Chapter 4, I show how this explains the distribution of *find*. In Chapter 5, I offer an explanation for most of

the puzzles I have presented here, given the sources of subjectivity.

1.6 Conclusion

At the opening of this chapter, we saw that Touchstone and the knight's disagreement about the pancakes and the mustard being *good* or *naught* relied on an assumption that such statements have a truth value. However, if we consider *good* to mean 'tasty', such statements don't seem to be objectively true: the disagreement is faultless and the knight should not lose his honor by staking it on the truth of this assertion. In this chapter, I have given an overview of the major trends in the literature and the major puzzles that subjectivity poses for semantic theories involving context, extension and truth.

Most scholars try to extend models that establish the truth conditions of a preposition to subjective propositions like that of the pancakes and mustard in one of three ways: assuming that subjective propositions stand outside truth (they concern norms, as in expressivism, or language norms, as in metalinguistic negotiation; that is, Touchstone and the knight were disagreeing about how they should act about those pancakes or what *good* should mean); or they allow variation on content among speakers, as in contextualism, adding certain features to the content, as with judge-free contextualism, so Touchstone and the knight had different contents for the meaning of *The pancakes were good and the mustard was naught*; or they allow contextual indices that determine truth to vary, as with relativism, so Touchstone and the knight resolved to different people the reference of whom the pancakes are good to.

None of these explanations, though, fully explain why the knight felt justified in staking his reliability in telling the truth on the claim that the pancakes were good: a norm comes about through negotiation rather than through one person's reliability; and both differing contents and contexts should allow that what the knight said was true according to him, and thus he should not lose his reliability. And few of them explain why the knight does lose his reliability if instead of claiming that the pancakes were delicious when they were not, he claimed that the pancakes were fresh when they were spoiled.

Fewer scholars try to address how to tell apart subjective and objective propositions and predicates in this manner. For many, if there is a difference, it is taken for granted that subjective propositions have a certain property (like a special argument) or that most words have different speakers take the same index by default, as with relativism. Bylinina (2013, 2016) makes a worthwhile effort and finds three different sources for subjectivity, and Kennedy & Willer (2016) find two (with one being a subtype of the other). However, as I point out, and explain further in Chapter 4, Bylinina's model overgenerates, whereas Kennedy and Willer's focus on stipulation via language norms does not fully explain what it means for a predicate to be predictably counterstance contingent.

But what if we treat the knight's honor in a different way — as a commitment to integrity with his own knowledge and cooperation with his interlocutors rather than as reliability in giving true statements? In the next chapter, I explore more pragmatic effects of subjectivity that have been less discussed in the literature. This will lead the way to presenting my own model of the difference in Chapter 3.

CHAPTER 2

THE PRAGMATICS OF SUBJECTIVITY

BARRY: No, not nothing. What's wrong with the Righteous Brothers?

DICK: Nothing. I just prefer the other one.

BARRY: Bullshit.

ROB: How can it be bullshit to state a preference?

— from *High Fidelity* (2000)

2.1 Introduction

In the previous chapter, we looked at other scholars' views of faultless disagreement and subjectivity, which mostly revolve around standard examples like the following:

(1) a. Sue: Avocados are delicious.
Zoe: No, they're not. (both are right)

b. Sue: Avocados are vegetarian.
Zoe: No, they're not. (Zoe is wrong)

In the conclusion, I pointed out that most previous analyses have one crucial gap: they take the difference between subjective and objective predicates for granted. It is stipulated that there is some difference between *delicious* and *vegetarian* that allows for faultless disagreement to happen with one but not the other, yet it's hard to tell what that difference is, since these adjectives behave very similarly in the syntax otherwise. The goal of this dissertation is to fill that gap.

In this chapter, I start by refining the picture. I show a number of other patterns that occur with subjective predicates (predicates of personal taste, aesthetic predicates, moral predicates and modal predicates) besides the conventional faultless-disagreement pattern. I show how these patterns lead us to certain conclusions about what must have been asserted

and added to the common ground. The patterns are the following: disagreement as update rejection (Section 2.2); the case of faulty disagreement (Section 2.3); the Acquaintance Inference requirement (Section 2.4); and the uptake requirement (Section 2.5).

This will lead us to conclude, I argue in Section 2.6, that the correct model for how subjective propositions p update the common ground between speakers A and B is that they update it with a conjunction of p with respect to A and p with respect to B. However, that is also true, if trivially so, for objective propositions, and so objective propositions only differ from subjective ones in that the evidence is treated as the same for speaker and listener.

We will then discuss evidence in Chapter 3.

2.2 Disagreement as update rejection

Given that the most common definition of subjective predicates in the literature is as those that allow faultless disagreement, let us take a closer look at what this disagreement entails. In this section, I will present evidence to show problems with the traditional understanding of disagreement that many scholars espouse, both on the contextualist and the relativist side.

We have an intuition of when people are disagreeing. Moreover, there is an intuition that when two people disagree, they cannot both be asserting something true: their claims are exclusionary.

(2) **Definition**: Two statements are exclusionary if and only if, when evaluated in the same context (world, time, indexicals), at least one of them must be false.

(We will ignore trivial examples like one of the statements being a logical contradiction like 'X and not X.')

This certainly sounds plausible with the typical examples of disagreement:

(3) Sue: 27 is a prime number.

Bob: No, 27 isn't a prime number.

(4) Sue: There is life on other planets, somewhere.

Bob: No, there's no life on other planets, we are alone in the universe.

Both of these pairs of claims are exclusionary. A number cannot be prime and not prime at the same time, and the universe either contains or does not contain life outside of Earth. It may very well be that both claims are false:

(5) Sue: The capital of California is San Francisco.

Bob: No, the capital of California is Los Angeles.

Sue and Bob are disagreeing and are both factually wrong (the capital of California is Sacramento).

Both contextualists and relativists seem to accept this model that disagreement means exclusionary content, and thus there must be some way to assign one claim in a faultless disagreement the value 'true' and the other one the value 'false.' However, most models to encode this end up fairly contrived, involving a judge argument or judge index. As I mentioned above, such models have many flaws, and do not encompass some real evidence about disagreement and exclusionary content. I devote this section to discussing recent experimental evidence from Khoo & Knobe (Forthcoming 2016) for disagreements about moral actions that shows that disagreement and exclusionary content do not have to go hand in hand, and in fact disagreement is better viewed as rejection of common ground update.

As Khoo and Knobe point out, there may be an intuition of disagreements without exclusionary claims, such as about proposals (which do not have a truth value):

(6) Sue: Let's go for a coffee.

Zoe: No, let's go for a beer. (based on Khoo & Knobe (Forthcoming 2016), ex. 4)

I can come up with a similar sense of disagreement about imperatives:

(7) a. Mom (to child): Clean your room!
Dad (to same child): Practice piano!

b. Mom (to child): Clean your room!
Child: No, I won't!

Imperatives do not have truth values — yet in the first case, we get the intuition that the mother and the father are disagreeing with each other as to the child's first priority, while in the second case, the child is in evident disagreement with the mother.

Khoo and Knobe did a series of experiments that showed that in the case of moral disagreement, people's judgements about exclusionary content systematically come apart from their judgements about disagreement. Given the example of a hypothetical transgressor, Dylan, who tested a new-bought knife by stabbing a passerby with it, subjects were presented with dialogues in which two speakers disagreed as to whether what Dylan did was morally wrong.

(8) a. A: What Dylan did was morally wrong.
B: What Dylan did was morally permissible.

Subjects were told that A was an American college student in all cases; however, the identity of B varied across conditions, with B being either another American university student, a member of a warlike Amazon tribe, or an extraterrestrial. They were asked whether one of the two speakers had to be wrong.

Khoo and Knobe found that in the case where B was of the same culture as A, B's statement was viewed as an exclusionary claim with respect to A; however, subjects were divided when B was a member of a distant culture, and in the case where B was an extraterrestrial, subjects generally disagreed that one of the two speakers had to be wrong. Thus, Khoo

and Knobe have evidence that a judgement of disagreement is possible without exclusionary content.

Khoo and Knobe theorize that the intuition of disagreement that subjects have is not about the incorrectness of one speaker's claim (as a model assuming exclusionary content would predict) but about rejection of the speaker's claim from the common ground. They tested this in a followup experiment, where subjects read the following vignette in one of three conditions, and were asked to judge either the Incorrectness or the Rejection interpretation of the vignette:

(9) Sam is having a discussion with one of his classmates (in other conditions, Sam and his classmate are two individuals from a warlike Amazon tribe, or two extraterrestrials). Eventually, the conversation turns to a recent event. A person named Dylan bought an expensive new knife and tested its sharpness by randomly stabbing a passerby on the street. Sam says, about this case, "Dylan didn't do anything morally wrong."
As it happens, Jim is listening in on Sam's conversation, and believes that Dylan did do something morally wrong. Jim jumps into Sam's conversation and says, "No, Dylan did do something morally wrong."

a. Incorrectness: Given that Jim and Sam have different judgments about this case, we would like to know whether you think at least one of their judgments must be incorrect, or whether you think both of them could actually be correct. Please tell us to what extent you agree or disagree with the following statement: *Since Jim and Sam have different judgments about this case, at least one of their judgments must be incorrect.*

b. Rejection: Given that Jim and Sam have different judgments about this case, we would like to know what you think about Jim's response to Sam's claim. Please tell us to what extent you agree or disagree with the following statement: *Since Jim and Sam have different judgments about this case, it was appropriate for Jim*

to reject Sam's claim by saying "No".

Subjects consistently agreed with the Rejection interpretation more than they did with the Incorrectness interpretation, with judgements of Rejection being consistent across the three conditions (of Sam and Jim being of the same culture, of different cultures, or Sam being an extraterrestrial), while Incorrectness judgements significantly declined the more different Sam's culture was from Jim's. This supports Khoo and Knobe's theory that disagreement, at least about moral judgements, is much more about rejection of the claim from the common ground than it is about a conflict in truth values. Several other followup experiments also supported this claim.

I wish to give more prominence to Khoo and Knobe's view about disagreement involving rejection of the claim from the common ground, rather than about exclusion based on truth value (if not necessarily adopt it outright; I want the conclusions I come to in Chapter 3 to be as theory-neutral as possible, as I show in Chapter 5). Such a view allows me to sidestep the problem of relative truth and its attendant issues of where the judge index comes from and what does it mean to be true *tout court*. However, there are two problems that I need to resolve.

One is that Khoo and Knobe take the difference between moral and non-moral judgements for granted when they show, in another experiment, that the Rejection interpretation of disagreement is not significantly preferred to the Incorrectness condition for non-moral judgements (the example they give of a non-moral judgement is the factual claim "Napoleon always used to go into battle on a helicopter," a claim made by an extraterrestrial in their test vignettes). However, distinguishing between non-moral/objective claims and moral and other subjective claims is the topic of this dissertation, and so I have to address this. In Chapters 3 and 4, I discuss the differences in evidence for moral and other subjective claims.

The other is that if Khoo and Knobe interpret disagreement about moral (and other subjective judgements) as being about rejection of the claim from the common ground rather

than about truth or falsity, how should we interpret cases when *true* or *false* is explicitly used in a disagreement most readers would judge as faultless, such as the one we have already seen about the merits of a food show:

(10) (Example in an online chat between University of Chicago students, used with permission, October 2016):
A: *Diners, Drive-Ins and Dives* is a fantastic show.
B: Literally false.
C: Literally true.

B certainly sounds like she is rejecting A's claim from the common ground, and C sounds like he is rejecting B's. However, I will claim that B and C are misusing the terms *true* and *false* or assigning them a different meaning. Our English-speaking culture, at least, does not have many established paths to rejecting a proposed update other than judging it false. "It's true but I'm not accepting it" is strongly marked; "Let's not talk about this" is perhaps the only other plausible option, but one cannot alleviate the prospect of being fed a food one dislikes by asking not to talk about it, so "This food is not tasty" is required as a way of rejecting that pressure to eat it.

There are in fact situations where rejection is simply infeasible, and as subjective propositions strongly interact with those, I turn to them in the next section.

2.3 Faulty disagreement

The cases of "faulty disagreement" (term coined by Gunlogson & Carlson (2016); called "faultless disagreement" by Dechaine *et al.* (2014); I will use the term "faulty") are where contradiction in general is not felicitous at all:

(11) Faulty disagreement:
Sue: This ride is fun for me.

Zoe: #No, it isn't.

Compare this to the following example:

(12) Sue: I have a headache/I'm cold/I find this ride fun.
Zoe: #No, you don't/#That's not true.

Although faultless disagreement has been a hot topic for investigation since well before Kölbel (2002) gave the phenomenon a name, its relationship to faulty disagreement has not been investigated very widely. As the case of Dechaine *et al.* (2014) shows (they label *to me* constructions as "faultless disagreement"), some scholars may still find the two confusing, while most others focus on faultless versus objective disagreement, with faulty disagreement ignored. However, there is a close relationship between predicates that license faultless disagreement and those that license faulty disagreement, and it is worth investigating. For example, the subjective attitude verb *find* in English, widely noted since (Saebø, 2009) to embed subjective propositions, creates contexts where any disagreement would be faulty. (I devote Chapter 4 to discussing *find* and proposing a reason for this.) In the following example, Zoe's disagreement is infelicitous:

(13) Sue: I find this ride fun.
Zoe: #No, it isn't.
#No, you don't/#That's not true.

Similarly, adding "to me" or "for me" forces any disagreement to become faulty, although as Stephenson (2007a,b) points out, in English, such prepositional phrases are only acceptable with a subset of subjective predicates and their choice is idiosyncratic. We see this example as also infelicitous:

(14) Sue: This ride is fun for me.

Zoe: #No, it isn't.

We add an additional factor to this: that if the disagreement is about a subjective proposition anchored to someone else, it is treated as an objective proposition, and disagreement about it is not faultless. I illustrate this with an example similar to Gunlogson & Carlson (2016):

(15) Sue: Avocados are tasty to Bob.
Zoe: No, they're not!

This is an objective disagreement; either Bob believes that avocados are tasty, or he does not. Notably, Sue does not need to either have tasted avocados herself, or have even seen Bob enjoy them or tell her he does. She may assert this based on simple rumour or thrice-removed hearsay.

As Gunlogson & Carlson (2016) point out, the reason that Zoe disagreeing with Sue is faulty is that Sue has much stronger evidence than Zoe does. Predicates of personal taste, such as *tasty* and *fun*, have as their best evidence direct experience. So in a situation where Mary asserts that the ride was fun for her, John cannot contradict her because she has much better evidence than he does; he would have only indirect evidence (say, her screaming in apparent terror). This makes the following dialogue infelicitous:

(16) Mary: This was fun for me.
John: #No, it wasn't.

The infelicity is quite similar to pain reports or sensory reports, which are also infelicitous to disagreement:

(17) a. John: I have a headache.
Mary: #No, you don't.
b. John: This blanket feels soft to me.

Mary: #No, it doesn't.

They note that disagreement is only faultless, when the propositions are stated generally, if the speakers have equal evidence — for example, if Mary and John both rode the roller coaster, the dialogue is felicitous:

(18) Mary: That was fun!
John: No, it wasn't.

If Mary rode the roller coaster, but John sat it out and met her at the gate, the dialogue would be infelicitous:

(19) Mary: That was fun!
John: #No, it wasn't.

Mary has a much stronger acquaintance with the experience of riding the roller coaster, and much stronger evidence for her position in the case of (19).

This leads me to a deeper discussion of what this kind of acquaintance entails, in the next section on the Acquaintance Inference. Before I proceed there, I will take a look at the larger picture of how faulty-disagreement propositions and subjective propositions may fit together. Faulty-disagreement propositions form a much larger set than subjective propositions, although the two are related. We see the following examples of categories of discourse where disagreement is faulty:

(20) *Expressive content:*
Sue: Oops!/Ouch!/Yum!
Zoe: #No, that's not true./#You're mistaken.

(21) *Pain and cold reports:*
Sue: My head hurts.

Zoe: #No, that's not true./#You're mistaken.

(22) *Not-at-issue content:*

Sue: Mary, who is a World Champion, won the race.

Zoe: No, she didn't win/#No, she isn't a World Champion.

(23) *Evidential content:*

Sue: Mary apparently/evidently won the race.

Zoe: No, she didn't win./#No, it's not apparent/evident (to Sue).

All of these faulty-disagreement utterances have felicity conditions but they do not seem to have truth conditions: it is infelicitous to assert you are in pain if you're not in pain (Kaplan 1999), it is infelicitous to assert X with a hearsay evidential if you actually saw it (Faller, 2002), but what is going on there would not be called *lying*.

I do not believe that one can have a generalized theory of all of these faulty-disagreement propositions. Korotkova (2016) tries to connect subjective content and evidential content, and I discuss this in Chapter 3 on evidence, and at length in Chapter 5. Murray (2011) and Faller (2002) analyze evidential content as being not-at-issue, trying to link not-at-issue content and evidential content, but I am not sure their analysis can extend to subjective content. I will treat the possibility of connecting not-at-issue content and subjective content as beyond the scope of this dissertation, and leave it as an interesting avenue of future research.

In the next section, I elaborate on what it means for Mary, who had ridden the roller coaster, to have better evidence than John, enough evidence to assert that riding the roller coaster is fun.

2.4 The Acquaintance Inference

One notable difference between subjective and objective predicates is that if we have not seen Paris, we can assert that it is the capital of France, but we cannot assert that it is beautiful. Several researchers (Pearson, 2012, 2013; Bylinina, 2013; McNally & Stojanovic, 2015) have previously observed that for a speaker to assert a predicate of personal taste, they must have experienced it in the relevant way; for example, one must have tasted a food in order to pronounce it delicious or disgusting, or seen a film or picture in order to pronounce it beautiful or ugly, as the following examples show:

(24) a. Avocados are delicious, #but I've never tasted one.

b. The Chicago Picasso is beautiful, #but I've never seen it, even in photos.

c. *Andrei Rublev* is a great movie, #but I've never seen it.

This seems to apply to predicates of personal taste and aesthetic predicates, but does not apply to moral judgements, even though those could also undergo faultless disagreement. One needs not have experienced a moral transgression in order to find it immoral, since one can felicitously say things like the following:

(25) Arson is wrong, as is robbery. Thankfully, I haven't experienced either.

This requirement that the speaker must have direct experience (of the relevant kind) of the object of a subjective predicate in order to assert a proposition with that object and predicate is referred to as the *Acquaintance Inference*. This section concerns itself with recent research on the nature of this acquaintance, as I present Ninan's (2014) arguments about it. I will leave the section with Ninan's conclusion, and reserve my own answer to it for Section 2.6.

What is puzzling about the Acquaintance Inference is that, Ninan (2014) shows, it is not a presupposition. Personal acquaintance with the object in the case of *tasty* is not in fact presupposed.

For one, the standard model of presuppositions is based on sentences like the following:

(26) a. The Prime Minister of Egypt gave/didn't give you that ring. (Presupposition: Egypt has a Prime Minster.)

b. John stopped smoking/didn't stop smoking. (Presupposition: John has smoked at one time.)

These standard presupposition contexts resist cancellation of the presupposition when the context is positive:

a. #The Prime Minister of Egypt gave you that ring, but Egypt doesn't have a Prime Minister.

b. #John stopped smoking but he never smoked at all.

However, as Ninan (2014) points out, there is a special emphasis pattern, in negative contexts, that allows for the presupposition to be cancelled:

(27) a. John didn't *stop* smoking – he's never smoked a cigarette in his life!

b. The prime minister of Egypt didn't give you that ring, *for Egypt has no prime minister.*

Ninan argues that if the Acquaintance Inference was a presupposition, then one would expect it can be cancelled in negative contexts in a similar fashion. However, this gives bizarre results:

(28) a. ??The lobster rolls at Neptune Oyster aren't *tasty* – I haven't even tasted them!

b. ??St. Mark's Basilica isn't *beautiful* – I've never seen it! (Ninan 2014, exx. 21-22)

Ninan comments, "These are odd: [b], for example, seems to suggest that something cannot be beautiful unless the speaker has seen it — a rather perverse interpretation of the dictum

that 'beauty is in the eye of the beholder.' But if the Acquaintance Inference is a matter of presupposition, it is difficult to see why it can't be canceled via this mechanism."

Second, Ninan points out that there are other well-known "presupposition holes" besides negation — that is, constructions where a presupposition still projects, like the antecedents of conditionals, questions, and modals. Standard presuppositions like the one associated with *stop* project over them: in each of the following cases, it must be true that John has smoked in the past.

(29) a. If John stopped smoking, his doctor will be happy.
b. John might stop smoking.
c. John probably stopped smoking.
d. John must have stopped smoking.
e. Did John stop smoking?

But the Acquaintance Inference does not preserve any of these cases: in Ninan's example *The lobster rolls are tasty* has the inference that the speaker has tasted the lobster rolls, but none of the following examples do:

(30) a. If the lobster rolls are tasty, I'll have two.
b. The lobster rolls might be tasty.
c. The lobster rolls are probably tasty.
d. The lobster rolls must be tasty.
e. Are the lobster rolls tasty?

Pearson (2013) also observes this fact about epistemic *must* and tries to explain it as a kind of indirect evidential. However, Ninan criticizes this approach: this would predict that the Acquaintance Inference would be much easier to cancel, by simply making the fact that the speaker has indirect evidence explicit: *These cookies are tasty, but I haven't tasted them.*

In an effort to define the Acquaintance Inference in some other way, Ninan brings in the notion of an autocentric context and the Knowledge Norm, which is a way of defending the Maxim of Quality:

(31) **Acquaintance Principle (AP)**: Normally, in an autocentric context c, s_c knows (at t_c in w_c) whether $[[o \text{ is tasty}]]_c$ is true only if s_c has tasted o prior to t_c in w_c.

(32) **Knowledge Norm**: For all contexts c, s_c must assert p only if s_c knows p at t_c in w_c.

Or in plain language, the only way to know that something is tasty is to taste it, and we assert things only if we know them (else one is lying or bullshitting). Importantly, the Knowledge Norm does not itself rely on the assertion being subjective or objective.

In Ninan's explanation the Acquaintance Inference is an epistemic implication, which cannot be cancelled or targeted by specially-marked presupposition-cancelling negations, similarly to Moore-paradoxical sentences:

(33) a. ??It's raining, but I don't believe it's raining.
b. ??It's raining, but I don't *know* it's raining.

Ninan points out, however, that the account of the Acquaintance Inference as an epistemic implication has its own problems: it does not explain how it disappears under modal or future operators, as in the examples:

(34) a. The lobster rolls might/must be tasty, but I haven't tried them yet.
b. The lobster rolls will be tasty when they're done. It's a shame that they contain arsenic and will have to be destroyed.

Ninan says that the biggest problem with the Acquaintance Principle, though, is why is it required: "The defender of AP needs to come up with a plausible account of the meaning of

taste predicates which helps to explain why autocentric taste propositions cannot be known unless the agent has the relevant kind of first-hand experience (and why exocentric knowledge is not subject to this requirement)." (p. 307).

An additional issue, which Ninan does not consider as he focuses solely on predicates of taste, is why the Acquaintance Principle does not seem to apply to moral judgements, even though those are generally treated as subjective, and we can felicitously assert things like the example I repeat again:

(35) Arson is wrong, and so is robbery. Thankfully, I've never experienced either.

I present a sketch of an explanation for this in Section 2.6 and elaborate on it in the next chapter. For now, I summarize this section as follows:

Subjective assertions, at least for predicates of personal taste and aesthetic predicates, seem to require that the speaker be acquainted with their object in some relevant way. This requirement is not a presupposition, seems to be related to Moorean effects, and disappears in modal and future contexts. This requirement, combined with the model of disagreement as update rejection and the option of disagreement being faulty, leads to subjective propositions behaving in certain ways in dialogue, which I will discuss in the next section.

2.5 The Uptake Requirement

In this section, I fit together the model of common ground update, and disagreement as common ground rejections, with the acquaintance principle.

The Stalnakerian model of common ground update holds that speakers update the common ground with propositions that they are committed to. Assertions are proposals to update the common ground; disagreements are rejections. By this theory, if Sue asserts, "My brother is arriving tomorrow" or "London is the capital of England," and Zoe does not object, Zoe is now aware and committed to the truth of Sue's brother's arrival (and the

presupposed fact that Sue has a brother, if that was not already common ground knowledge) or the truth of England having situated its head of government in London. Zoe can now felicitously relate to Chloe, "London is the capital of England" or "Sue's brother is arriving tomorrow," if she judges Sue as a reliable person who is neither lying nor confused, even if Zoe had never been to England or met Sue's brother.

If Sue asserts in Zoe's hearing, "Avocados are delicious", there are three options for Zoe's response, which crucially depend on Zoe's knowledge state and attitude. She may have tried avocados and loved them, in which case she can felicitously agree, "Yes, they are!" She may have tried avocados and thought they were disgusting, in which case she would reject the assertion, "No, they aren't!" Or she may have never tried avocados. In that case Zoe cannot felicitously reject Sue's claim that avocados are delicious.

So at first glance, in the last case it seems that the update *Avocados are delicious* entered the common ground, encountering no obstacles to doing so. However, Zoe cannot now felicitously tell Chloe, "Avocados are delicious" if Zoe herself has never tasted avocados, even if she considers Sue's taste perfectly reliable. Zoe may tell Chloe, "Sue thinks/believes/tells me avocados are delicious," or mark it with a number of other evidentials, but unlike "Sue's brother is arriving tomorrow" or "London is the capital of England", Zoe cannot assert the bare phrase if she has heard and accepted it from Sue. This is what Ninan (2014) defines as the autocentric requirement and Acquaintance Inference, as discussed in the previous section.

If we posit that the content of Sue's update was *Avocados are delicious,* then we cannot explain why the Zoe who has not tasted avocados can't just take this up and felicitously assert it herself. On the other hand, if we posit that the update was *Sue believes avocados are delicious,* then we have no explanation for why the Zoe who has tasted avocados and disliked them would reject this, as what Sue likes and does not like to eat should not affect Zoe, so a fact about it should safely be in the common ground. Even if she had vehemently argued with Sue, "No, avocados are not delicious, they are disgusting!", Zoe can still later

say to Chloe, "Sue thinks avocados are delicious, can you believe it?" The fact of Sue's belief has been updated to the common ground, even if the content of it has been rejected.

Farkas & Bruce (2010) propose that the common ground is more fine-grained than Stalnaker has suggested, and deal with a similar problem by using the concept of a Table, a discourse component that records the content of assertions and polar questions until they have been answered or accepted/rejected. Every time an item is placed on the Table, it projects a *projected set* of future common grounds in which the item is decided. Essentially, Farkas and Bruce have the following components to a discourse structure K:

- A set of discourse commitments: the propositions DC_X for each participant X that X is publicly committed to, which are not shared by the conversation's participants.
- The common ground: a set of propositions *cg* that are shared by the conversation's participants.
- The Table representing the issues to be resolved, which is a stack of pairs of syntactic objects and their denotations.
- The projected set *ps* which is the set of future common grounds projected by elements of the Table.

To simplify, I will follow Umbach (2016) in omitting the syntactic objects, although unlike her, I will incorporate the projected set. I will indicate any items placed on the Table as denotations.

In this model, if Zoe does not have the grounds to reject or accept the item *Avocados are delicious*, it would stay on the Table until she does. The issue with that model is that Farkas and Bruce define, "As long as there are items on the Table, there are issues that need to be dealt with. A conversation is in a *stable state* when its Table is empty; a conversation is at a natural endpoint only when it is stable." However, taken without elaboration, that would predict that as long as Zoe cannot either take up or reject *Avocados are delicious*, Sue and

Zoe's conversation is in a unstable state. But in real life, we often hear people assert their tastes and opinions which we do not have the experience to accept or reject, and speakers know this, as illustrated with such common assertions as the following:

(36) a. The new film is boring and badly written.

b. Slovenia is spectacular.

c. The performance I saw last night was great.

We can assert these things without expecting that our interlocutor has the experience to agree or disagree, if they have not yet seen the film, have never been to Slovenia, or missed the performance. Our interlocutor would generally ask us for more detail, and could tell other people about our opinion of the film/vacation destination/performance, couching it in some form of hearsay evidential:

(37) a. Sue says the new film is boring and badly written.

b. According to Tamara, Slovenia is spectacular.

c. I hear that performance was great.

We do not get the sense that us not being able to agree or disagree with the speaker's assertion puts the conversation in an unstable state, yet we cannot relay their assertion directly as true.

Thus the Table model works excellently for objective statements but not for subjective ones. What is placed on the Table when a speaker asserts a subjective proposition cannot be the (denotation of the) subjective proposition itself, as this cannot then be repeated by anyone who did not reject it.

Malamud & Stephenson (2015) elaborate on Farkas and Bruce's model based on how tag questions ("It's X, isn't it?" "It's Y, is it?") and rising intonation ("What is the capital of New York state?" "It's Albany?") affect the pragmatics through modifying the force

of commitment that the speaker makes. (Independently, Wolf (2015) provides an analysis of gradable force of assertion in his dissertation work, and comes to similar conclusions.) They add projected commitments as independent components of the conversational scoreboard, distinct from public commitments: "Unlike F&B's system, this allows for moves that give tentative commitments (by adding propositions to the speaker's projected, rather than present, commitments), or to offer the speaker's best guess of commitments of other participants (by adding to others' projected commitment sets)" (p. 288.)

They also add a set of Common Standards to the conversational scoreboard, setting negotiation about standards for vague predicates as separate from negotiations about propositions in the common ground. Thus they can model agreement, negotiation and disagreement about vague predicates, whether a particular shade of paint is *red* or whether someone is *tall*, which F&B's model has trouble doing.

Malamud and Stephenson do not subscribe to the Acquaintance Inference as firmly as most other scholars on predicates of taste; they affirm,

(38) The relationship between taking on a discourse commitment towards a proposition involving taste and actually being in a position to make that taste judgment is indirect and asymmetrical in the same way. For instance, if person A has seen person C and judged C to be attractive, A is certainly in a position to commit herself to the proposition that C is attractive. But if A then tells B that C is attractive, then even if B has no basis for a judgment herself, she is still in a position to commit herself to the proposition that C is attractive, provided she has some reason to defer to A's judgment. – (Malamud & Stephenson, 2015), p.278.

At the same time, in their examples on predicates of taste, the infelicity relies on violation of the Acquaintance Inference, as I show with one of their examples:

(39) Context: A and B are gossiping. A doesn't know anything about B's neighbor.

B says, blushing, 'You've GOT to see this picture of my new neighbor!' Without looking, A replies:

a. #A: He's attractive, isn't he?

b. ^{OK} A: He's attractive, is he?

c. ^{OK} A: He's attractive?

d. #A: He's attractive.

((Malamud & Stephenson, 2015), ex. 2)

The failure of the Acquaintance Inference is why A cannot assert (d). Even when her statement is modified by the Reversed Polarity tag *isn't he?*, A is committing to having seen enough of B's neighbour to judge whether he is attractive. But by modifying the statement with a Same Polarity tag *is he?*, or via rising intonation, A can place the taste judgement about B's taste in her projected commitments, but not in her public commitments.

Beltrama (2016) points out that subjective statements have an inquisitive characteristic, which would be further evidence for projected commitments. "[S]ubjective assertions always need to be explicitly confirmed to ensure that agreement has been reached. Failure to confirm can be seen as a sign of the interlocutor's being agnostic, or even as a sign of disagreement, where the interlocutor is refraining from expressing dissent via politeness, but never as a sign of agreement." (Beltrama (2016), p. 46.) He continues, "Disagreement with subjective assertions, in some cases, might be in fact *welcome* [italics original]. In other words, a speaker might just be trying to elicit their interlocutor's perspective without particular argumentative proposals." He illustrates the requirement for confirmation with examples like the following (his ex. 65):

(40) a. Greg: My dad will beat your dad.

Fred: [Silence]

Greg: # Good, we are on the same page.

b. Greg: The movie was awesome.

Fred: [Silence]

Greg: # Good, we are on the same page.

He illustrates the welcoming of disagreement with his example 67, which contrasts subjective and objective statements:

(41) a. Greg: The movie was awesome.

Fred: No, I didn't like it.

Greg: Ahah, that's hilarious! I knew you were the kind of person that doesn't like these movies.

b. Greg: The movie started at 8 pm.

Fred: No, it didn't. It started at 9 pm!

Greg: # Oh, that's interesting! We have different perspectives on this.

Thus, as Beltrama points out, the purpose of a subjective assertion is to elicit others' experiences; it has an inquisitive character as well as an assertive one:

(42) Specifically, I propose that assessments containing subjective predicates partially overlap with both factual assertions and polar questions. On the one hand, they have an informative nature, as they convey information about the speaker's view on a particular issue. Similar to assertions, then, they add p as judged by the speaker to the list of their discourse commitments. On the other hand, they are inquisitive, as they summon the other participants' view on the proposition, raising the issue as to whether the proposition is taken to be true by the collective judge of the conversation and should therefore be added to the Common Ground. On a broader level, the upshot of this proposal is that the dichotomy between informative and inquisitive assertions is mediated by a third, mixed type of movement. (Beltrama (2016), p.48)

(43) I capture this component of subjective assertions by suggesting that, in a similar way to polar questions, they partition the Table in two equivalent alternatives. In one such alternative, all participants in the conversations agree to take p to be true, thus projecting the addition of the proposition to the Common Ground (by assertion norm). Assuming a conversation with A and B as the only participants, I represent this option as $[p_A \wedge p_B]$. In the other alternative, it is not the case that a consensus around p is reached by all participants. This happens if at least one participant does not judge p to be true. Assuming again a two party conversation, I represent this alternative as $\not[p_A \wedge p_B]$, where the addition of p to the Common Ground is projected in the PS. (Beltrama (2016), p.50)

The distinction that p_S and p_Z have a chance of being different is what in fact explains the uptake requirement — the sense that if Zoe does not explicitly agree with Sue about Sue's statement of p, then she is disagreeing with her. Explicitly rejecting proposals from the common ground bears a social cost. In much of European-based culture, and in many other cultures, we do not like saying no. The difference between saying no to a subjective proposition and saying no to a factual proposition is that the option of rejection is a much more live option with subjective propositions — another person may well have different knowledge. Thus, we read deliberate silence as such a rejection, while to a factual proposition, avoiding explicit rejection yet disagreeing would be odd (but plausible in certain very polite contexts: "I knew he was totally wrong, but I didn't want to make a scene").

Therefore we have the following model, which fits with the general pattern of assertion: when Sue asserts p to Zoe, that *Sue believes p* or *p is true to Sue* is automatically added to the common ground. However, what is on the Table or in the projected commitments for negotiation, partitioning the projected set into two possibilities, is that p is true in both Sue and Zoe's experience, that is $p_S \wedge p_Z$, or not. Zoe may reject it by asserting that p is not true to her, p_Z is false, and therefore the whole conjunction fails and we are in the $not-(_S \wedge p_Z)$

common ground.

This makes sense in order to account for what Zoe may afterwards assert to a separate party (Chloe, in our toy world). She may always repeat that p_S is true ("Avocados are tasty to Sue"), regardless of what she thinks about p in general (believes p, believes *not-p*, or lacks information). She may repeat p or *not-p* only if she herself has the relevant information for p or *not-p*.

This kind of conjunction of group members' perspectives, that *tasty* means *tasty to the whole group*, had already been considered and dismissed, by Lasersohn (2005, 2017), but due to a different notion of what disagreement means and an adherence to the notion of exclusionary content. Lasersohn suggests the following counterargument: suppose John and Mary are riding on a roller coaster, and Mary is enjoying it but John is not. The following dialogue may take place:

(44) John: This is not fun.
Mary: Oh yes, it is!

Lasersohn discusses this as follows:

(45) Here, as before, we feel that John and Mary are contradicting each other. We can account for that intuition by claiming that both their utterances involve indexical reference to the same group, which contains both John and Mary as members. But then it becomes very hard to make sense of Mary's response. John is riding the roller coaster; hence he is in a perfect position to know that the roller coaster is not fun for him. Under the assumption that something is fun for a group only if it is fun for all members of a group, it follows directly that the roller coaster is not fun for any group containing John, including whatever group is implicitly referred to in John's and Mary's utterances. Mary is therefore in no position to refute the claim that the roller coaster is not fun for the group. What justification could she have for

> contradicting John's claim? By contradicting him, she must be acting irrationally, or ignoring what he said, or claiming to know his own mind better than he does himself, or something similar. But in fact she does not seem to be doing any of these things; she is simply expressing disagreement with John's claim that the roller coaster is not fun. (Lasersohn (2017), p. 22.)

However, under the analysis of disagreement as update rejection, John and Mary's disagreement can be seen to make sense. What John has proposed on the Table for update to the common ground is the conjunction *This is not fun for John* $\wedge$ *This is not fun for Mary.* Mary is rejecting this by asserting that it is indeed fun for Mary, making *This is not fun for Mary* false and therefore the entire conjunction false.

But, we must note, this is also true of objective predicates — if we take an objective predicate p to have p_S and p_Z be identical, it becomes the trivial case of a subjective predicate. All of us have a *p is true/false to me* for any proposition we are informed about, and we can only assert, confirm or reject those propositions we are informed about (if we are being cooperative speakers). In a sense, the only distinction here is that p_S and p_Z have a chance of being different. In all other ways, the more we've dug into the difference between subjective and objective predicates, the more we've realized that they are very nearly the same.

2.6 Analysis

Let's summarize what we've seen so far. We have seen three issues about faultless disagreement that rely on each other:

- 1. A speaker's disagreement with p may be an update rejection, rather than an assertion of the truth of *not-p*.
- 2. In order to assert, or reject, a subjective proposition, a speaker must be herself

acquainted with the subject of that proposition (in the relevant way). But that acquaintance is not presupposed by the standard norms of presupposition.

- 3. If an interlocutor lacks acquaintance with the subject, she cannot reject a subjective proposition p asserted by Speaker X, but neither is the common ground successfully updated with p. *Speaker X believes p* is successfully added to the common ground, though, and cannot be rejected, because a speaker's evidential basis for *p is true to the Speaker* cannot be contested.

All of these items, however, are true of all propositions, subjective and not, simply due to the Knowledge Norm and the felicity conditions on assertion — we assume that speakers believe the things that they assert, and speakers have grounds for rejecting what they reject. If Sue asserts, *Paris is the capital of England*, the common ground between her and Zoe is automatically updated with *Sue believes Paris is the capital of England* and she can later relay to Chloe, "Sue thought Paris was the capital of England — can you believe it?" even as Zoe, herself better geographically informed, would reject the update *Paris is the capital of England* from the common ground by replying "No, it's not!" or similar. So if that is true of all propositions, why does it seem particularly significant in the case of subjective ones?

The answer lies in two conditions. First, the acquaintance must be in the relevant way, which for most subjective propositions is experiential, as we saw above with the Acquaintance Inference. Second, for a subjective proposition, an interlocutor cannot argue with or challenge speaker's evidential basis for *p is true to the Speaker*. We saw this above with faulty disagreement. What muddled evidence Sue has for the idea that Paris is the capital of England is easily overridden by a mass of other evidence (world maps, gazetteers, the experience of being in Paris and noticing that travel in France and not England is involved and that France and England are difficult to confuse). Paris being the capital of England may be "true" to Sue, but we do not respect that truth the way we do with subjective propositions.

Let us rearrange those facts to get the following properties of subjective predicates:

(46) a. A subjective proposition requires a certain form of relevant direct knowledge (which form varies from proposition to proposition but is mostly experiential) for its assertion conditions.

b. If other speakers have the form of relevant direct knowledge that contradicts it, they are justified in rejecting the subjective proposition from updating the common ground.

c. If other speakers lack the direct knowledge, they cannot accept or reject the subjective proposition's update into the common ground.

d. That the speaker of the subjective proposition p believes p is automatically added to the common ground.

e. If the subjective proposition has a judge/involves the experience of someone other than the speaker — *Avocados are tasty to Bob,* as uttered by Sue — then it is treated just like an objective proposition and the knowledge base for it does not have to be direct.

So far, this looks exactly like the behavior of objective predicates and their assertion conditions, except for one factor — that *other speakers can have relevant direct knowledge that contradicts the asserted proposition.* This is not true of objective propositions. For objective propositions, if two speakers disagree, at least one speakers' evidence must be bad (they are mistaken, they were lied to, etc.). But we accept this for subjective predicates — and it may be culture- or context-dependent whether we accept this or not, as I explain much more deeply in Chapter 3.

This relevant direct knowledge, as we see, is mostly experiential. We talk about the experience of seeing a painting or a film, of hearing a piece of music, or tasting a food, or of interacting with a person (either in real space or via communication devices; I may form a subjective impression of an actor or celebrity I have never seen and who does not know that I exist). However, we also treat modals and moral judgements as subjective propositions —

disagreements about which team is likely to win the World Series or whether a particular purchase is necessary, or whether harassment or fraud is worse, are treated as faultless. In Chapter 4, I argue that there is an experiential dimension to modals as well, but for the purposes of this chapter, I am going to speak of the knowledge basis that causes one to believe a particular modal claim simply as relevant direct knowledge, or *evidence.*

And as Gunlogson & Carlson (2016) explain, this relevant direct knowledge is equally valid for both speakers for a subjective proposition — if it is adequate. If Sue has seen the film and Zoe has merely read reviews, Sue's appraisal of the film's merit is generally treated as better than Zoe's, but if Sue has very little coherent knowledge of films while the reviews Zoe read are of very high quality, the balance of power between their judgements may shift.

For relativists and contextualists, this relevant direct knowledge is introduced by a judge, which they have encoded as an index and as an argument, respectively. This allows either the context or the content part of the meaning of a subjective predicate to be different for the two speakers. But that makes the difference between subjective and objective predicates greater than it actually is; as we have seen, in their felicity conditions, they behave almost identically. I argue for a much simpler solution: this relevant evidence requirement is true of all propositions, subjective and not. It's just that for objective propositions, we expect that speakers will have the same evidence, and having it, will come to the same conclusions. For subjective propositions, we allow that different speakers have different evidence (mostly experiential) and thus will not come to the same conclusions. However, culture and circumstances play a role in whether we allow this; as I show in Chapter 3, experts may be more stringent as to whether different experiential evidence would be accepted or treated as faulty.

Thus, the distinction between objective and subjective propositions is that for subjective propositions, the evidence for different speakers is not the same, and if it is, it still cannot be compared.

The exact nature of what this evidence must be, I will discuss in the next chapter.

Even as the evidence cannot be compared, the uptake requirement for subjective proposi-

tions is a mechanism for speakers' judgements to converge while relying on different evidence. As argued by Barker (2002, 2012) and also Sundell (2010) and Plunkett & Sundell (2013), faultless disagreement often serves to establish community norms, negotiating, metalinguistically, where the standards and norms are for the subjective predicate to apply. Beltrama (2016) makes a similar observation, that subjective propositions have an inquisitive character, seeking out consensus (or not, if a conflict would be more attractive to audiences).

I argue (as I have argued before, in less detail, in (Vardomskaya, 2014)) that objective predicates are precisely those for which there is a reliable consensus of what evidence matters (to distinguish from a reliable consensus as to whether propositions containing them are true: we do not know whether there is life on other planets, but we know what it would take to prove it). For subjective predicates, and propositions containing them, there is no reliable consensus as to what evidence matters, and how much, and what does not.

I will conclude this section, and this chapter, by noting a potential distinction between two ways I've stated the problem:

(47) Speakers' use of subjective propositions may differ (but be equally valid)

a. because they have different (but equally valid) evidence; or

b. because they have the same perceptual evidence but come to different conclusions.

For many subjective propositions, such as those of taste, it is hard to tell what the difference between those is. Humans generally assume that avocados taste the same to everyone, the way colors and textures are the same for everyone; only recently has research (Knaapila *et al.* , 2012) shown that people with different genes actually have different taste perceptions, such as possession of a particular gene causing one to perceive cilantro as having a soapy taste. If Sue and Zoe both tasted cilantro but Sue concluded that it was delicious but Zoe concluded that it was disgusting, and we had assumed that their taste experiences were exactly the same, then we would believe that they have the same evidence lead them to

different conclusions. But if we know that Zoe has the soapy-cilantro gene, then we must interpret their taste experiences as being objectively completely different, just like we allow that a blind and a sighted person in the same situation would have different experiences.

In Chapter 3, I argue that in the case of (b), if two people have the same perceptual evidence, they categorize it differently (in both classificatory categories like *This is a malbec wine* and evaluative categories like *This is a good malbec wine*) and thus their mental representations of it are different.

This accounts for the Acquaintance Inference being required for predicates of personal taste, but I point out again that it is not required for moral propositions, as one can say:

(48) Arson is wrong, and so is robbery. Thankfully, I haven't experienced either.

I expand on this problem in Chapter 3 on evidence and in Chapter 4 on the phenomenology of experience, but for now I present a sketch of the idea. Experiences of injury are assumed to be more similar across human beings than experiences of taste; how I feel when having a broken leg is assumed to be pretty much like how you would feel in the same condition, and so pain reports do not require the Acquaintance Inference. It is much easier for a doctor to identify the illness when comparing a patient's description of the symptoms to other patients' descriptions, than it is for even a skilled oenologist to identify a wine given only a taster's description of the wine and comparison to other tasters' descriptions.

Thus, we can very easily put ourselves in someone else's shoes when it comes to having wrong done us, especially since we feel confident in extrapolating from other experiences: we may not have experienced armed robbery, but we've experienced having a toy stolen from us in kindergarten, and we may not have experienced arson but we have experienced the pain caused by our property being accidentally destroyed, and can extrapolate that this would feel very bad if done maliciously. Guided by this tendency towards empathy, societies have a much stronger incentive to come to consensus about moral judgements than they do about taste judgements, so that they can form laws and conventions of behavior. It is much more

in society's interest to ban arsenic in food, which does clear and evident harm, than it is in its interest to ban avocados. So in essence, the reason that we do not have an acquaintance principle pertaining to moral judgements is that we consider imagined evidence of moral harm to be enough.

2.7 Conclusion

In this chapter, I have looked at how subjective propositions behave in discourse, using four features of subjective propositions that have been identified by other writers: disagreement as rejection, the notion of faulty disagreement, the Acquaintance Inference, and the uptake requirement. I note, though, that these features end up being very similar to how objective propositions behave in discourse: objective propositions also require that the speaker has relevant knowledge in order to assert them and if one has much greater evidence than the other speaker, disagreement is faulty; objective propositions can also be accepted or rejected by the speaker's interlocutors, and objective propositions still automatically update the common ground with the fact that the speaker believes them. The main difference that I identified between objective and subjective predicates, then, is that speakers' evidence for the proposition can differ but be weighted equally, if it is driven by direct experience — the evidence cannot converge and be shared.

In the next chapter, I will look at what determines the evidence for subjective propositions, as well as what determines the evidence for objective propositions. I will derive from this what subjective propositions are.

CHAPTER 3

THE EVIDENTIALITY OF SUBJECTIVITY

"Do you know, Watson," said he, "that it is one of the curses of a mind with a turn like mine that I must look at everything with reference to my own special subject. You look at these scattered houses, and you are impressed by their beauty. I look at them, and the only thought which comes to me is a feeling of their isolation and of the impunity with which crime may be committed there."

— Sir Arthur Conan Doyle, "The Adventure of the Copper Beeches" (1892)

3.1 The problems we need to solve

In the previous two chapters, we looked at multiple characteristics that distinguish subjective propositions from objective ones. Here I provide a brief review of the features my model needs to account for. First, there are the well-known features of subjective propositions that I discussed in Chapter 1:

(1) Faultless disagreement:

Sue: Avocados are tasty.

Zoe: No, they're not! (neither is viewed as saying something false)

(2) Faultless disagreement in the comparative:

Sue: Avocados are tastier than carrots.

Zoe: No, carrots are tastier than avocados.

(3) Embedding under *find*:

Sue: I find avocados tasty/?vegetarian.

(4) Modification by *to X/for X* propositional phrases:

Sue: Avocados are tasty to me.

In Chapter 2, I added a number of pragmatic features of subjective propositions that have

been less documented, which I review below:

(5) Disagreement as rejection from the common ground rather than as contradiction in truth value (Khoo & Knobe, Forthcoming 2016):
Sue: Avocados are tasty, so let's have some with our toast.
Zoe: No, they're not! Let's not have any!

(6) Faulty disagreement patterns when one speaker has much weaker evidence than the other (Gunlogson & Carlson, 2016):

a. Sue: Avocados are tasty to me.
Zoe: #No, they're not.

b. Sue: I find avocados tasty.
Zoe: #No, you don't.

(7) The Acquaintance Inference (not a presupposition!) (Ninan, 2014):
Sue: Avocados are tasty, #though I've never tried one.

(8) The uptake requirement:

a. Sue: Avocados are tasty.
Zoe: [silence]
Sue: I take it you don't agree. / #Good, we're on the same page.

b. Sue: Avocados are tasty.
Chloe: Avocados are tasty; #I've never had one, but Sue said so.

We have seen in Chapter 2 that the pragmatic features of subjective propositions boil down to the following:

(9) a. A subjective proposition requires a certain form of relevant direct knowledge (which form varies from proposition to proposition but is mostly experiential) for its assertion conditions.

b. If other speakers have the form of relevant direct knowledge that contradicts it, they are justified in rejecting the subjective proposition from updating the common ground.

c. If other speakers lack the direct knowledge, they cannot accept or reject the subjective proposition's update into the common ground.

d. That the speaker of the subjective proposition p believes p is automatically added to the common ground.

e. If the subjective proposition has a judge/involves the experience of someone other than the speaker — *Avocados are tasty to Bob,* as uttered by Sue — then it is treated just like an objective proposition and the knowledge base for it does not have to be direct.

In Chapter 2, I noted that all of these look exactly like the behavior of objective predicates and their assertion conditions, except for one factor — that other speakers can have relevant direct knowledge that contradicts the asserted proposition, which knowledge is most often experiential. This relevant direct knowledge, which I am calling *evidence*, is the focus of this chapter.

In the next section, I will take a close look at what disagreement entails through the filter of evidence. Once I set up the terms for disagreement, the rest of the chapter is structured as follows. In Section 3.3, I relate the Acquaintance Inference and the pragmatics of subjectivity to direct evidentiality, and I split the topic of "different evidence" into differences in perception and differences in categorization. In Section 3.4, I discuss the growing information on human perception differences. In Section 3.5, I analyze differences in how humans categorize their perception and how that separates experts from amateurs. In Section 3.6, I put the pieces together: I explain how differences in perception and categorization account for the Acquaintance Inference and for the spectrum of subjectivity. I conclude in Section 3.7.

3.2 Breaking down disagreement

Let us look back on what a disagreement means, spelling it out. In all cases, we have two people, A and B, who review evidence e_A and e_B and form beliefs p_A and p_B, which they state as a proposition. To simplify the discussion, I will use p for both the proposition the speaker utters and the belief the speaker holds, thus assuming that the speaker is not lying and she is not misusing words (e.g. to give a famous example, claiming that she has "arthritis of the thigh," not knowing that the standard definition of arthritis is as a joint condition).

We will assume, as a first step, *evidence for p* to mean *a set of one or more propositions that increase the probability that p is true*, which is the definition used by McCready & Ogata (2007); Krawczyk (2012); McCready (2015).

For propositions, we assume that the evidence e_A is equivalently reliable to the evidence e_B for any two speakers A and B in order for them to be felicitous in asserting p_A and p_B. "Equivalently reliable" does not mean "the same" — for example, A could know that Paris is the capital of France through having been there and seen the government buildings, while B could know it through consulting an atlas. But we assume that an atlas is good enough, and therefore A and B's evidence is equivalently reliable.

We also assume that if p_A and p_B have the same uttered form, e.g. *Paris is the capital of France*, they reflect the same belief we can call p. There are numerous possible ways this can fail, such as homophones (*John saw Mary at the bank*) or structural ambiguities (*I saw a woman with a spyglass*), but we have pragmatic conditions such as follow-up questions established for resolving this, as Kennedy (2011) explains.

What does this mean for subjective propositions? Again, e_A and e_B should be equivalently reliable — which for predicates of personal taste ties to direct experience, and I explain why in the following sections — but need not be the same. If A tasted avocados ten years ago in Israel, and B tasted them last week in Mexico, we still assume that they have both fulfilled the felicity conditions to assert *Avocados are (not) tasty*. Whereas if B has never

tasted avocados, she simply does not fulfill the felicity conditions to assert it.[1]

However, for subjective propositions, we have A and B with equivalently reliable evidence e_A and e_B, that leads them to felicitously assert p_A and p_B, which contradict each other. That means that p_A and p_B, depending on your model of truth, may not both be true in the same world. Here I can show that different interpretations of what faultlessness means in disagreement would target different aspects of that statement (note that the following list does not mean that all items on it are orthogonal; some depend on each other).

(10) The requirements for faultless disagreement:

a. Speaker A felicitously asserts p_A from evidence e_A, and Speaker B felicitously asserts p_B from evidence e_B.

b. p_A and p_B both have the same references.

c. p_A and p_B are contradictory/not truth-compatible/cannot both update the common ground.

d. e_A and e_B are reliable: e_A should for all reasonable people lead to the belief p_A and e_B should for all reasonable people lead to the belief p_B.

e. e_A and e_B are considered equally reasonable.

If we interpret faultless disagreement as objectivism — B (without loss of generality) is simply lying or misinformed — then we are addressing (a); B should not have felicitously asserted p_B. If we interpret faultless disagreement as a misunderstanding, in the style of Stojanovic (2007), we are addressing (b): p_A and p_B do not actually have the same references, e.g. either A and B were talking about different things, or they have different meanings for words like *tasty*; this is what contextualist perspectives address. If we adopt a relativist perspective, we address (c): we relativize truth so that p_A and p_B are not contradictory.

1. The problem of why tasting just one avocado, or even one bite of an avocado, is enough to conclude *Avocados are (not) tasty*, even though it's possible that the avocado bite was atypical, is known as the *generality problem* in epistemology. Addressing it is beyond the scope of this dissertation, but see (Comesaña, 2006).

If we interpret faultless disagreement as based on flawed evidence, e.g. asserting *X is tasty* when we have not tasted X, we are addressing (d): If e_A is reliable, e_B is not and should not lead to the belief p_B for all reasonable people.

My position is that subjective propositions are those for which (e) is unprovable: we do not know that e_A and e_B are equivalent, but we treat them both as reliable.

As I argued in Chapter 2, I will follow a model of (c) in which p_A and p_B cannot both update the common ground to interpret disagreement; I will rely on the unprovability of the equality of e_A and e_B to determine faultlessness; even though I present some ways in which e_A and e_B are not equal, I will show that the pragmatic requirements of a common ground require us to treat them as if they are. Thus I am arguing for a model that is more expressivist than a truth-conditional model. Because e_A and e_B are equally reasonable, we cannot have an exclusive model of truth. Thus we are expressing something else — a pragmatic position — when we assert p_A or p_B.

3.3 Reliability and evidentiality

Let us look first at the reliability condition: that if A and B state p_A and p_B, A's evidential basis e_A for p_A and B's evidential basis e_B for p_B should be treated as equally reliable. Otherwise, we have cases of faulty disagreement. In the straightforward case where p_A and p_B are conventionally objective statements, e.g. *This table is wooden* or *Paris is the capital of France*, we know that if two people disagree, assuming both believe that their statements are true, one person's evidence is not reliable. Maybe A had been lied to by slick salespeople that the table is genuine Brazilian rosewood, when it is actually just plastic lumber. Maybe B had dozed off during geography classes and misremembers the gazetteer.

The issue of how evidence e would lead to belief p for all reasonable people is the main concern of the philosophical subfield of epistemology.[2] Epistemology concerns itself with how we may know things: how facts about the world get translated to (hopefully true) beliefs

2. I thank Anubav Vasudevan for useful discussion that helped me consolidate my ideas in this chapter.

in the brain. It is clearly not within the scope of this dissertation to be an epistemological dissertation; thus I will simplify the path between facts in the world and beliefs in the brain to two stages — humans' *perceptions* with their senses and their *categorizations* — classifying the perceptions as belonging to this or that category (both substantive: this is a beer; and evaluative: this is a good beer).

We know that subjective predicates have a particular epistemological feature which is the Acquaintance Inference: that it is infelicitous to assert either a subjective predicate or its negation without firsthand knowledge, as in the following examples.

(11) a. Avocados are tasty/are not tasty #even though I've never tried one.

b. This movie is terrific/awful #because the critics told me so, but I haven't seen it myself.

c. This book is great/terrible #though I haven't read it, but all books by this author are great/terrible.

Explicit evidentials make this particularly clear in Japanese, where predicates of personal taste have a first-person restriction (Bylinina, 2013) as shown in (a). Second- and third-person uses must be modified with an evidential, as shown in (b):

(12) a. watasi-ni-(totte)-wa / *John-ni-(totte)-wa kono keikii-wa oisii
I-DAT-TOTTE-TOP / John-DAT-TOTTE-TOP this cake-TOP tasty
"This cake is tasty to me / to John."

b. John-ni-(totte)-wa kono keikii-wa oisii noda / ni tiganai
John-DAT-TOTTE-TOP this cake-TOP tasty EVID / there's.no.mistake

"The cake must be tasty to John."

These examples illustrate that for predicates of personal taste, only direct evidence, the kind that involves perception, is reliable. Interestingly, hearsay and indirect evidence may be sufficient for *belief* in the validity of a subjective predicate, as the following felicitous

examples show:

(13) a. I believe avocados are tasty/are not tasty, even though I've never tried one.

b. I believe this movie to be terrific/awful because the critics told me so, but I haven't seen it myself.

c. I believe this book to be great/terrible, though I haven't read it, but all books by this author are great/terrible.

The embedded propositions, however, seem to be interpreted exocentrically — *I believe avocados are tasty to Bob* seems a reasonable if stilted assertion if one is pondering what to serve Bob for dinner — or as generics, and thus not treated as subjective per se. We can look at belief-embeddings with *to me* and *find*-embeddings with *I* as the subject to confirm that there isn't an exocentric usage, and we discover that we can affirm beliefs as long as they are about modal uses of subjective propositions:

(14) a. I believe avocados would be/#are tasty to me, even though I've never tried one.

b. I believe I #(will) find this movie to be terrific/awful because the critics told me so, but I haven't seen it myself.

c. I believe I #(will) find this book to be great/terrible, though I haven't read it, but all books by this author are great/terrible.

Such sentences are perfectly felicitous with a modal, but not without it — thus, one can hold the belief that in every possible world where one tastes avocados, or sees the movie or book, the avocados are tasty or the movie or book is great.

So indirect evidence is sufficient for *belief*, or modalized belief, but not for assertion. The experiential component is transparent and presupposed for *find* embeddings (see Chapter 4 for much more) and *to me*-modified verbs, but these do not give rise to faultless disagreement.

(15) Sue: I find avocados tasty.

Zoe: #No, you don't./#That's not true./?You're mistaken.

Similar expressions that have an experiential presupposition, as Pearson (2013) points out, are perception verbs like *look*, as shown by the following example:

(16) Chloe looked tired today, #but I haven't seen her.

However, those do not seem to give rise to faultless disagreement intuitions:

(17) Sue: Chloe looked tired today.
Zoe: No, she didn't.
(Objective disagreement: we're assuming that Sue's or Zoe's perceptions are off.)

A particularly curious example of experiential predicates that do not give rise to faultless disagreement is pain reports. As mentioned by Korotkova (2016), they are subjective, and therefore non-challengeable. People have differing experiences of pain, and there are people who are completely insensitive to it. However, faultless disagreements about pain seem very odd:

(18) Sue: Being hit with a hammer is painful.
Zoe: #No, it's not!

Even if Zoe happens to have a neurological condition causing insensitivity to pain, so being hit with a hammer really is not painful for her, this is still a very odd dialogue.

Similarly, although pain is a subjective experience, we do not have an Acquaintance Inference arising on the bare assertion of something being painful. A biological male speaker who cannot possibly ever have the experience of giving birth can still felicitously assert *Giving birth without an epidural is painful.* This, however, may be an difference in degree rather than kind: it may be generally accepted that certain experiences like giving birth or

stubbing one's toe are painful, while others that are not commonly accepted to be painful but are for some people with some conditions (swallowing, sitting in a split, etc.) are treated as subjective, with assertion requiring the Acquaintance Inference.

Subjective expressions that give rise to faultless disagreement but do not have an experiential component and do not seem to require the Acquaintance Inference include vague expressions such as *tall* and *rich* and moral expressions such as *wrong*, as in the following examples:

(19) a. Feynman was tall, but I haven't actually seen him; I just know he was over 6'.

b. Arson and robbery are wrong, but I thankfully haven't experienced either.

These expressions do require a knowledge norm; one cannot assert that Feynman is tall without knowing anything about Feynman's height, and one cannot assert that arson and robbery are wrong without knowing what they are. However, the perception of faultless disagreement is still there.

Bare assertions of predicates of personal taste appear to be in the middle: they have an experiential component, but it is not a presupposition, and they license faultless disagreement, but only if both participants have the relevant experience. We find that the example about winning the debate, which has faultless disagreement in the case of a presidential debate, does entail the Acquaintance Inference in that case:

(20) a. Sue: Obama won the debate last night [in the 2012 US election], #but I didn't see the debate.

Zoe: No, Romney won it, #but I didn't see the debate. (Faultless disagreement)

b. Sue: Sharon won the debate last night [at the World Universities Debating Championships], but I didn't see the debate.

Zoe: No, Wallis won it, but I didn't see it either. (Objective disagreement).

Notably, this is not a presupposition either, as I will analogize with Ninan's (2014) *The lobster rolls are tasty* example, which shows that putting *tasty* in the antecedent of a conditional does not require that the subject had tasted it, unlike well-known presuppositions like *stop*:

(21) a. If the lobster rolls are tasty, we'll have some. But we haven't tried them yet.

b. If Obama won the debate last night, I will vote for him, but I didn't see the debate.

Thus the evidential requirements for predicates of personal taste and evaluative predicates seem stricter than for any other class.

The issue of how evidentiality interacts with reliability has been explored extensively by Krawczyk (2012) and McCready (2015), but both of them overwhelmingly concern themselves with evidential marking on objective statements such as *The keys are in the car* and with indirect and hearsay evidentials rather than direct ones. Viewing evidentiality through an epistemological lens, indirect and hearsay evidentials do pose more intriguing puzzles. Krawczyk does spend some time considering the position first put forward by Faller (2002) about Quechua evidentials where Faller argues that when a speaker uses a direct evidential (or leaves the sentence unmarked in languages where the direct evidential is the unmarked form), he or she has the *best possible grounds* for asserting a proposition.[3]

We thus see that the Acquaintance Inference forces the best possible grounds, in fact the only possible grounds, for asserting a subjective proposition to be those of direct sensory

3. I will note here that I am not the only researcher looking at the relationship between subjectivity and evidentiality. In her dissertation, Korotkova (2016) concludes that subjectivity drives evidentiality: she argues that evidentials describe a first-person subjective experience and are associated with a mental state that describes the acquisition of evidence for p. Although Korotkova presents very interesting data, and I agree with her arguments that the long-assumed split between epistemic and illocutionary evidentials in the evidentiality literature should not be maintained, I do not agree with her that it is subjective experience that should be treated as basic and evidentiality treated as secondary. I say it is the other way around. I seek to prove that it is evidential differences that make the difference between subjective and objective propositions — it's not that we are driven by our sense of self to gather evidence, it's that we demarcate our self with evidence that only we can have — because that explains how the same word can be objective in one context and subjective in another, which Korotkova's model cannot do. I discuss this more extensively in Chapter 5, where I present several arguments against Korotkova's conclusions.

perception. That would be the case only where it is reasonable to expect different people's grounds to assert it to differ so much that only one's own perception and judgement can matter. So we have that the subjectivity in propositions is rooted in speaker differences.

Thus differences in beliefs are either differences in *perception* of the evidence, or differences in *categorization* of the evidence. Differences in perception, I will categorize as *defective taste* and discuss in Section 3.4. Differences in categorization, I will categorize as *bad taste* following Lasersohn (2017) and discuss in Section 3.5.

A primary distinction between these is that differences in perception would generally require medical intervention to change, if they can be changed at all. Differences in categorization, however, can be changed through training, education and persuasion. For example, a difference in how two people perceive music may be due to the fact that one is hard of hearing, which is a perception difference, whereas it can also be due to the fact that one has better relative pitch and is better at identifying out-of-tune intervals and chords, which is a categorization difference (relative pitch can be taught; absolute pitch, on the other hand, would be treated as a perception difference). Similarly, a difference in how two people judge a cilantro-containing dish may be due to one person bearing the gene that causes cilantro to have a soapy taste; this is a perception difference. It can also be due to one person not being used to this type of cuisine, and not sure what to think of it; this is a categorization difference.

3.4 Perception, sensory issues and defective taste

Intuitively, we treat 'subjective' predicates as precisely those where testimony is trustworthy only if the testifier herself is trustworthy. That is what is usually implied when we say, "Well, that's subjective" in discourse, in response to someone's assertion of a taste or an opinion. We are not calling the speaker a liar, but we are asserting that we will not accept their testimony into the common ground. How do we learn what is subjective and should be treated this way?

One reason for us to learn this is that taste and other perceptions can differ. Recent science has confirmed that bearers of certain genes will predictably perceive flavors differently, the most famous being a gene that makes cilantro taste soapy to the bearer (Knaapila *et al.* , 2012).

A person who has lost the ability to hear high pitches would have a different experience of the Chicago Jazz Festival than someone who still has full-spectrum hearing, and a deaf person would have a different experience still (the Chicago Jazz Festival provides ASL interpretation on the main stage, so evidently there is demand for it). A color-blind person would have a different experience of landscapes and paintings from one who has full vision, and a tetrachromat (someone with four types of color receptor cones in the eye instead of three) would have a different experience as well. We recognize that these are differences of perception due to anatomical and physiological variation.

However, color and sound are not generally treated as subjective (see Kennedy & McNally (2010)). I argue that a key distinction in whether a property is treated as objective, and differences in perceiving it are treated as sensory differences, is the use of computers and other devices to measure it.[4] A pocket tuner or smartphone app linked to a microphone can measure the frequency and thus calculate the pitch of a sound. A color picker can identify colors by either the hue, saturation and lightness values, or the amount of cyan, magenta, yellow and black. Note that for the latter, colors are still vague: with the help of tools we may identify a certain color as having a hue of 34, a saturation of 94, and a lightness of 45, but that does not help us decide whether this counts as 'yellow' or 'orange' (these are the values for the color known as 'School Bus Yellow'). But we treat the value as determined by a computer as objective, and thus "darker" or "lighter" or "sharper" or "flatter" or "out of tune" are objective properties. Difficulties in perceiving them despite training are considered sensory perception difficulties (color-blindness, tone-deafness).

4. A book that extensively explores this topic is *Objectivity* by Lorraine J. Daston and Peter Gailson, MIT Press, 2007, but I was made aware of it too late to incorporate it into the body of this dissertation.

We do not yet have devices that can similarly assign an objective value to taste, although we can measure the amount of sugar, salt, flavor esters, and other components in a food. However, we know that a simple measurement of the amount of sugar in a substance is only partly correlated with how sweet the substance tastes to people: for example, the amount of salt in the substance would counter it, which is why salted caramel tastes less sweet than pure caramel. It is a question for phenomenologists and psychophysicists whether codifying tastes into a device is impossible, or whether we simply do not have current technology to do it. But we know that taste results from the interaction of many different flavors and odors with the sensory receptors we have. And we've learned that those differ.

Recent investigations have suggested that our perceptions of smells and tastes vary widely among human beings. The phenomenon of *specific anosmia*, where the patient cannot detect the odor of a particular substance, is actually very widespread, suggesting that most of us have an olfactory "blind spot" where we cannot smell a certain chemical, due to genetics or lack of training.

This has been previously known to be genetically linked for asparagus metabolites, which cause urine to have a sulphur-like smell after eating asparagus. About 40% of humans can detect this smell[5] and the ability to perceive it is associated with about 871 genetic variations on chromosome 1 (Markt *et al.* , 2016). Croy *et al.* (2015) report that cases of people unable to smell vanilla, though detecting all other odors, had been known since 1893; about 0.1% of people cannot detect butyl mercaptan which is the characteristic smell of skunk spray; 2-3% cannot detect isovaleric acid (a smell that is a major component of foot odor but can add characteristic notes to wine in small quantities) while 7-12% cannot smell the larger molecules of pentadecalactone (a synthetic musk). They report,

(22) "For nearly all tested odors, some specific anosmics could be identified. Thus, each odor needed to be tested in numerous subjects in order to capture the relatively

5. Including baseball great Babe Ruth, who is said to have rejected asparagus salad at a society function by uttering some version of "No, thank you, it makes my urine smell."

> small prevalence rate – apparently, as long as enough subjects were tested, there would be always some specific anosmia to any odor. [...] If indeed almost all smells are possible targets of specific anosmias, and if the overlap of specific anosmia is rather low, it seems plausible that everybody should have a specific anosmia to some odor. As a consequence, while the frequency of anosmia to any single odor is rather small, the general phenomenon of specific anosmias seems quite normal: Specific anosmia might be considered to be the rule, rather than the exception of olfactory processing." (Croy *et al.* (2015), p. 181)

Based on testing 1600 participants on progressively stronger dilutions of 20 characteristic scents (including chemicals associated with musky, flowery, minty, citrus, honey and sandalwood scents), Croy *et al.* found that for all but four of the odors, there was at least one person who fulfilled the criteria for specific anosmia, and the probability of having anosmia to at least one of the 20 odors was over 51%. They extrapolate that "the prevalence for specific anosmia to at least one out of 50 given odors amounts to 86.9% and 98.3% if 100 odors are taken into account."

This is an extremely interesting result, with major ramifications for comparing tastes. It is well-known that smell is a major component of how foods taste, so it is natural that a person with specific anosmia for a chemical contained in a certain food would have a different experience of eating that food than a person who can perceive the scent. For example, 11 people in Croy et al's study could not detect the smell of I-carvone, which is minty. These people would almost certainly have a different experience of mint tea or mint candies.

In humans, smells are determined by over 400 receptors (Croy *et al.* , 2015), while color is determined by only three (possibly four for certain rare people). This allows for much more varied differences between different people's smells and tastes (leading to specific anosmia and ageusia — taste-blindness) than different people's color perceptions (color-blindness). This biological fact explains why taste, largely determined by smell, is generally treated

as much more subjective than sight. People can fairly easily compare perception of red and green and determine whether they see the same thing, whereas a person who does not perceive a chemical like I-carvone, pentadecalactone, or asparagus metabolites may go her whole life without realizing that there is a gap in her perception.

If specific anosmia is so prevalent, while most of us do not know that we have it, that means that many if not most disagreements about taste and smell *are* actually based on differing perceptions, already known for the taste of cilantro. We must also consider that there is the flip condition of specific anosmia, specific hyperosmia, which means extra sensitivity to a particular scent. General hyperosmia is commonly associated with pregnancy symptoms according to anecdotal data and folklore, although scientific evidence for this is sparse and inconclusive; there is some debated evidence, however, that hyperosmia can occur in individuals with temporal lobe epilepsy or migraines (Cameron, 2014). Genetically-linked specific hyperosmia has been documented for isovaleric acid (Menashe *et al.* , 2007) and several other chemicals, although that is also under-investigated. In general, we can conclude that people with specific hyperosmia would also have differing perceptions of certain foods, just as people with specific anosmia would.

Which means that we need to consider whether the following is faultless:

(23) Sue (who does not have the soapy-cilantro gene): Cilantro is delicious.
Zoe (who does have the soapy-cilantro gene): No, cilantro is disgusting!

Knowing what we do now, we would judge Sue and Zoe to have different perceptions. This leads to a valid disagreement but one that is not about truth, but about updating the common ground. We cannot judge an objective fact about whether cilantro tastes delicious or disgusting *tout court*, because that relies on both facts that can be established about the food — the chemicals in cilantro that are perceived by the genetically-coded taste receptors — and each judge's smell taste receptors.

This leads us to an interesting conundrum. We generally assume that when two speakers

disagree, both have reliable evidence. As I laid out in the previous section, e_A and e_B are reliable: e_A should for all reasonable people lead to the belief p_A and e_A should for all reasonable people lead to the belief p_A. The evidence would be based on perceiving. We would not accept that a speaker can assert "Avocados are tasty" without having tasted avocados. But it is not general knowledge that one may taste avocados *without being able to perceive some key component of their smell and taste*. Without that key component, does e_A still count as reliable?

I believe that we are left with saying "Yes" by the requirement that e_A should reasonably lead to p_A. A normal person, tasting cilantro or any other plant and finding that it tastes of soap, *should* reasonably find it disgusting (although they may not, for whatever reason; people can like soapy tastes). The difference in taste here is actually covered by the assumption that the reliable evidence is equal for both. We do not judge that a color-blind person has equivalent evidence as a person with full color vision in looking at a landscape, a flag or a painting. However, mainstream culture is far more aware of color-blindness than of specific anosmias. It may be that as we become more aware of genetic differences in taste and smell perception, we will see fewer "faultless" disagreements about the taste of foods and more conversations like the following:

(24) a. Sue: Cilantro is tasty.

Zoe: I have the soapy-cilantro gene, so cilantro is not tasty *to me*, but I allow it's tasty to other people.

b. Zoe: This wine is delicious.

Sue: I'm hypersensitive to isovaleric acid, so it's not delicious to me, it tastes a bit like sweaty socks. But okay, not everyone has that problem, so it's probably tasty to you.

Differences in perception leading to "defective taste" can explain "faultless" disagreement in a lot of cases, but they cannot explain everything. Depending on the gene that codes for

perception of cilantro, or asparagus metabolites, one's taste and smell experience of these should remain constant from birth to death. Barring a brain injury or an illness that affects the receptors, there should be no change in one's taste, and no reason for a speaker to retract an earlier claim. However, tastes do change over time and as MacFarlane (2014) emphasizes, people do feel they have to faultlessly retract earlier taste claims. Tasters can take courses or training and gain more refined palates, whereas a negative experience with a food can color one's perception of its taste forever after. Even specific anosmia can be overcome, as I discuss in the next section. And we cannot generalize to all subjective predicates: genetic differences in activated smell receptors cannot explain the difference between an art critic or figure skating judge and an amateur viewer. Moreover, as food researcher Martin Yeomans remarked, "It is striking how little genetics predisposes humans to like or dislike food flavors" (quoted in (Vanderbilt, 2016), p. 20). Differences in perception do not account for differences in preferences even initially.

These all suggest that what differentiates people with different tastes is not just (or even mainly) differing perceptions, but differing ways of analyzing what they perceive, which may change with time and training, and which forms the topic of the next section.

3.5 Categorization, expertise and bad taste

Among the scholars working on subjectivity, MacFarlane (2014) is perhaps the most scrupulous in looking at how tastes change and why we may need to retract statements we have previously made about taste. This is summarized in his "fish sticks" example:

(25) When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I've changed my mind about the tastiness of fish sticks. So, if someone said, "But you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that

would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I didn't mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane (2014), pp. 13-14)

MacFarlane uses this example to argue for the relativist view, particularly for the assessment-sensitive relativism that he favors (as I discussed in Chapter 1) and against the contextualist view. He explains in his own words how the example of changing one's mind about fish sticks is an application of a Retraction Rule that he posits is necessary for discourse:

(26) Retraction Rule: An agent in context c_2 is required to retract an (unretracted) assertion of p made at c_1 if p is not true as used at c_1 and assessed from c_2.

(27) An example will help show how. Let c_1 be a context centered on ten-year-old Joey, who loves fish sticks. According to both R [relativist models that include the Retraction Rule] and C [contextualist models], the proposition that fish sticks are tasty is true as used at and assessed from c_1. So the Reflexive Truth Rule tells us that Joey is permitted to assert that fish sticks are tasty. Let us suppose that he does. Now consider another context c_2 centered on Joey, ten years later. As a twenty-year old, Joey no longer likes the taste of fish sticks. Here R and C diverge. According to R, the proposition that fish sticks are tasty is false as used at c_1 and assessed from c_2, so by the Retraction Rule, Joey is now required to retract his earlier assertion. According to C, by contrast, the proposition that fish sticks are tasty is true as used at c_1 and assessed from c_2, and Joey need not retract. The practical difference between C and R lies in what they imply about Joey's obligation to retract his earlier assertion. (MacFarlane (2014), pp. 109-110)

The question that most concerns this dissertation is what has changed in Joey's mind or perceptions that made him change his mind about fish sticks. We can assume that, barring

injury, his sensory organs have not radically changed in ten years. Many more adults than can be accounted for by illness or injury discover that they lose their taste for childhood loves like fish sticks, chicken fingers, high-sugar milk chocolate or gummi worms. The way MacFarlane himself explains the change is that he has been exposed to a broader range of tastes. Thus, in comparing the taste of fish sticks to the taste of high-quality seafood, he is forced to categorize them as inferior, even though they may taste exactly the same as they did in his youth.

This is the issue of categorization, and as I suggested in the previous section, this, rather than perception, is what distinguishes the expert from the amateur.

The difference between two speakers who have the same sensory data but have different ways of analyzing it is illustrated in this example from Comesaña (2006), discussing the generality problem in epistemology. Comesaña suggests the scenario of two people looking at a many-sided polygon. Ms. Accurate has a remarkably excellent number sense, what we normally use to instantly gauge the number of objects in front of us without individually counting each one. For most of us, this only works for up to five objects or so, while Ms. Accurate has it developed for much higher numbers, so she can look at the polygon and tell at once that it has 386 sides. Mr. Sloppy, on the other hand, is a normal person like most of us, with a number sense that only works up to five objects or so; thus when he looked at the polygon, he just randomly guessed that the number of sides is 386, and got the correct answer by sheer coincidence. Comesaña remarks:

> If Ms. Accurate's has the same experiential input as Mr. Sloppy, then her superior capacity can only be explained in terms of how she gets from this input to the belief. But, if that is the case, then we can just build this way of getting from the input to the belief into the description of the type. The relevant type-schema will then no longer be *producing a belief that p based on evidence E*, but rather *producing a belief that p based on evidence E by using method M*. Applied to our case, this would yield the processes *producing a belief that the figure in*

front of you has 386 sides based on the fact that it looks to you as if it has a lot of sides by using the method of discriminating exactly how many sides it appears to have (for Ms. Accurate), and *producing a belief that the figure in front of you has 386 sides based on the fact that it looks to you as if it has a lot of sides by using the method of venturing a wild guess as to how many sides it has* (for Mr. Sloppy). Again, the reliability that is plausibly associated with these two methods coheres nicely with our intuitions about the case.(Comesaña (2006) p. 42, italics original)

Comesaña's evidence E and method M map fairly well, as I see it, to my perception-categorization distinction. To take a more plausible case than Comesaña's somewhat outlandish one, consider three people hearing the same short passage of recorded saxophone music.[6] Alice likes music in a casual way, but has no musical training; she listens to what is on the radio. Beatrice has extensive training in classical piano, theory and harmony and has absolute pitch, but knows very little about jazz. Clarice spent many years in the jazz world, although she does not have absolute pitch. They utter the following:

(28) Alice: This is a nice piece of music. I think it's a saxophone playing. Sounds like jazz, but that may be because of the saxophone.

Beatrice: This is a nice piece of music. This is an F major ascending broken triad played on a tenor sax.

Clarice: This is a nice piece of music. This is Sonny Rollins playing. I know the distinctive tone he had on the saxophone. If we had a longer passage, I'll probably give you the track and the album.

Alice, Beatrice and Clarice all have the same perceptions of the sensory data — they heard the same recording, perhaps even in the same condition. But their differences in experience

6. This a fictionalized version of several real conversations I have had with jazz musicians and musicologists.

(and possibly natural ability, if one believes absolute pitch to be an inborn talent) lead them to very different beliefs. Both Beatrice and Clarice have beliefs that are a superset of Alice's — they agree that the music is by a saxophone and is pleasant to the ear. Beatrice and Clarice's beliefs do not contradict each other. But Beatrice has experience and training (as well as perception) that Clarice doesn't, and Clarice has experience and training that Beatrice doesn't, and that allows them to sort music into categories that Alice has no access to.

What differentiates Beatrice's and Clarice's impressions from those of taste, though, is that they are matters of fact: everyone else with absolute pitch should be able to identify that it is an F major chord, as could a digital tuner. Whether it is Sonny Rollins or not should be confirmed by the liner notes, and most other people as familiar with his oeuvre as Clarice should also be able to identify Rollins's distinctive tone. Is that true of prototypically subjective opinions, where "there is no fact of the matter"? How can people learn these?

Recall Croy *et al.* (2015) from the previous section, who were able to extrapolate that the probability an individual having an odor they are not able to detect approaches 1 as the number of odors approaches 100. Even more interestingly, in a follow-up study, Croy and colleagues selected participants who were anosmic to specific scents, gave them bottles of the scent they could not detect, and instructed them to sniff it twice a day for two months, thus training their ability to detect it. All 25 participants improved in their ability to detect the scent, showing that their specific anosmia was not due to congenital inability. Which means that taste and even smell can be taught.

The popularity of wine-tasting and other tasting courses adds further support the idea that taste can be taught. The perception had been there all along, presumably, but speakers did not know how to analyze it. Valentin *et al.* (2007), also cited in (Vanderbilt, 2016), found that the difference between beer experts and amateurs lay not in *detecting* more odor components of beer, but in *recognizing* more components of beer. Experts and novices did equally well (or equally poorly) at detecting novel odors, showing that it is not the experts'

finer perception that makes the difference.

What differentiates experts from ordinary people is that their tastes also tend to converge. Watching the judging of the Great American Beer Festival, Vanderbilt (2016) remarks, "Experts are people who have the same opinions as other experts" (p. 214). A study of the reviews on RateBeer.com similarly concluded that one of the ways you can distinguish novice reviewers from expert reviewers was that "we find that experts' ratings are easier to predict than those of beginners; and we find that experienced users agree more closely when reviewing the same products" (McAuley & Leskovec (2013) p. 2). The training of judges for judged sports such as figure skating or gymnastics also aims for a goal of fair evaluation, that an athlete should expect to get the same score for her performance even if she gets different judges — that is, being as objective as possible.

Explicit categorization helps with this. Gymnastics and figure skating have published and publicly available codes of points — lists of admissible moves the athletes could make to fulfill program requirements, the standard form of execution of each move, possible faults in the execution of the move, and the point value deducted for each fault. As Chris Grabowecky, a gymnastics judge at the 2004 Olympics, explained, "When you are normally judging, you'd script what element they did and then what the deduction was so if somebody asked, you could say, 'On this element, I took this amount.' " (Myers (2016), p. 82.)

The difference between how experts view a performance and how amateurs view it, sorting what they see into different categories, is clearly illustrated in the most infamous judging case of Grabowecky's career, at the 2004 Olympics, when celebrated Russian gymnast Alexei Nemov was given a score (by a panel of judges including Grabowecky) that put him out of medal contention in the high bar event, causing the audience to boo so much that it delayed the competition for ten minutes.[7] "He [Nemov]," said Grabowecky, "had errors in his routine that the public doesn't recognize as errors because they're not judges. But he

7. Grabowecky re-evaluated Nemov's score at the insistence of the head judge and raised it by one tenth, which still left the gymnast out of medal contention. The controversy was one of the many factors that led to the sport of gymnastics switching from scores out of 10 to an open-ended score (Myers, 2016).

had those beautiful release skills and they go, 'Wow, that guy should win.' But they don't see that he's catching his release close, bending his knees here, taking a step on the landing." (Myers (2016), p. 86.) The errors were there to be perceived all along, but amateurs did not recognize them as important, while judges did.

Similarly to gymnastics judges (although without Olympic medals on the line), official panels of taste testers list the elements: they identify odors and flavor profiles from a list and assign a numerical score to a number of values such as "amplitude" (how well-balanced the flavors are) (Gladwell, 2004), and these score sheets can be compared; the goal of a tasters' panel is to converge on the profile, tasters compensating for each others' specific anosmias and other differences. As the website of Aroxa, a company that manufactures training flavor samplers for taste panels[8], explains it,

> "an experienced sensory assessor can identify and name (using widely understood, unambiguous terms) as many as 100 different characteristics in a given product type. Furthermore, they can give an estimate of the intensity of each flavor note. To the untrained sensory assessor, this looks an impossible task. However, to those who are trained in professional sensory assessment of product quality, it is relatively straightforward...When it comes to perception of odors, tastes and mouthfeel characteristics, each of us lives in a world of our own. There is no such thing as the 'average assessor'. Rather than fight these differences in sensitivity among individuals it is best to embrace them. Professional tasting is 'crowdsourcing' in action. It is the average response of a group of professional tasters that we're interested in."

Note the use of "widely understood, unambiguous terms": I argue that what is essential to categorization working to make judgements more objective is having enough vocabulary. Valentin *et al.* (2007) emphasize this for assessing beer; describing the level of training

8. http://www.aroxa.com/about-train/, accessed June 14, 2017.

their beer experts had, they note, "They had developed a consensual vocabulary, along with the ability to identify and quantify the majority of the basic flavors commonly found in beers...Previous studies on odor memory showed that associating a label to an odor improves its retrieval." Having the vocabulary, their study and studies they cite show, helps in retrieving the necessary perceptual categories. But also, I add, having explicit vocabulary helps in sharing the required categories with others; this allows expert judgements to converge.

This is also supported by another linguistic argument from Majid & Burenhult (2013), who compare English to the Jahai language spoken on the Malay Peninsula. They describe the rich vocabulary the Jahai have for describing odors:

> "The Jahai, a group of nomadic hunter-gatherers in the mountain rainforests along the border between Peninsular Malaysia and Thailand, have a lexicon of over a dozen verbs of olfaction that are used to describe a wide array of odors. These are 'basic' smell words (cf. Berlin & Kay, 1969): they are monolexemic and psychologically salient; they are not source-descriptors, nor are they restricted to a narrow class of objects (Burenhult & Majid, 2011). For example, the term *ltpcit* is used to describe the smell of various flowers and ripe fruit, including intense smell of durian, perfume, soap, *Aquillaria* wood, and bearcat (*Arctictis binturong*, which, according to Wikipedia, smells like popcorn)."

Majid and Burenhult tested speakers of English and of Jahai on describing odors and colors, measuring for rate of agreement (what key terms were repeated across descriptions) and length of description, with concise descriptions indicating the speaker knew how to describe it at once, while long descriptions meaning the speaker is struggling for words. They concluded, "Whereas English speakers grappled to find words for odors, Jahai speakers could name odors with the same conciseness and level of agreement as colors." The Jahai language (and, I argue, the subset of the English language used by experts in odor and taste) helps agreement among speakers about their perceptions, because they can categorize more easily.

With tasters' panels, we have sufficiently fine categorization compensating for deficiencies or differences in perception. And we find that as differences in perception serve to bring taste judgements farther apart, better and more explicit categorization serve to bring them more closely together. Taste deficiencies lead to subjectivity, and good taste leads to objectivity — or at least, objectivity within a given group.

In the next section, I return to the issues of disagreement, given what we now know.

3.6 Predicting subjectivity

3.6.1 Back to faultless disagreement

The goal of this dissertation is to predict which propositions, in which contexts, would be treated as subjective. We defined a subjective predicate as one which allows for faultless disagreement. In Section 3.2, we then broke down disagreement to the following criteria, which I reproduce here:

(29) The requirements for faultless disagreement:

a. Speaker A felicitously asserts p_A from evidence e_A, and Speaker B felicitously asserts p_B from evidence e_B.

b. p_A and p_B both have the same references.

c. p_A and p_B are contradictory/not truth-compatible/cannot both update the common ground.

d. e_A and e_B are reliable: e_A should for all reasonable people lead to the belief p_A and e_B should for all reasonable people lead to the belief p_B.

e. e_A and e_B are considered equally reasonable.

Although a deep exploration of how evidence leads to beliefs is beyond the scope of this dissertation, I took two factors to be crucial: sensory perceptions and internal mental categorizations. If we allow for requirements (a) to (e) to be true (we have disagreement in good

faith, that is, both A and B are neither lying nor misinformed nor deluded but they cannot agree), then we have A and B differing in their perceptions or their categorizations. Both of these map to reliability and reasonableness.

In Section 3.3, we saw that perceptions can differ, and in fact perceptions of smell and taste can differ much more than perceptions of sight, due to the phenomenon of specific anosmia being much more widespread, varied and complex than its equivalent in the visual field, color-blindness. This explains why predicates describing taste and smell are much more acceptably subjective than visual or auditory predicates. People's perceptions genuinely do differ. But because this fact is little-known, most people assume that their perceptions of smell and taste are reliable and base their beliefs on those perceptions.

Even if a speaker knows that her perceptions are non-standard (e.g. she knows she is anosmic or hyperosmic to a certain flavor to the point that foods containing it are unpleasant for her while others would find them pleasant), although technically her perceptions are unreliable for everyone else, this does not mean that she can accept that she is wrong about her taste. This is a case of her not allowing update of the common ground with someone else's tastes, even though she knows her own tastes are abnormal. Suppose that I am hypersensitive to the bitter agent in coffee. I am aware of this fact and that my perceptions are nonstandard. I can (and if my goals are to avoid unnecessary unpleasant experiences, I should) felicitously assert:

(30) Coffee is not tasty (even though I know a lot of people believe it is).

I am not asserting a truth. Nor am I asserting a falsehood; I am not mistaken about my perceptions and I do not intend to deceive about them. I am making the pragmatic move of rejecting the assertion that coffee is tasty from the common ground, as Khoo & Knobe (Forthcoming 2016) show, because otherwise I may be left with nothing but coffee to drink, and knowing that my perceptions are nonstandard perceptions and my taste is 'defective' will not make it any less unpleasant an experience. Essentially, my proposal is of an expressivist

semantics here (as supported by arguments such as Gutzmann (2016)).

Thus we can make the following prediction:

(31) **Principle 1:** The more perceptions differ, the more subjective the predicates based on those perceptions are.

with the corollary that predicates dealing with smell and taste, where genetic variability makes perceptions differ quite a bit, would thus be more subjective than predicates dealing with visual or auditory evidence.

However, we have seen that there are other differences among human beings: not only in what we perceive, but also in what our minds make of it, or how we categorize our perceptions. We discussed this in Section 3.5, and saw that many perception differences are actually learned categorization differences, and most evaluative judgements rely on categorization. Similarly to differences in perception, it is logical to have differences in categorizations. So we can make a second prediction:

(32) **Principle 2:** The more categorizations differ, the more subjective the predicates based on those categorizations are.

This prediction explains how people with the same perceptions (e.g. people with full color vision looking at a painting) could have different judgements about the quality of the painting (or even about what is portrayed on it, if it is an optical illusion or a Rorschach test). But we have added complexity to this by looking at the factor of language and the factor of expertise. In Section 5, we saw that experts tend to converge in the categorizations they make when perceiving their subject of expertise, be it beer, wine, art, or gymnastics performances. Similarly, having words to describe categories allows for people to come to consensus on what they perceive, or rather, how they categorize what they perceive. They essentially become more objective. Thus, we have the following principle:

(33) **Principle 3:** The more categorizations are explicitly the same, the less subjective predicates based on those categorizations are. (Or: categorizations can be made the same if you name them right.)

Therefore, when perceptions converge and categorizations are shared, disagreements shouldn't be faultless.

3.6.2 Metalinguistic negotiation, the Acquaintance Inference, and the spectrum of subjectivity

Explicit categorization becomes very similar to metalinguistic negotiation, argued as a model of subjectivity by Barker (2002, 2013); Sundell (2010); Plunkett & Sundell (2013). It is about establishing for both speakers what words mean and how they apply to the given object in the given context. The gaining of expertise, and of the associated vocabulary, can be viewed as either a way of sorting out perceptions by associating them with words (what I argue, based on cognition studies like Majid & Burenhult (2013); Croy *et al.* (2015), or the sorting out of words by associating them with perceptions, as the metalinguistic negotiation model puts it. Essentially, there isn't a practical difference between the two; it all depends on what you start with.

However, Barker, Sundell, and Plunkett do not account for differences in perception (for one thing, a large part of their argument hinges on *tall*, which cannot have a difference in perception due to being a measurement in one dimension). The evidence from Croy *et al.* (2015) and other studies on individual sensory differences show that perception differences in taste and smell among speakers cannot be ignored. And indeed, we find that we need differences in perception to account for Ninan's (2014) Acquaintance Inference requirement.

If we agree on what perceptions are described by what words, then we should be able to relay a judgement of taste or valuation made by someone else and accept it as our own. However, most people's judgements are that the following in the case of *good book* is somewhat

awkward and most people find it infelicitous:

(34) Arundhati Roy's new book is good (#even though I haven't read it, but someone I agree with on the definition of *good book* has).

We should be able to do this as well as we can relay a judgement of exceeding a height standard, which is felicitous as the following example shows:

(35) Richard Feynman was tall (although I haven't seen him personally).

But we will have some trouble with that being a felicitous expression for *good book*, and even more for the example involving taste, which is infelicitous as the following example shows:

(36) Avocados are delicious (#even though I haven't tasted them, but someone I agree with on the definition of *delicious* has).

Why would you need to have perceived the object with your own perceptions? Because perceptions of taste differ, we know from experience, even though we may not be aware that perceptions of taste and smell are linked to genes. So to verify that our own perceptions match those of other people, we need to have tasted the object.

In a genetically homogeneous community, where everyone has the same tastes and the same experiences with foods, possibly it wouldn't occur to anyone that tastes might be subjective. Thus, in such a community, people would be expected to take others' word for whether foods are delicious, the same as they would take others' word for mathematical facts. Note again that pain reports are "subjective" but are not subject to the Acquaintance Inference: we accept others' word for whether experiences are painful, such that a biological male who has never given birth and never will can still assert felicitously, *Giving birth is painful*. This, I argue, is because we assume pain experiences to be homogeneous.

But we do not live in such a community, and both genetic predisposition and experiences

affect people's perceptions of taste and smell differently — this being further complicated by the fact that, unlike the Jahai, most of us lack the communication tools to determine *how*.

We thus have an explanation for the Acquaintance Inference. This would predict that the less perceptions (though not categorizations) differ about a particular predicate, the less the Acquaintance Inference is required. We may judge a person as tall without actually seeing the person, as long as we know a salient fact about their height — either its exact measurement, or whether it exceeds the height of another person or relevant object. The categorization difference between speakers in this case is whether we place this height measurement into the extension of *tall (for the comparison class)* or not.

On the other hand, for evaluative judgements, such as those for a book or work of art, the requirement for the Acquaintance Inference is not, I argue, because perceptions differ. It is because evaluative judgements involve multidimensional predicates (as I discussed in Chapter 1). The multiple dimensions that they involve are too many to be reliably relayed in a hearsay account, unlike the explicit multiple dimensions of *big* or those categorized in a list of symptoms and signs like the dimensions of *healthy*. There is a risk that another speaker could have missed noting a dimension that is very relevant to you. This would not count as a perception difference per se, due to it not resting on biological differences and changing with training and experience. However, it would be radical enough a categorization difference that it is *de facto* a difference in perception: which components you pay attention to are not that different from which components you perceive. In both cases, uncertainty and variability lead to differences in judgement.

Note that for moral judgements, the Acquaintance Inference is not required. One does not need to have experienced either stealing or being stolen from in order to felicitously assert *Stealing is bad*. Thus, I argue, they are also purely about categorization differences, unless one argues that having a strong fear reaction to certain crimes as opposed to others (having experienced them in a frightening context, e.g. having been robbed at gunpoint, or having had nightmares about them following a horrifying account or film) is a perception

difference.

What about differences in modal statements? I argue that these are due to differences in categorization: we categorize a possible world as probable or improbable based on differing *information* — which is not the same as differing perception. For teleological or bouletic modals, we are categorizing possible worlds as more or less desirable or necessary based on our mental goals or desires. Where differing perceptions come into play is in how we experience this desire or need, as we shall see in Chapter 4.

Knowing what we now know about the requirements for differing perceptions and differing categorizations, we can group predicates in order along a spectrum of subjectivity:

(37)

Predicate Type	**Differs in Perception**	**Differs in Categorization**
Objective	No	No
Dimensional	No	Yes
Modal	No	Yes
Moral	No or Partly	Yes
Evaluative	Partly	Yes
Personal Taste	Yes	Yes

This also helps explain the counterstance model that Kennedy & Willer (2016) use to explain the distribution of *find* and *consider*. They argue that the distribution of these verbs is determined by whether there is a counterstance available – a way that the world could be — that supports coordination by stipulation. What I have described as categorization differences seem to me precisely the same differences that support coordination by stipulation. What is categorization but stipulation that these things go into these categories, and what is learning different categories that others agree to but coordination? Thus the verb *consider*, as a first sketch, has to incorporate categorization, while *find* has to incorporate differences in perception as well (although I devote the next chapter to further discussion of *find*, as well as discussing the counterstance model in further depth).

We have yet to use differences in perception to explain why the Acquaintance Inference, as Ninan (2014) notes, is not a presupposition. It would seem plausible that the use of predicates where perceptions differ across speakers would presuppose that the speaker's own perception is involved. As Ninan argues, this is not the case; at least, the Acquaintance Inference does not give rise to the familiar patterns of presupposition projection that I discussed in the previous chapter, like cancellability under negation (*#The lobster rolls aren't* tasty — *I haven't even tried them!*) or projecting over conditionals, modals and questions, none of which require the Acquaintance Inference. (As we will see in the next chapter, the verb *find* with subjective predicates does presuppose the Acquaintance Inference, passing all of the classical tests.)

Ninan concludes that a better, though flawed account, is to have the Acquaintance Inference be based on epistemic requirements — a Knowledge Norm that a person may know p only from first-hand knowledge. I explain this as due to the fact that since perceptions differ, only verifying with one's own perceptions would do. However, as Ninan points out, the statement *The lobster rolls must be tasty (but I haven't tried them)* poses a puzzle for the Knowledge Norm. If we have *must* ϕ, we should be able to deduce ϕ, if *must* ϕ entails ϕ. However, von Fintel & Gillies (2010) argue that *must* is a marker of indirect evidence. Hence, we have a clear solution to Ninan's concern: indirect evidence is infelicitous with assertion of predicates of personal taste, because due to differing perceptions, only use of direct perception would do. But *must* is felicitous with such statements. Thus, *must* ϕ is allowed while ϕ itself is blocked.

This can also extend to *will* in such expressions as *The lobster rolls will be tasty,* which do not require the Acquaintance Inference to be acceptable. As Klecha (2012) argues, *will* is a modal — it refers to a possible world rather than the actual future. Thus, we include as a criterion for this possible world where *The lobster rolls will be tasty* that we also have a tasting experience in that world that satisfies that they meet our judgements for tastiness. We can verify this:

(38) a. The lobster rolls will be tasty tomorrow. (Acceptable)

b. The lobster rolls will be tasty tomorrow, #but no one will eat them because this week's batch of lobster has just been recalled. (Infelicitous)

c. We are holding a luncheon tomorrow. #The lobster rolls are tasty tomorrow. (Infelicitous).

We see that (b) is infelicitous without the prediction of a tasting event in the possible world that *will* is providing an accessibility relation to. If we use the English future present in (c), which is not a modal, then even though *We are holding a luncheon tomorrow* is perfectly fine if the certainty of the event is high, *The lobster rolls are tasty tomorrow* is not, because the tasting event has not happened — even if we are certain it will happen in the future.

A distinction between differences in perception and differences in categorization thus both provides support for the metalinguistic negotiation model of subjective propositions, and establishes its limits: it will only fully apply to issues where categorization differences play a role. It also predicts the strength of the Acquaintance Inference to be correlated with differences in perception (though not differences in categorization), and helps fill a gap in Ninan's epistemic model. This is a major step forward. I will leave further discussion of the implications for Chapter 5.

3.7 Conclusion

In this chapter I have looked into what it means to have an evidential basis for the assertion of a subjective predicate. I laid out the structure for faultless disagreement, and showed that different models of subjectivity assume that different requirements for faultless disagreement fail. I then analyzed differences in predicates as due to either differences in perception (based on conclusions from recent cognition studies and analyses from the world of professional tasters) or differences in categorization. These combine to form the evidential basis. Acknowledging differences in perception allowed me to explain the Acquaintance

Inference.

Accounting for the pragmatic and contextual factors that lead to subjective predicates is a major step forward. In Chapter 5, I will discuss how this affects most subjectivity models on the market right now. However, I will devote Chapter 4 to using the model to explain one particular well-known construction associated with subjectivity: the case of the subjective *find*.

CHAPTER 4

THE SUBJECTIVE FIND, EXPERIENCE AND MODALITY

'I thought you did,' said the Mouse. '—I proceed. "Edwin and Morcar, the earls of Mercia and Northumbria, declared for him: and even Stigand, the patriotic archbishop of Canterbury, found it advisable—"'

'Found what?' said the Duck.

'Found it*,' the Mouse replied rather crossly: 'of course you know what "it" means.'*

'I know what "it" means well enough, when I find a thing,' said the Duck: 'it's generally a frog or a worm. The question is, what did the archbishop find?'

—Lewis Carroll, *Alice's Adventures in Wonderland* (1865)

4.1 Introduction

As I argued in the previous chapters, for the most part there seem to be no syntactic distinctions between subjective and objective predicates. The difference, as I proposed in Chapter 3, is pragmatic and contextually motivated. But many languages have a verbal construction that can felicitously embed clauses with subjective but not objective predicates, what Saebø (2009) dubbed a subjective attitude verb:

(1) a. John finds the soup delicious/disgusting/?green/?vegetarian.
    b. John considers/?finds the soup vegetarian.
    c. Anna finds John annoying/admirable/?Canadian/?blond/?tall.

Any model of subjectivity needs to explain what such a verb is selecting for. If we claim that it selects for a silent pro-like argument in the content, as Saebø (2009) does, the model is easy. But inflexible: it cannot identify why *find* may select for predicates in some contexts but not others, as in this example from Pedersen (2012):

(2) (During the 2012 US presidential election debates):

A: I find Obama won the debate.

B: Really? I find Romney won.

Compare this with a context where *find* is infelicitous:

(3) (During the World Universities Debating Championship):

A. #I find Sharon won the debate.

B. #Really? I find Wallis won.

Saebø's contextualist model, assuming that *find* has a judge argument, cannot explain why there is an argument in one example here but not in the other when *debate* and *win* are the same word each time.

Since in this dissertation I am taking a view much more reliant on context, as I have argued in the previous chapters, I need to explain the consistency of judgements about *find* another way. I will explain that *find* has a direct experience requirement and as a consequence, selects for predicates where differing perceptions or categorizations are in play.

Stephenson (2007b) argues that *find* is an attitude verb evidential based on direct experience, but she does not clearly define direct experience and thus her arguments are subject to criticism, such as Saebø (2009). In this chapter, I close the gap where previous literature has loosely defined both *find* (not giving a clear diagnostic between the attitude verb meaning and the discovery meaning) and its direct experience requirement.

I show that the properties of direct experience allow for *find*'s selection of a subset of subjective predicates, those that rely on direct experience and only direct experience as their evidence, and I explain why previous criticisms of direct experience being sufficient are not really criticisms. I show how this direct experience can include modal adjectives such as *plausible, difficult* and *likely*, which have not been previously considered in the subjective literature but which *find* readily embeds. Following Gunlogson & Carlson (2016) and a number of researchers on evidentials, I show that this explains why, similarly to evidential

content, propositions embedded under *find* cannot be disagreed with. I also use a definition of experience as necessary and sufficient to explain why *find* can embed modals, but only of a certain type.

This chapter is structured as follows: In Section 4.2, I establish the terrain by giving some brief diagnostics of how to distinguish Subjective-Find from the *find* that means discovery; this will prove crucial in order to argue against previous analyses. In Section 4.3, I present the evidence for Subjective-Find as a verb selecting for subjective predicates, and argue that the selection for subjectivity is an implicature that can be cancelled. I explain the direct experience requirement, loosely for the time being, and explain how it differs from the Experiencer theta role. I argue that having direct experience as both a necessary and sufficient condition explains the behavior of *find* in discourse, and I address some prior criticisms of this explanation.

In Section 4.4, I turn to *find* embedding modal adjectives, and explain which conversational backgrounds allow *find* and which do not. I argue that all the conversational backgrounds that allow *find* in embedding are those that involve a feeling as part of the evidential basis, such as a feeling of likelihood, difficulty or desire. In Section 4.5, I create a more precise definition for the direct experience necessary for *find* that allows for modals. I conclude in Section 4.6. The Appendix looks at *consider*, showing that it is actually two different verbs and distinguishing between them.

4.2 Diagnostics for subjective-find as opposed to discovery-find

In some languages, some of the subjective attitude verbs are only derivationally related to 'discover', such as German: although *finden* is also used for discovery, 'to uncover' is *herausfinden* (Malte Willer, p.c). And other languages have verbs that only have a subjective use, with no relation to 'discover' at all: Norwegian *synes* (Saebø, 2009), Swedish *tycker* (Coppock, 2015) and Mandarin Chinese *jue de* (Marianna Zhang, p.c.).

However, English *find*, French *trouver* and Russian *naxodit'* are homophonous with the

verb meaning *to discover*, so I will use this section to establish clear diagnostics for Subjective-Find in English, showing that it differs in its syntactic distribution from Discovery-Find, and allowing us to identify it when we turn to its semantics and pragmatics in Section 3. This keeps us from making a circular argument: "How do we know that a predicate is subjective? Because it can be embedded under *find*. How do we know that this is the subjective kind of *find*? Because it embeds subjective predicates." In subsequent sections, it will allow us to show that certain arguments against the experiential properties of Subjective-Find are actually using Discovery-Find in their examples.

In this section, I show that Subjective-Find differs from Discovery-Find in Aktionsart or aspect, in passivization and in null subject embedding. I will not give a syntactic explanation for how these features of Subjective-Find should be modelled. For the purposes of this chapter, I am simply using them as diagnostics.

4.2.1 Aktionsart

Subjective-Find differs from Discovery-Find in its aspect or Aktionsart in the sense of (Vendler 1967). As the tests of Dowty (1979) show, Subjective-Find is consistently a stative, while Discovery-Find is an eventive achievement:

(4) Only non-statives occur in the progressive:

a. Maria is finding the treasure missing as we speak (Discovery, non-stative)

b. ?John is finding the soup delicious as we speak (Subjective, stative)

(5) Only non-statives occur in pseudo-cleft constructions:

a. What Maria did was find the treasure.

b. ?What John did was find the soup delicious.

(6) Non-achievements can be modified by *for [time]* while achievements take neither *for [time]* nor *in [time]* but when coerced as accomplishments, take *in [time]*:

a. Beatrice found the answer in an hour/?for an hour.

b. Betty found her dinner cold ?in an hour/?for an hour.

c. John found fish disgusting for a month/?in a month.

However, the picture gets more complicated cross-linguistically: the Russian verb meaning 'find' (both the discovery and subjective senses) would take either the imperfective atelic *naxodit'* or the perfective telic *najti* form. One would expect that the *naxodit'* would correlate with subjectivity, as statives and activities map to imperfectives in Russian. However, although small-clause (instrumental case) constructions with imperfective form and subjective meaning are acceptable from judgements:

(7) Ja naxozhu jeto vkusnym.
I find-IMP-1sg this delicious-INSTR
'I find this delicious.'

they are completely absent from the Russian National Corpus, with all of 1337 small-clause tokens (across all numbers and tenses) being in the perfective aspect, including 153 judged as subjective, as they embed predicates like 'necessary, ugly, delicious, smart, funny' — but are perfective. I do not currently have an explanation for this.

4.2.2 Passivization and null subject operators

We will add two further features that distinguish Subjective-Find from most attitude verbs as well. Unlike *consider* or Discovery-Find, Subjective-Find is hard to passivize:

(8) a. The treasure was found by Maria under the floorboards.

b. ?Abalone is found delicious by many gourmets.

c. Abalone is considered delightful by many gourmets.

Similarly, unlike *consider* or Discovery-Find, Subjective-Find cannot have a null operator in

subject position:

(9) a. The treasure was easy to find.

b. Jason was difficult to consider admirable.

c. ?Jason was easy to find annoying.

This runs counter to the analysis of Svenonius (1994), which grouped *find* and *consider* together (as well as *regard* and others) and opposed them to the *want, love, need* class of verbs. Two of Svenonius's indicators were passivization and null subject operators:

(10) a. Mary is considered loyal to her friends.

b. Santa Cruz is regarded as a great place to spend the winter.

c. *That man is wanted off my ship.

d. *Metallica is loved turned up loud.(Svenonius (1994), ex. 6, p. 91)

(11) Null subject operators, *consider* class:

a. Who do you like without really considering admirable?

b. That's the kind of evidence you can use without proving irrefutable.

c. Mikey is easy to consider intelligent.

d. That kind of evidence would be difficult to show invalid.

(12) Null subject operators, *want* class:

a. *Jeff is easy to want arrested.

b. *That's the kind of dog that it's difficult to like in the house.

c. *Who do you despise without really wanting arrested?

d. *That's the kind of evidence you can use without fearing refutable. (Svenonius (1994), exx. 16-17. p. 94-95)

However, we see that *find* does not pattern with *consider* in those two tests. This may be

due to the direct experience presupposition, which requires a subject capable of experiencing and so cannot take null or covert subjects; other verbs that require an Experiencer argument, such as *suffer, enjoy*, can take null subjects:

(13) a. That's the kind of film it's easy to enjoy drunk.

b. That's the kind of film you can watch without enjoying.

c. This is an illness it's easy to suffer from.

d. Who do you despise without suffering in their presence?

The acceptability of these suggests that it is not an Experiencer argument that plays a role, so I will withhold judgement.

As the Aktionsart diagnostic has the crispest judgements, in my perception, I will generally be using that one, leaving to future research how the other diagnostics may follow or result from the properties of Subjective-Find. In the remainder of the chapter, where there is no risk of confusion, I will refer to Subjective-Find simply as *find*.

Now, armed with solid tests to distinguish Subjective-Find from constructions that are actually about discovery, we turn to the properties that make Subjective-Find so interesting for the study of subjectivity in general, and account for them.

4.3 Subjective predicates and the experience requirement

4.3.1 Subjective predicate embedding

First, I will present the evidence about Subjective-Find's most famous property: the fact that it embeds subjective predicates, and seemingly only subjective predicates, as Saebø (2009) observed for English, Norwegian and German, Ducrot (1975) for French, and many other researchers on subjectivity followed. We see this contrast in the following examples:

(14) a. Sue finds the soup delicious.

b. #Sue finds the soup vegetarian/green.

c. Sue finds John fascinating/annoying/intelligent/tall.

d. #Sue finds John Canadian/vegetarian/six feet tall.

This is often compared to *consider*, which does not have this contrast:

(15) a. Sue considers the soup vegetarian (because it contains shellfish, but not the meat of birds or mammals).

b. Sue considers John Canadian (because he has lived in Canada for years).

c. The airline considers/#finds this bag heavy, so you will have to pay an extra baggage fee.

d. I consider/#find the Earth to be flat.

Hirvonen (2014) argues that *find* imposes a further requirement, that the predicate has to be gradable as well. She uses this to rule out a problematic sentence that Saebø (2009) has proposed:

(16) ?Homer finds himself gay.

Hirvonen argues that it is the fact that *gay* is not gradable that makes this sentence infelicitous under *find*. (I will return to this sentence in Section 4.6 and offer a different argument.) However, Bylinina (2013) has substantial data on predicates that can be embedded under *find* but are not gradable, specifically, extreme adjectives like *gorgeous, gigantic, huge* or *microscopic*. They lack comparative forms altogether, but are fine under *find*:

(17) I find the cinematography in this film gorgeous/?more gorgeous than in *Lawrence of Arabia*.

Gradability is not a sufficient condition for infelicity under *find* either: *spotted* is a well-

known gradable adjective (discussed extensively in (Kennedy & McNally, 2005)), but it is entirely infelicitous to say:

(18) ?I find this egg spotted.

Thus, I will reject gradability as a possible candidate for Subjective-Find's selection properties, and concentrate on subjectivity. My explanation for *find*, which I use the rest of this chapter to defend, is that what seems to be selection for subjectivity is actually an epiphenomenon of an attitude verb selecting for a direct experience evidential basis, which I will first explain in the next subsection.

4.3.2 A direct experience requirement

It has been noted in the literature as far back as (Borkin, 1973) that Subjective-Find, especially in the reduced form without an inflected verb, requires that the subject be capable of thought and feeling. Borkin provides the following contrast:

(19) a. The survey found this chair to be comfortable.

b. ??The survey found this chair comfortable.

The use of *find* in (b) is infelicitous because surveys do not have thoughts or feelings and cannot experience comfort. This is a noted contrast with Discovery-Find, as shown in the following examples; I verify that this is Discovery by using a cleft construction:

(20) Discovery:

a. The study found the effect significant.

b. What the study did was find the effect significant.

(21) Subjective:

a. #The study found the effect pleasant.

b. ??What the study did was find the effect pleasant.

Stephenson (2007a), Pearson (2013), Bylinina (2013) and Hirvonen (2014) note than the subject of *find* must have direct experience of the object in question. We see this in the following examples:

(22) ??Anna finds the cat food tasty, because the cat ate it all up.

(23) a. Max finds beer delicious #but has never tasted it.

b. I find La-Z-Boys very comfortable #just from how they look in ads.

This experience may have happened before the statement time, rather than simultaneously:

(24) a. I last tried avocados ten years ago, and I find them disgusting.

b. I heard this band play back in 1987. I find them great.

Let us take a quick initial sketch of what it must mean to have "experience of the relevant kind" by the subject. Evidently, it varies according to the object. For example, to experience a food, one must have tasted it; to experience a written work, one must have read it, to experience a visual work, film or performance, one must have seen it, and to experience a piece of music, one must have listened to it. Some kind of sensory perception is involved (at first glance).[1] The use of *find* obeys those constraints, as shown in the following infelicitous examples:

1. A good question, that I do not at this time have a solid answer to, is whether the experience has to be conscious. There are many hits on Google for the string *find sleeping (uncomfortable/difficult/etc.)*. Sleep is the prototypical example of an unconscious state; however, most of the Google hits describe sleep negatively. Thus, they seem to be talking about a conscious experience of *not* sleeping, and the subject being awake enough to note that they were feeling discomfort or difficulty:

(i) a. I find sleeping on my back uncomfortable.
b. I find sleeping difficult.

So I would tentatively answer yes, the experience has to be conscious, but I would need further research to confirm this.

(25) a. Max finds beer delicious #but has never tasted it.

b. Athena finds Scorcese's oeuvre boring #although she has never seen a single film of his.

c. Olivia finds the poetry of Jean Racine delightful #without having read it.

d. John finds the music of Bartok repellent #despite having never heard it.

e. Sam finds the cake beautiful/delicious#(-looking), having just seen it in the bakery window.

Stephenson (2007a) incorporates this into her analysis by giving the following definition for *find*:

(26) $[[find]] = [[think_G]]+$ 'and this is caused by x having a direct experience of p in w.'

However, she does not elaborate on what direct experience means, and resorts to saying,

> Of course this notion of direct experience is a bit vague, but for present purposes the idea is this: whether something tastes good to someone depends on the internal psychological state of that individual, and the only way to have direct experience of an individual's internal psychological state is to be that individual. (Stephenson (2007b), p. 61)

In Section 4.5, I will refine this notion. *Find* cannot refer to just any internal psychological state, else we would have to judge *??I find 5 prime* to be acceptable, since a state of conscious knowledge that 5 is prime should be a psychological state. For now, we will see how far we can get with just a loose sense of 'direct experience.'

We saw in Chapter 3 that in situations where perceptions and categorizations of the evidential basis for a particular proposition differ considerably between speakers, the evidence is required to be direct — one must perceive with one's own senses rather than rely on someone else's differing senses. In the rest of this chapter, I elaborate further as to what that entails.

This experience requirement is a presupposition for propositions embedded under *find*, unlike the bare assertions of subjective predicates (Ninan, 2014) discussed in Chapter 2. Even when *find* is negated, the requirement is still present, as in the following infelicitous examples:

(27) a. Max doesn't find beer delicious #but has never tasted it.

b. I don't find La-Z-Boys comfortable #just from how they look in ads.

In the following example, we further apply the presupposition tests based on Karttunen (1973); Stalnaker (2002); Simons *et al.* (2010); Potts (2014), and confirm that the direct experience perception requirement projects as a presupposition under modals (a), conditional antecedents (b), and questions (c). Furthermore, we verify with (d) that cancelling the presupposition works in much the same was as does presupposition cancellation for such predicates as *stop* (*John didn't* stop *smoking — he never smoked at all!*).

(28) a. Max might/should/must have found the film appealing #but he has never seen it.

b. If Max found the film appealing, he would talk about it to all of his friends, #but he has never seen it.

c. Did Max find the film appealing? #No, he has never seen it.

d. Max *didn't find* the film appealing — he has never seen it! (Presupposition cancellation).

All of these confirm that the direct experience requirement for *find* is a presupposition.

4.3.3 The experiential presupposition and the Experiencer theta role

Here, we must distinguish between a presupposition of *experience*, defined flexibly for the time being, and the Experiencer theta role, one of a set of theta roles used in syntax and

semantics to explain the selection properties of predicates. (I am going to use the capital letter to distinguish the Experiencer theta role from an experiencer meaning a subject who experiences, even though some of my sources do not use that distinction.) While even lay thinkers about language can intuitively agree as to what *experience* is, or at least what should and should not count as a direct experience, there is no widespread agreement even among theorists as to what a theta role is.

Bylinina (2013, 2016) presents a range of evidence that predicates of personal taste like *tasty* and *fun* (but not dimensional adjectives like *tall* or evaluative adjectives like *smart* and *beautiful*) project an Experiencer argument that has overt syntactic realizations in multiple languages. These include prepositional phrases *to X* and *for X* in English and the dative experiencer NP in Russian, Hungarian:

(29) a. These avocados are tasty/#plump/#black to me.

b. Mne jeta kniga interesnaya/#tolstaya/#chornaya.
me-DAT this book interesting/thick/black
"This book is interesting/#thick/#black to me."

Hungarian has a dative Experiencer argument. It can be applied to both subjective predicates such as *important* and vague predicates such as *tall*, but Bylinina (2013, 2016) argues that in the two cases, the semantic properties differ: when the dative Experiencer appears with *important*, it has the properties of an argument, whereas when it appears with *tall*, it has the properties of an adjunct.

(30) Hungarian (Bylinina, 2013, 2016):

a. János-nak fontos-ak vagyunk
John-DAT important-PL be.1PL
'John finds us important.' (Argument)

b. János-nak minden felńott magas
John-DAT every adult tall
'For John, all grown-ups are tall.' (Adjunct)

In Japanese, *tasty* is restricted to first-person usage, as shown in the example below in (a). If used in the second or third person, it must be marked with an evidential, as in (b) .

(31) Japanese (Bylinina, 2013, 2016):

a. watasi-ni-(totte)-wa / *John-ni-(totte)-wa kono keikii-wa oisii
I-DAT-TOTTE-TOP / John-DAT-TOTTE-TOP this cake-TOP tasty
"This cake is tasty to me / to John."

b. John-ni-(totte)-wa kono keikii-wa oisii noda / ni
John-DAT-TOTTE-TOP this cake-TOP tasty EVID / there's.no.mistake
tiganai

"The cake must be tasty to John."

Evaluative predicates such as *beautiful*, *smart*, *boring*, or *foolish*, Bylinina argues, do not have Experiencer arguments because they do not take either *to X* prepositional phrases in English nor dative experiencer NPs in Russian.

(32) a. ?John is smart to me.

b. ?Mne Vanya umnyj
me-DAT V. smart
Intended: "Vanya is smart to me."

The predicates with explicit Experiencer arguments seem to take *find* quickly and easily. To Saebø (2009), *find* also takes predicates that do not have an explicit Experiencer argument. However, McNally & Stojanovic (2015) argue that this may be coercion. They use corpus data to argue that *find* primarily selects for experiencer semantics (which they do not identify with a morphosyntactic indicator as Bylinina does, but treat simply as personal experience), and argue that *find* does not easily take evaluative predicates such as *mediocre*:

> A search of the British National Corpus revealed only 9 uses (out of 2353 total tokens of 361 different adjectives) with *good/better/best*, 4 with *beautiful*, and 1 with *pretty* in the complement to evaluative *find*; there are none with *bad,*

> *mediocre, great, excellent, awesome, super, gorgeous, handsome*, and *ugly*. In contrast, the adjectives that occur most frequently with *find* (over 20 times) are *difficult, hard/harder, easy/easier, useful, helpful, impossible, necessary, interesting, attractive*, and *strange*, none of which are explicitly expressive of positive or negative attitude or evaluation. The general failure to find evaluative adjectives under *find* strongly suggests that their evaluative component is not based directly on personal experience. (McNally & Stojanovic (2015), p. 9-10)

McNally and Stojanovic also point out using examples that *find* is odder than *consider* in assertions involving evaluative or moral judgements, and carries a strong experiential flavor:

(33) a. ?I find Miró's mosaic on the Rambles mediocre.
b. I find lying bad/worse than stealing.

(34) a. I consider Miró's mosaic on the Rambles mediocre.
b. I consider lying bad/worse than stealing.

They point out that *I find lying bad/worse than stealing*, though acceptable, strongly implies that the speaker has made their evaluation about lying on the basis of specific experiences of doing it, while the evaluation of Miró's mosaic is based on a comparison with other works of art that the speaker has experienced.

This leads McNally and Stojanovic to at least infer that *find* is *solely* about experience/the Experiencer. They are not very specific about what the Experiencer is. Bylinina is, but her account has several problems. She draws an equivalence between the direct experience requirement and the Experiencer argument, and claims that predicates of personal taste are the only ones that have both, as summarized in the quotations below:

> As discussed in some detail in section 3 above, PPTs show a 'direct experience requirement' — a statement like *This cake is tasty* can be made only if the judge has had direct perceptual experience involving the subject (i.e. has tasted the

> cake). For me this means that PPTs lexically include a reference to an experience event as part of their semantics, while POS-DAs do not. As experience events involve an experiencer, an extra argument gets projected in the case of PPTs but not POS-DAs. (Bylinina (2016) p. 17)
>
> However, one parameter that we have seen seems to correlate with the presence or absence of an extra argument — it is the reference to an experience event (and thus the presence of an experiencer [sic] argument) as part of PPT semantics, but not anywhere else. (Ibid., p. 23)

To claim an exact equivalence between experience events and Experiencer arguments is too strong a claim. First of all, although evaluative adjectives like *beautiful* do not have syntactic Experiencer arguments in the languages that Bylinina looks at, they do have a direct experience presupposition with *find*:

(35) a. Athena finds van Gogh's "Starry Night" beautiful, #but has never seen it.

b. ?Van Gogh's "Starry Night is beautiful to Athena.

c. ??Afine "Zvezdnaya noch" krasivaya
Athena-DAT Starry Night beautiful
(Intended) "Starry Night" is beautiful to Athena.

Bylinina says explicitly that "I do not have an explanation for this split [that some speakers accept *Anna finds her bowl of pasta big* and similar uses of *find* with dimensional predicates, while others do not], nor do I have an account of the subcategorisation properties of *find*." (p. 17.) But we see that predicates that do not have Experiencer arguments, which Bylinina would therefore predict should not have direct experience requirements, do have them under *find* (and also have an Acquaintance Inference requirement, as discussed extensively in Chapters 2 and 3).

Thus we see that the direct experience requirement must not be the same thing as the Experiencer, although I will tentatively assume that in general predicates with an Experiencer

argument also have a direct experience presupposition, so Experiencer implies experience but not vice versa. To extensively test that assumption is beyond the scope of this thesis.

Moreover, not all predicates of personal taste in all languages have an Experiencer argument according to Bylinina's tests. The predicate *tasty* has since Lasersohn (2005) been the prototypical example of a predicate of personal taste in English, and gets the Experiencer PP (according to Bylinina) *tasty to me*. There should be no reason why its Russian equivalent, *vkusnyj*, does not do the same. However, unlike *interesnyj* 'interesting,' the predicate that Bylinina uses for her Russian examples, the predicative adjective form of *vkusnyj* does not take the Dative NP that she claims is the indicator of an Experiencer:

(36) a. Mne jeta kniga interesnaya
me-DAT this book interesting
"This book is interesting to me."

b. ??Mne jetot sup vkusnyj
me-DAT this soup tasty
Intended: "This soup is tasty to me."

Dative NPs are only acceptable with the adverbial form of *vkusnyj*, *vkusno*, referring to a particular experience event.

(37) While eating soup:

a. Mne vkusno
me-DAT tasty-ADV
"(It's) tasty to me."

However, Bylinina's claims generalized to all predicates of personal taste, without drawing such exceptions. Bylinina does not use *vkusnyj* in her examples, but does not comment on this lexical gap, either.

I cannot explain why *vkusnyj* has this restriction, although it may be related to a distinction pointed out by Gunlogson & Carlson (2016) between the predicates *tasty* and *fun* that are usually grouped together and treated alike in the subjectivity literature. Gunlog-

son and Carlson point out that *fun*, along with *boring, nerve-wracking, comfortable*, etc. is predicated on a characteristic experience (an event) pragmatically associated with its object, rather than the object itself:

(38) The roller coaster was fun.

a. = The ride (on the roller coaster) was fun.

b. = Riding the roller coaster was fun.

(39) a. That book is boring.

b. =Reading that book is boring.

(40) a. The wait was nerve-wracking.

b. =Waiting was nerve-wracking.

(41) a. That chair is comfortable.

b. =Sitting in the chair is comfortable.

However, *tasty* does not pattern in the same way, but in a different way, similar to *stinks* and *beautiful*, in being conceptualized as a property of the object, not the experiential event of perceiving or consuming it:

(42) The chilli was tasty.

a. ≠ ??The consumption of the chilli was tasty.

b. ≠ ??Eating the chilli was tasty.

(43) a. The garbage stinks.

b. ≠ ??Smelling the garbage stinks.

(44) a. The view of the ocean was beautiful.

b. ≠ ??Viewing the ocean was beautiful.

But to complicate the picture yet further, only some of the predicates in the *fun* class have the same pattern in Russian, allowing both the object and the event to be described using the predicate ('boring' = *skuchnyj* and 'comfortable' = *udobnyj*, and 'fun/engaging' = *zanimatel'nyj* do, but 'fun-ADV' = *veselo* does not).

I do not have a full explanation for this, but I use this as evidence to show that Experiencers, however they manifest, are more complex than they may sound from Bylinina's (2016) discussion, and equating the direct experience presupposition with the Experiencer argument is not a reliable move. Experiencer theta role contexts seem to be a proper subset of contexts which involve a direct experience. Thus Experiencer theta roles would imply direct experience, but not vice versa, and Bylinina's model overgenerates. For the purposes of this chapter, I will refer to the direct experience presupposition only, and leave the question of how it is cashed out for predicates that do not have an Experiencer argument for further research.

4.3.4 Two earlier models relying on contingency

Bouchard 2012: Subjective contingency presupposition

To further buttress my claim that *find* is about experience, in this subsection I analyze two recent proposals about *consider* and *find* that do not involve experience. I show that one is seriously flawed, while the other, in its current form, undergenerates. The first is Bouchard (2012), who follows a relativist view of subjective predicates involving a judge index, and looks at the presuppositions of *find*. However, the way he analyzes *find* is that it selects for subjective propositions (he has no opinion on the sources of subjectivity, but takes it as read that some predicates are subjective while others are objective, and is concerned with how to model this distinction). He posits that everything that is not subjective in the complement of *find* is presupposed.

But the examples he uses involve full inflected clauses under *find* — *John finds that*

Mike gave a great class yesterday — which poses problems for his analysis, as many of his examples are actually of the Discovery-Find. His model presents nothing that a model based on experience cannot account for, and in fact overgenerates because it would predict that the subject would find p true, if there are accessible worlds where p could be true and accessible worlds where p could be false, even if the subject has no experience with p. But this is contrary to most speakers' intuitions about *find*, as I will show using Bouchard's own examples (60) and (61):

(45) a. John finds that Mike gave a great class yesterday.

b. Mike gave a class that was great for John yesterday.

Bouchard sees the distinction between the two sentences as (a) having the proposition *Mike gave a class yesterday* be presupposed, while (b) has it asserted. I reply that there is a narrower presupposition for (a), which is that Mike gave a class yesterday, *and John was in it*. Let us assume that the class Mike gave was on John's favorite topics, or provided missing data or sources for John's work, or had jokes in it perfectly matched to John's sense of humour, but John was not in class. Thus (45)(a) is infelicitous while (45)(b) is still felicitous.

He notices that the presuppositions for *find* only work if it is possible for a judge to yield true for the complement clause, and a judge to yield false. He codifies this into his Subjective Contingency Presupposition:

(46) Subjective Contingency Presupposition (Bouchard, 2012):
Presupposition for $x\ finds\ p = (\exists w'.[w' \in Acc(w) \& p_{w',j} = T]) \& (\exists w''.[w'' \in Acc(w) \& p_{w'',j} = F])$

(where $Acc(w)$ is the set of accessible worlds, w', w'' are worlds in it, and j is a judge).

Bouchard relies on "the fact that the non-subjective part of the complement clause is

assumed to be true by the speaker" (p. 180). However, this does not explain the infelicity of the following:

(47) John finds that Mike gave a great class yesterday #but John missed the class.

Bouchard's model does not predict this; there are accessible worlds in which *The class was a great class for John/The class would have been a great class for John* is true, and accessible worlds where it is false. But the direct experience presupposition actually encompasses Bouchard's proposal that everything in the complement of *find* except the subjective predicate is presupposed. For a subject to find X p/that X is p, she must have direct experience of X. Thus it is presupposed that she has an access relation to X, which gets the same results as Bouchard's model.

Bouchard does not believe that *find* has an evidential component of direct experience; he argues that it is not *find* that is evidential, but predicates of personal taste that are. For example, in the famous case

(48) John finds the cat food tasty.

he claims that it is the *tasty* that makes us assume John has direct experience for this. However, as Ninan (2014) shows (see Chapter 2 for extensive discussion), personal acquaintance with the object in the case of *tasty* is not in fact presupposed, while as I confirmed above, it is presupposed with *find*.

Bouchard's other counterargument is the fact that this example under *find* can only be read de re:

(49) John finds that you married someone with a strange name. (Bouchard (2012), ex. 113)

Thus, de dicto readings do not arise with *find*, which he claims is because the DP cannot

scope outside the complement clause, the way he analyzes *think*:

(50) a. John thinks that someone is a genius (because he found a piece of clever machinery and, without knowing who built it, assumes that whoever built it is a genius)

b. John finds that someone is a genius (can only be read as John having a specific someone in mind).

However, first of all, the example in (50-b) is a Discovery-Find (*What John finds is that someone is a genius* is perfectly acceptable, while something like *?What John finds is Mary/someone brilliant*, as a cleft formed out of *John finds Mary/someone brilliant*, is not).

Second, it is not necessary to rule out de dicto readings by using a movement analysis; a direct experience requirement would do. Thus we see that in modal contexts, a de dicto reading for an indefinite DP is available:

(51) John finds some woman having killed JFK implausible.

Both readings are available: that John has a specific woman in mind (say, Sarah Jane Moore, who attempted to assassinate Gerald Ford), and that *any woman* having killed JFK is implausible; in fact, the latter reading is preferred, in my view. Thus, Bouchard's argument, insofar as I understand it, does not work.

Thus, Bouchard's arguments rest on a fallacy of equivocation, and do not extend to modals. His dismissal of *find* having an experience requirement has too many holes.

Kennedy & Willer (2016): Counterstance contingency

As I explain in Chapter 1, Kennedy & Willer (2016) explain *consider* and *find* in terms of what they call *counterstances*, which are similar to Bouchard's model of accessible worlds where p is true and accessible worlds where p is false. They point out that a speaker uttering

Kim considers Lee a vegetarian or *Kim finds Lee fascinating* implies that it is in some sense "up to the subject (Kim)" how the words *vegetarian* or *fascinating* are used — Kim has a choice among several discourse options.

> [W]e propose that [subjective attitude verbs'] felicitous use in discourse presupposes the existence of distinct contextually provided discourse alternatives that arise from users awareness of the contingency of how contextual underdetermination of meaning is resolved. We label these discourse alternatives COUNTERSTANCES.

I will quote their model as follows:

> Instead, we propose that [the distinction between 'subjective' and 'objective' predicates] corresponds to a pragmatic distinction between those aspects of the common ground that discourse participants take be grounded in objective facts of the world and those they take to be (to a certain extent) arbitrary matters of linguistic practice. To model this distinction, we embellish our context model with a function k that tracks the contingency of the stipulations involved in achieving an information state. k takes an information carrier s and derives a set κ(s) of s's COUNTERSTANCES: each such counterstance agrees with s on its factual information but disagrees on contextually salient decisions about linguistic practice. (p. 5)

They explain the distinction between *consider* and *find* as *consider* involving ordinary counterstance contingency, and *find* involving what they call *radical* counterstance contingency — that which cannot be resolved using coordination by stipulation, e.g. *Let's agree to consider Lee vegetarian* but not *Let's agree to find Lee fascinating.*

> According to our analysis, $find$ semantically presupposes that its prejacent is radically counterstance contingent, which is to say that it is contingent across all

> partitions of the relevant counterstance space in which the parameters of meaning that support coordination by stipulation are fixed. In order to satisfy this presupposition, it must be the case that the meaning parameters that *resist* [italics original] coordination by stipulation and on which the truth of the prejacent depend are indeterminate in the context, since fixing these parameters to a particular taste, experience, stance, perspective, or 'judge' (to use the term most familiar from the literature) eliminates variability in the truth of the prejacent within each partition of the counterstance space. (p. 15-16)
>
> To say that proposition is presupposed to be radically counterstance contingent then, is to say that neither the facts of the world nor the conventions of linguistic practice that support coordination by stipulation provide a basis for asserting or denying the prejacent, and further that the experiential/perspectival factors relevant for evaluating its truth are indeterminate. Given this degree of uncertainty, what could provide the evidential basis for an assertion or denial of the at-issue component of the meaning of a find statement, namely that the attitude holder believes the radically counterstance contingent proposition expressed by the prejacent? The evidential basis for such an assertion or denial, we claim, is that the attitude holder has experience of the sort that supports their belief, which the case of e.g. *tasty* is experience of the taste of the relevant stuff, in the case of fun is participation in the relevant event, and so forth. (p. 16).

The counterstance model is effective at explaining *consider* in particular, but relies on indeterminacy of the evidential factors. Where does the indeterminacy come from? My model adds to it by pointing out that the indeterminacy arises out of the variability among speakers. If we refine the experiential presupposition and the notion of experience well enough, then we have no need for counterstances for an analysis of the phenomenon. Counterstance effects and subjective contingency will simply fall out of the fact that human experiences

differ. In Chapter 5, I explain how my analysis and the counterstance model get the same result.

Neither of these models do much better than simply working with *find* as an attitude verb with an evidential requirement of direct experience. Nor do they (evidently) predict the role of *find* in the discourse, which I address in the following subsection.

4.3.5 Discourse effects of experience

Recall the prototypical example of faultless disagreement from previous chapters:

(52) Sue: The soup is delicious.
Zoe: No, it's not!

Sue can follow up with a *find* clause.

(53) Sue: The soup is delicious.
Zoe: No, it's not!
Sue: Well, I find it delicious.
Zoe: #No, it's not!
Zoe: #No, you don't.
Zoe: Well, I don't.

Once Sue has embedded her assertion of the soup's deliciousness under *find*, it is no longer felicitous for Zoe to argue with her. Disagreement with a *find* clause with a first-person subject becomes an example of what Gunlogson & Carlson (2016) call *faulty disagreement*, discussed extensively in Chapter 2. Moreover, this comes at a cost: it seems odd for Sue to assert that the soup is delicious *after* she has asserted that she finds it so. Once she has retreated, she cannot advance again.

(54) Sue: Well, I find the soup delicious.

Zoe: I don't find it delicious.

Sue: #The soup *is* delicious.

Note that this is only true when the subject is the same as the speaker; if Sue and Zoe are talking about someone else's tastes, they can still felicitously (and in fact, objectively rather than faultlessly) disagree:

(55) Sue: Well, John finds the soup delicious.

Zoe: No, he doesn't!

Here, Sue and Zoe's disagreement is objective: if we ask John whether he finds the soup delicious, or observe his behavior, one of Sue or Zoe will be proven right, and one will be proven wrong.

Gunlogson & Carlson (2016) use the experiential presupposition to explain this (they do not discuss *find* directly, concentrating on *fun* modified by a prepositional phrase like *fun for me/fun for Mary*, but their observations hold true for *find* as well). They argue that in the case of predicates of personal taste, like *fun*, the person experiencing the event in question (the experiencer, not to be confused with the Experiencer theta role discussed above) is the foremost authority on his or her experience. Thus faulty disagreement arises because to disagree with another's experience is presumptuous. When Sue said *I find the soup delicious*, she framed it as her own experience, which is better evidence than Zoe's, so Zoe cannot disagree with her.

This is consistent with *find* being associated with a direct-experience evidential basis; as shown by Faller (2002), Murray (2011) and Korotkova (2016), among many others, in languages with obligatory evidentials, disagreeing with evidential content is also infelicitous.

When we look at uses of *find*, we need to make the essential assumption that people's experiences differ, and are personal; they cannot be overwritten by another person's order, request, or bargaining (and indeed, Western culture views attempts to do so as repulsive,

"thought policing" or "brainwashing" and a violation of people's autonomy). Thus experience would naturally not be subject to coordination by stipulation.

Since we experience most things, comes the natural question, why don't we use *find* all the time? If I perceive with my eyes that the glass is empty, why is *?I find the glass empty* strange?

Hirvonen (2014) tries to explain that that *find* is weaker than a bare assertion, and thus the infelicity of *I find the glass empty* is by conversational implicature:

> There's a perfectly natural Gricean explanation. A *find* judgment is weaker than a judgment that attributes the property (e.g. "The glass is empty"). But when one sees an empty glass one knows the glass to be empty. To say that one merely $finds$ it empty would be odd from the point of view of communication, and thereby it would probably trigger the expectation that some implicature is intended. (Hirvonen (2014), p. 63)

However, Hirvonen's explanation is too weak. The oddness of *I find the glass empty*, which Hirvonen would assume implies that one does not know for sure, cannot be cancelled as easily as standard implicatures like *some* implying not *all* or *or* implying not *and*, as we can see in the following example:

(56) a. Some of the cookies are gone; in fact, all of them are.
b. You can have coffee or tea; in fact, you can have both.
c. #I find the glass empty; in fact, I know it is.
d. #I find cookies tasty; in fact, I know they are.

Although implicature certainly plays a role, I am taking another route. I argue that use of *find* in the discourse is felicitous if direct experience is sufficient as well as necessary. Thus classically objective predicates like *prime* would not work, as there is no way of directly *experiencing* a number like 13,497 as being prime — looking at it would not do! Same with

examples like Bylinina's example below:

(57) #I find Vera a five-year-old girl.

This is infelicitous because just looking at Vera or interacting with her is not sufficient, as one may be mistaken about her age (or even gender). On the other hand, we can have the example:

(58) I find Vera a typical five-year-old girl.

For this, interacting with Vera is sufficient experience, as long as the speaker knows already she is a five-year-old girl and compares the experience with other known five-year-old girls he or she has previously encountered. Thus *typical* would vary with different speakers' experience and would thus be subjective.

Thus, explicit categorizations, which reduce or eliminate subjectivity as we saw in Chapter 3, also reduce or eliminate the felicity of *find*, as experience becomes unnecessary (hearsay becomes more reliable). Lacking experience of the object (saying you find X *tasty* without having tasted X) would also make the use of *find* infelicitous (experience is necessary). Using a predicate that does not require experience (*prime*) would be infelicitous.[2]

We have already seen that direct experience is necessary for most genuine uses of subjective *find* and previous models do not cover this (we have not yet explained how this accounts for modal adjectives, but I will do so in Section 4). Let me show how this accounts for one common criticism of *find*, and show that the criticism is baseless.

2. It may be an artifact of Western culture and its interpretation of the Maxim of Quantity that direct experience is not considered a requirement for most predicates. According to Dechaine *et al.* (2014), in the discourse of the Plains Cree, many propositions that would be asserted and felicitously disagreed with if translated into English are stated in a form known as the root indexical clause, which, like *find*, relies on one's own experience and does not allow disagreement.

4.3.6 Experience being sufficient as well as necessary

Let us address the problem that Saebø (2009) sees with Stephenson's definition of *find* as involving direct experience. He points out "However, by the same token, [(59)] should be good, as Homer can have the most direct evidence possible of his own sexual orientation."

(59) ?Homer finds himself gay.

I explain that (59) is infelicitous because *gay* is actually a compound concept and is not experiential. Certainly one can have experience of being sexually attracted to men, and a Google search shows many examples of *find* embedding that phrase:

(60) a. Homer finds himself sexually attracted to men.
b. I sometimes find myself attracted to men, other times to women.
c. That's why I find myself attracted to men at times.
d. I find myself sexually attracted to men, but I don't like to date them.

Even though one may suggest that the original sentence could have been a Discovery-Find — *Homer discovered himself to be attracted to men* makes sense — we can confirm that this is the Subjective-Find, because in most of these examples, *discover* does not make sense:

(61) a. Homer discovers himself to be sexually attracted to men.
b. ??I sometimes discover myself to be attracted to men, other times to women.
c. ??That's why I discover myself to be attracted to men at times.
d. ??I discover myself sexually attracted to men, but I don't like to date them.

But we notice that *gay* does not occur in the same contexts as *sexually attracted to men* does, the former being an individual-level predicate while the latter is a stage-level predicate. Homer may be attracted to men at a club party; he almost certainly is not attracted to men while at home alone doing his laundry; but he is gay throughout.

I claim that *find* is acceptable where the experience is sufficient for the predicate to apply. Being sexually attracted to men cannot be sufficient to count as gay — or else every heterosexual woman would also count! Even if we add that *gay* must mean *attracted to men* + implicitly *male*, that still includes bisexual men. We must define *gay* as *exclusively or almost exclusively attracted to men* + (implied) *male*. The more factors we add, the more the experiential dimension of *gay* shrinks. If *gay* is a conjunction of factors, then *find* is ruled out by Saebø's own evidence that conjoining a subjective (experiential) and a non-subjective (non-experiential) predicate under *find* is impossible:

(62) a. You must be handsome and below 45.

b. #She finds him handsome and below 45.

c. She finds him handsome and pleasant to be with.

We verify this impossibility of conjoining two predicates, one subjective and one not, for the components of the meaning of *gay*:

(63) a. Homer finds himself sexually attracted to men.

b. #Homer finds himself sexually attracted to men and male.[3]

This resolves one well-known complication of the presupposition of direct experience: that "?Homer finds himself gay" is not really a counterexample. But we still have to deal with another one that, to my knowledge, has not be discussed in the literature. At first glance, the direct-experience presupposition does not account for how *find* can take modal adjectives. Accounting for *find* embedding modal adjectives is the topic of the next section.

3. According to some views of gender, gender identity is experiential. I would predict that a holder of such views should judge (d) as more felicitous and judge a sentence like *Even while she was a top men's decathlete, Caitlyn Jenner found herself female* as acceptable, but I cannot comment on this further.

4.4 Find and modality

Previous work that discussed the properties of Subjective-Find and accounted for it via experiential semantics or using another model has focused almost solely on predicates of personal taste, evaluative adjectives, and vague predicates (Bylinina, 2013; Hirvonen, 2014; Saebø, 2009). Other than Stephenson (2007a), researchers discussing Subjective-Find have not talked about its embedding epistemic modals. The goal of this section is to fill this gap in the literature.

4.4.1 Normative modals

Bylinina (2013) does look at modal constructions. However, she is concerned with judge-dependence and with normative modals, such as the following:

(64) a. John finds this book too long (for me) to read in one day.
b. John finds this a long book (for me) to read in one day.
c. John finds this book a bit long (for me) to read in one day.

She analyzes these as lacking a 'judge argument' or an Experiencer argument, because the subject of *find* is different from the person doing the reading. I have already discussed that Bylinina confuses direct experience with the Experiencer argument, and we can see that these kinds of normative modals do require the subject to have some experience with the object, as we can verify:

(65) John finds this book too long for me to read in one day, #but he has never read it.

She does not have anything to say about the modal constructions we are going to look at in this section.

4.4.2 Modal adjectives

Find is bad with some modal verbs, while acceptable to others, as the following English and German examples show:

(66) a. ?I find it might rain.

b. I find it ought to rain by now.

c. I find we ought to visit Grandma.

d. Ich finde wir sollten die Oma besuchen
I find we ought the grandma visit
"I find we ought to visit Grandma."

e. ??Ich finde wir mussen die Oma besuchen
I find we ought the grandma visit
"I find we must visit Grandma."

For now, I note that this infelicity may be the reason why previous researchers have not concerned themselves with modal adjectives. However, COCA shows that $find\,X\,ADJ$ is used with the following modal adjectives:

(67) *find X acceptable, conceivable, convincing, credible, doubtful, implausible, impossible, inapplicable, inconceivable, inconsequential, likely, necessary, unconvincing, unfathomable, unsurprising, unlikely, unthinkable*

For *find/finds/found it [ADJ] that*, there were the following modal adjectives:

(68) *acceptable, apt, conceivable, credible, doubtful, implausible, inconceivable, likely, plausible, unimaginable, unlikely*

At first glance, the experiential presupposition seems not to cover modals. How can one use experience of an event in judging it likely or unlikely when in many cases the event has never happened and so could not have been experienced by anyone?

(69) Samir finds the explosion of the Yellowstone Caldera likely/unlikely.

Some scientists believe the volcanic caldera under Yellowstone National Park has the potential to erupt explosively, causing extreme devastation. This has never happened before, and therefore neither Samir nor anyone else could have experience of that event itself. However, we would judge the example felicitous.

We have several choices about how to proceed from here. One option is that there are actually two further types of Subjective-Find that are polysemic or homophonous. I've already shown that there are two verbs, of differing aspect features and behaviors, that both sound like "find" in English; why can't there be even more?

I reject this analysis for two reasons. One is the principle of Occam's Razor: I would like not to multiply my analyses beyond what is necessary. Two different *find*s are complicated enough; I do not want three (or more). Furthermore, if there were two different types of Subjective-Find, distinct according to the kind of argument they take in the small clause, one would predict that trying to conjoin both a predicate of personal taste and a modal adjective under the same *find* would be infelicitous, or at least a highly marked example of zeugma. This is the case for Subjective-Find versus Discovery-Find, as the following marked example shows:

(70) ??She found the book fascinating and under the stairs.

However, this is not the case when conjoining both a predicate of personal taste and a modal adjective under *find*:

(71) a. She found the story fascinating and plausible.

b. I find the idea of Trump being elected incredible and terrifying.

Although this may be stilted to some speakers, to my ear it is much, much better than (70). For this reason, I will assume that *I find this fascinating* and *I find this plausible* involve the

same verb, with the same lexical semantics.

We can further verify this with German, which uses *finden* for Subjective-Find and *befinden* for a judge *find* like *We find the defendant guilty*, and can combine both modal and evaluative predicates under *finden*[4]:

(72) a. Sie findet die Erklaerung faszinierend und plausibel.
she finds the account fascinating and plausible
"She finds the account fascinating and plausible."

b. Ich finde einen moglichen Wahlsieg Trumps unglaublich und/aber auch furchteinflssend.
I find a possible election Trump's improbable and/but also terrifying
"I find Trump possibly being elected improbable and/but also terrifying."

Another option, similar to the option that Giannakidou & Mari (Forthcoming 2016) pursue in their analysis of the subjunctive, is to have a two-tier system: if Subjective-Find embeds a predicate p that involves a subject having an experience, it invokes the direct experience presupposition; if it embeds a modal adjective, it does not. This would be similar to a ranked set of constraints in terms of Optimality Theory — if p is non-veridical (in Giannakidou's terms, modals are non-veridical), then the experiential constraint is ruled out.

Although there are certainly other constructions in many languages which require a two-tier analysis of this sort, I do not believe that we need it for this analysis. Instead, I am going to show that modal adjectives, at least, the kind that Subjective-Find takes, do involve a direct experience presupposition, so we do not need to treat them separately from predicates of personal taste and evaluative adjectives. It is just that the direct experience that they require, rather than being sensory, is of having experienced a mental process and associated feeling. Thus, they fall in the category of having speakers' differing perceptions of their evidential basis, as discussed in Chapter 3.

4. Thanks to Malte Willer for German judgements.

In the next few sections, I will take a closer look at what modal types can and cannot be embedded under Subjective-Find, but first, I will provide a review of the most widely-accepted model for modal expressions, and how it may divide modals into types.

4.4.3 Background: Types of modality

For a very basic analysis of modals, I rely mostly on Klecha's 2014 dissertation on scalarity and modality, as it gives a clear overview of recent work on the subject and makes a good case for distinct terminology. I choose to use the same terminology as Klecha for the building blocks of modals. For the basic theory, Klecha relies heavily on Kratzer (1981).

The most important distinction between modals that are embeddable under *find* and those that are not is the *conversational background*. As Kratzer (1981) explains, there are multiple conversational backgrounds that different modals take. These are implicit elements determined by context that fix the domain of a modal. I provide the four main types of conversational backgrounds as listed by Klecha (2014) and discuss how they apply to Subjective-Find.

Deontic conversational backgrounds fix a set of worlds consistent with rules or information supplied via some authority.

(73) (Klecha, 2014): *Deontic Conversational Background*

a. Erik has to submit his taxes soon (according to the law).

b. Erik can submit his taxes soon (according to the law).

Metaphysical conversational backgrounds fix a set of worlds consistent with a set of information that describes the real world at a given point in time.

(74) (Klecha, 2014): *Metaphysical Conversational Background*

a. Drew is bound to be here next weekend.

b. Drew might be here next weekend.

Teleological conversational backgrounds fix a set of worlds consistent with a set of information describing the achievement of a particular individual's goals. *Bouletic conversational backgrounds* fix a set of worlds consistent with a set of information describing a particular individual's desires; these are often conflated with teleological conversational backgrounds.

(75) (Klecha, 2014): *Teleological Conversational Background*

a. Felicia has to take the F Train (to get to Brooklyn).

b. Felicia could take the F Train (to get to Brooklyn).

Epistemic conversational backgrounds fix a set of worlds consistent with a set of information known by an agent or agents at a given time.

(76) (Klecha, 2014): *Epistemic Conversational Background*

a. Chuck must have left Chicago by now.

b. Chuck might have left Chicago by now.

Klecha says that there is much debate in the literature on the nature and force of ability modals, and does not go further into details about them. However, we need to deal with them in order to discuss examples such as *I find calculus impossible*, so I introduce them here.

(77) (Klecha, 2014): *Ability Conversational Background*

a. Graham can solve a crossword puzzle in under 15 minutes.

Given this background on conversational backgrounds, we can see some patterns in which modals can and cannot take Subjective-Find:

(78) Modals that can, in general, take Subjective-Find:

a. Ability modals: I find calculus impossible.

b. Epistemic modals: I find the Pope converting to Judaism impossible/inconceivable.

c. Teleological modals: Our boss finds hiring a new analyst necessary.

d. Bouletic modals: We find fixing the roof desirable.

The latter two examples may sound somewhat stilted, but are not as strange as the following examples:

(79) a. Deontic modals: ??I find filing my taxes mandatory.

b. Metaphysical modals: ??Peet finds Katie bound to be there next weekend.

Let us first address the type of modals that *cannot* be embedded under Subjective-Find.

4.4.4 Deontic and metaphysical modals

Deontic modals that are about laws/rules cannot be taken by *find*; however even with *mandatory*, indisputably a deontic modal, we can come up with more and less acceptable examples.

(80) a. #I find filing your taxes mandatory if you're earning income in the United States.

b. I find a visit to the Art Institute mandatory if you're in Chicago.

The first example, (a), is a ruling based on law. The second, (b), is a recommendation based on emotion, disguised as a ruling. If we look carefully at the two uses of *mandatory*, we see that the second case does have an experiential component that the first case does not.

(81) a. I've never filed my taxes, but filing taxes is mandatory.

b. ??I've never visited the Art Institute when I was in Chicago, but a visit to the Art Institute is mandatory.

If we allow that tax evaders can still say pragmatically felicitous things, then the second sentence sounds odder than the first. As a listener, I want to respond to sentence (81-b),

"On what grounds do you rule that, then?" If the speaker has never visited the Art Institute, based on what experience does she issue such a strong recommendation as to be rhetorically equivalent to law? This is not a problem with the first case: we know that the ordering source for *mandatory* comes from the U.S. legal tax code. But the ordering source from the second case comes from the speaker's emotion. Indeed, to me at least, (81-b) would sound strange if uttered by Star Trek's notoriously emotionless Mr. Spock. Given our knowledge of what art museums are like, we understand that such strong emotion is only tied to experience — it relies on the perception of the art and the categorization of it as emotionally moving.

We see a similar case of infelicity with metaphysical modal adjectives under *find*, which I also verify with an Aktionsart test using *for years*:

(82) a. #I find Drew bound to be here next weekend.

b. I find I'm bound to have a headache every time it rains.

c. For years, I found myself bound to have a headache every time it rained.

Recall that a metaphysical conversational background fixes a set of worlds consistent with the information available about the real world at a given point in time. The information about Drew's arrival is all information available about the real world, that all speakers would agree on. *Bound to* does not reflect a particular experience. However, if we succeed in coercing *bound to* to be about private experiences, such as headaches, we get a much-improved sentence.

So far we've found that in cases when we can coerce Subjective-Find into taking a deontic modal or a metaphysical modal, we also have an experiential presupposition available. The speaker must be basing her deontic ruling on an experience the Art Institute in Chicago. In cases where the ordering source is based on law, where no experience is necessary (such as with the tax evader's confession), *find* is infelicitous as well. Similarly, where the ordering source is based on common knowledge of the real world, and no experience is necessary, *find* is also infelicitous.

Let us now turn to the set of modals that *find* does take easily.

4.4.5 Ability modals

Almost any adjective ending in *-able/-ible* denoting ability (either directly or metaphorically) is good under *find*. As McNally & Stojanovic (2015) note, deverbal adjectives carry over their Experiencer thematic role semantics from their root verb, as shown by them also being able to take *to X/for X* prepositional phrases, which are overt Experiencer arguments:

(83) a. We enjoyed the experience.
b. The experience was enjoyable for us.
c. We found the experience enjoyable.

Knowing from Section 3 that Experiencer properties imply direct experience (but not vice versa), we can use this Experiencer property to explain why *find* embeds ability modals that end in *-able/-ible*, such as *conceivable/inconceivable*, *believable/unbelievable*, as well as *credible/incredible*, *possible/impossible* and *plausible/implausible* — the latter being originally derived from Latin verbs meaning *credere* "to believe", *posse* "to be able" and *plaudere* "to approve" which do not have a verb form in modern English. (Here I will discuss *possible/impossible* as an ability modal, such as *I find calculus impossible*, which does not imply that there is no world in which calculus exists, but implies that I have tried calculus and found it extremely difficult. I will discuss the epistemic meaning of *impossible*, e.g. *Perpetual motion is impossible* in the following subsection.)

We see that all of these *able/ible* adjective involve some sort of experience; one could not use *find* under an ability modal unless one has tried it, so it is infelicitous to say:

(84) ??Two-year-old Emma finds calculus impossible.

We see that in this way, *find* is different from *consider/believe*. To *find something incon-*

ceivable, one must have had the experience of trying to conceive it, that is, trying to think about it (and failing).

(85) a. #I find married bachelors inconceivable.

b. I believe/consider married bachelors to be inconceivable.

We confirm that there is a direct experience requirement for *inconceivable* here, the experience being that of having thought about it, or having attempted to conceive of it (and failed), when we look at examples where the speaker has not thought about it, as in (a).

(86) I find the Pope converting to Judaism inconceivable #but I've never thought about it.

However, we see that the speaker having thought about it does not rule out the felicity of *inconceivable* so there must be another factor at work:

(87) I've thought about your hypothesis, and find the scenario you describe inconceivable.

Inconceivable here does not mean that you were unable to conceive of it (imagine it); it must mean that you attempted to imagine it, and experienced a certain feeling (of ridiculousness, etc.).

We can compare this to *incomputable*, as in the following example:

(88) (?)I find 4/0 incomputable.

If this sentence is felicitous (it isn't to some speakers), it implies that the speaker did attempt to compute 4 divided by 0, in some fashion, and generated an error. (If the sentence is infelicitous, that would be because experience is not necessary, as 4/0 is equally incomputable for all speakers.) We can further verify this by trying to have the subject of *find* be something capable of computing, but lacking feeling, such as a computer. We see that *find*

X impossible/incomputable is infelicitous with a computer as the object:

(89) a. #My calculator finds 4/0 incomputable.

b. #My vintage 1990 computer finds loading YouTube videos impossible.

(89-b)

(b) is still infelicitous even if we made an attempt to load YouTube videos on the computer. We see that most adjectives ending in *-able/-ible* that *find* embeds, with the exception of the epistemic meaning of *possible/impossible*, do indeed have a direct experience presupposition — that of having at least tried to conceive, to judge and approve, or to compute, and having an associated feeling. So far, we haven't actually needed to invoke any additional semantic hardware to account for modal adjectives being embedded under *find* — the direct experience presupposition would do.

4.4.6 Teleological and bouletic modals

We can deal with bouletic modals like *desirable* in the same way we dealt with ability modals above: we understand that *X is desirable* means that one has tried to desire it, and was able to do so; thus, one has an experience of feeling desire for it and this fulfills the experiential component. Similarly, for teleological modals like *necessary*, we can understand that a speaker stating that X is necessary means that she has perceived or felt a need for X, and has experienced that feeling.

We can test this, as before, by trying to come up with cases where there was no experience of feeling a need or desire:

(90) a. I find us hiring an analyst necessary/unnecessary, #but I haven't looked at our needs.

b. I find a bigger house desirable/undesirable, #but I don't have any feelings on

the matter.

Both of these sentences are self-contradictory, either with or without negation. Even finding something undesirable, or unnecessary, still seems to involve the experience of feeling something about it. So having a non-experiencer subject should be odd, and indeed it is:

(91) a. ?The business improvement study found hiring an analyst necessary.

b. The business improvement study found the hiring of an analyst to be necessary. (Discovery-Find)

(92) a. ?The survey of our needs found a bigger house desirable.

b. The survey of our needs found that a bigger house was desirable for us. (Discovery-Find)

Thus, we can see that in the sense of the direct experience presupposition, teleological modals and bouletic modals are very similar to ability modals in that they also involve the experience of trying something. In this case, it is the experience of feeling either a need or a want, or feeling its opposite: a sense that this would be a wasted effort, or a bad possible world to be in.

We have thus solved the problem of teleological and bouletic modals quite easily by finding their direct experience presupposition. We must now turn to the most complicated case, that of epistemic modals.

4.4.7 Epistemic modals

Epistemic modals, especially those concerning likelihood, such as *likely, probable, possible, impossible*, are conventionally modelled using probability, and we can use probabilities with them in natural language, as in the following example:

(93) Serena is 80% likely to win the match.

Previous models of epistemic modals, such as (Lassiter, 2011) and (Klecha, 2014), have discussed the gradability and scalarity of modal expressions — one can speak of some event as being more likely or plausible than another event. However, they work with two different kinds of probability.

One is more familiar from mathematics: it is the kind based on previous cases, the probabilities of other correlated events, and possible outcomes. We will call this *objective probability*. Classic examples are the probability associated with coin flips, dice rolls, or draws of a card. We know that with a fair coin, the probability of heads is 50% or 1/2; with a fair die, the probability of rolling a 4 is 1/6; and the odds of drawing four-of-a-kind in a five-card hand of poker are 1 in 4,164.

However, there is also *subjective probability*, people's sense of how likely something is, based on holistic impressions. For example, a poker player may believe that a four-of-a-kind is very likely to come up, even though the mathematical odds are high against it. Lassiter (2011) discusses how people's subjective probability may not match up to objective probability, based on experiments conducted by Yalcin (2009) and others: given the same objective probability but differing alternatives or focus, speakers would have differing estimates of likelihood.

In the following subsection, I argue that subjective probability and objective probability can be distinguished by embedding under *find* and subjective probability carries an experiential dimension of feeling that is what *find* is selecting for. Given what we have already discussed concerning *find* and subjective predicates, it seems intuitive that *likely* based on subjective probability would be acceptable under *find* while *likely* based on objective probability would not. My judgement seems to agree with that hypothesis, as the following examples show:

(94) a. #Sue finds it likely that the next flip will be heads.

b. #Sue finds it likely that Hillary Clinton will win the 2016 US election.[5]

We can further zero in on the assumptions underlying (b). After all, an election forecast is a one based on multiple sources of data, both objective, such as polls, and subjective such as "the word on the street," "the sense of the mood of the nation," and the opinions and projections of pundits, experts and ordinary voters.

Suppose that in the autumn of 2016, Maria is a native of a foreign country, has never been to the US, and knows almost nothing about US politics. Told by her boss to prepare an urgent report and forecast on the 2016 US election, she consults poll numbers (and no other data). The polls report that Hillary Clinton has a 80% chance of winning the election.[6]

(95) Maria (assuming this is translated and her language has an equivalent verb to *find* that works the same way): #I find it likely that Hillary Clinton will win the election.

We would think Maria's utterance somewhat odd, even in a casual setting like telling her spouse about how her work went that day. All she has are numbers and statistics; *find* seems strange.

Compare this to Miriam, in the United States. Miriam does not read poll numbers at all. But she is constantly talking to anyone she meets about the election, watching it on TV (yet always skimming or not paying attention when polls are posted). She can't put a number to it, but she can say the following:

(96) Miriam: I find it likely that Hillary Clinton will win the election.

At least some native speakers I consulted found overall that Miriam's utterance was more felicitous than Maria's.[7]

5. Most of the examples in this subsection were written in the autumn of 2016.

6. This was written well before the election on November 8, 2016, where despite models based on poll data claiming she had a strong advantage, up to 80% at times, Hillary Clinton did not win.

7. Some speakers commented that both seemed odd, given their limited information. Others, though, did

The subjective probability that we base the claim *I find it likely* on is tied to an experiential state. In Miriam's case, she is *feeling* a likelihood. Skilled (or at least, enthusiastic) players of poker or other games of chance may similarly describe their experience of probabilities as a feeling. Speakers very commonly use the phrase *It feels likely*, as in the following results from a Google search:

(97) a. It feels likely that I'll kill myself in my forties.

b. If you and your spouse have had four daughters in a row, it feels likely that you'll have a boy the next time.

c. Vulnerability refers to whether a person feels likely to develop a condition over time.

d. Counter this concern by realizing that just because it feels likely does not mean it is likely. (From "Anxiety 9 to 5: How To Beat Worry, Stop Second-Guessing Yourself and Work With Confidence" by Laura Kase, 2006.)

All of these discuss feeling rather than probability; in the case of (b), the real probability of a son after four daughters should still be approximately 0.5, not higher. Example (d) particularly presupposes that a 'feeling of likelihood' exists and is common (at least, to anxiety sufferers). This also suggests that *likelihood* can be an experiential state and a perception that can differ among speakers.

Interim summary: I have claimed that epistemic modals, evaluating likelihood, do involve an experience, and need to involve an experience, if they are to be embedded under *find*. We did not need to involve any additional modal architecture; from all the evidence, *find* is about experiencing feelings or mental states for modal adjectives, and experiencing sensory input for predicates of personal taste that rely on the senses.

Let us turn to the final main section, given what we now know of what the experiential

judge that both contexts were fine. Others sensed differences in register between the two speakers' usages, claiming that either Maria or Miriam was speaking in a more formal tone, but their stated opinions of the difference were not consistent. The difference is nuanced.

requirements of *find* must be.

4.5 Refining the notion of experience

In this chapter, I argue that all we need to explain the distribution of *find* is a well-defined notion of experience as a necessary and sufficient condition. As we saw in Chapter 3, the requirement for direct experience in the case of subjective predicates (a presupposition in the case of *find*; a different kind of inference in the case of bare assertions of subjective predicates) stems from only direct experience being viewed as reliable. This is due to humans differing in how they perceive the evidence for subjective propositions, and also in how they sort and categorize such evidence in their minds, based on previous education, experience and preferences. This provides a good explanation for predicates of personal taste like *tasty* and *fun*, where people's senses of taste and preferences about taste, and likewise their preferences about enjoyable activities, differ yet are equally valid. In the *Critique of Judgement* (1790), Immanuel Kant claims that predicates of personal taste (in his view, 'agreeable' or 'pleasant' (*Angenehmen*)) induce a feeling of pleasure or pain; and that we cannot make judgements of beauty on the basis of hearsay, but must "submit the object to our own eyes," that is, directly experience it.

> "If we judge Objects merely according to concepts, then all representation of beauty is lost. Thus there can be no rule according to which any one is to be forced to recognise anything as beautiful. We cannot press [upon others] by the aid of any reasons or fundamental propositions our judgement that a coat, a house, or a flower is beautiful. We wish to submit the Object to our own eyes, as if the satisfaction in it depended on sensation; and yet if we then call the object beautiful, we believe that we speak with a universal voice, and we claim the assent of every one, although on the contrary all private sensation can only decide for the observer himself and his satisfaction." (J. H. Bernard translation,

Book 1, Moment 2, Part 8)

According to Kant, if we judged objects according to concepts, they can be communicated by hearsay; however, a judgement of beauty depends on something more than a conceptual notion, but on sensation as well. To find something a certain way, such as beautiful, it also seems to induce a feeling. He does not extend this, but it can be extended that since human sensations and what provokes them and whether they are treated as good or bad differ, only direct experience would do.

However, the claim that all interaction with predicates of personal taste relies on feelings of pleasure or pain exclusively seems a bit reductive. Some actions, such as consuming habanero peppers for a spicy-food aficionado, or riding a frightening roller coaster for a thrill seeker, would induce both. Other subjective predicates may induce different feelings; *thought-provoking*, for example, seems subjective, and may easily be embedded under *find* — *I found the documentary thought-provoking, but Ben says it's not, it's shallow* — but does not seem to be clearly tied to either pain or pleasure. However, we can conclude that subjective predicates involve feelings, and it is the direct experience of the object that involves these feelings.[8]

Modals pose a complication to this. To find something tasty, I must know what it is like to have tasted it, but to find an event likely, or a scenario that I've thought about plausible, what do I have to know? We must rule out the first approximation, that I must know what it is like to have that event occur, but since the event has not occurred, the only way I can

8. Whether conscious thoughts induce phenomenological experiences — the phenomenal intentionality theory or PIT — is a hotly debated topic in the phenomenology literature; please see Bourget & Mendelovici (2016) for an extensive overview of the main positions. Pitt (2004) draws a distinction between the *representative content* and *propositional content* of a thought. "By the representational content of a thought, I mean those of its properties by virtue of which it represents (expresses) the proposition it does. The proposition it represents, in contrast, I shall call its propositional content." (Pitt (2004), p. 6) He argues further in the same paragraph that "a thought's having a particular representational content *is* its having a particular phenomenology" (italics original). We can apply this definition of representational content to *find*, that *find* requires as input precisely the representational content of the clause it embeds. To find a food tasty, we must be aware of the properties by which we know it is tasty. We can only access these properties through the direct experience of taste. Similarly, to find that a chair is comfortable, we must be aware of the properties by which we know that it is comfortable — softness, shape, etc. Just knowing the propositional content "This chair is comfortable" because someone told us or we read it in a study is not enough.

do it is by imagining. If imagining an event like a papal conversion is sufficient experience for me to find the event likely or unlikely, why is it not sufficient for me to imagine myself skydiving (which I've never done) to state where I find it fun or not? To summarize this chapter, our view of experience must account for the following distinctions:

(98) a. I find (myself) going skydiving likely (though I have never skydived).
b. #I find skydiving fun (though I have never skydived).

This, I claim, is because all the evidence I need for a judgement of likelihood is imagining (and is augmented by sensation) and direct experience of that psychological state. We allow that a feeling of likelihood of the event is accessible from imagining the event. However, the properties by virtue of which *skydiving is fun* or *riding a roller coaster is fun* is true must be perceived — one must know *what it is like* to skydive in order to have an experience of skydiving.

As we have seen from the discussion of modality, those modals that work with *find* involve a feeling, something perceived by the senses. Something *feels* likely, or desirable, or difficult or impossible; on the other hand, when the rules are clearly stated, one does not depend on feeling whether something is mandatory. So it makes sense for *find* to be an intermediate attitude verb and perception verb, reflecting experiences of perception and sensation. Moreover, these experiences must be both necessary and sufficient.

Thus, that *find* embeds subjective predicates is actually *an epiphenomenon* of the fact that direct experience is necessary for subjective predicates, because evidence given by hearsay is unreliable. We thus predict that *find* would not embed certain predicates that are otherwise subjective, because for them direct experience is not necessary or not sufficient; e.g. those where the subjectivity lies in a purely categorization distinction, such as vague predicates like *tall*: *?I find her tall* is markedly worse than *I find her pleasant*.

Also, we predict that *find* would select for those properties that rely on perception and for perceptions that differ from person to person, such as taste or a sense of plausibility. If

the properties by virtue of which p is true are not all perceived by the senses (*gay, extinct*), we do not have *find*. If direct perception of the properties is not necessary (*wooden, green*), we do not have *find*.

And this last paragraph brings this chapter back to the central thesis of this dissertation. Direct perception of the evidence is not necessary when experience does not vary (or is assumed not to vary) from human being to human being, thus we can trust another person's perceptions to be the same.

So we have learned the following:

(99) *Find* embeds these predicates where direct perception by the senses of their evidential basis is both necessary and sufficient. This includes modals where the evidential basis is feeling.

We would predict from this that there would be certain predicates that are subjective but not good under *find*. Dimensional predicates would not be good under *find* because dimensions can be measured; thus directly experiencing them is unnecessary. However, it is possible to coerce an experience of height or duration, if one's experience does not coincide with the measurement:

(100) I found the flight from Chicago to London longer than the flight from Chicago to Hong Kong (even though the flight to London is shorter, because I travelled first class to Hong Kong and economy to London).

Evaluative predicates would be good under *find* if one is talking about the feelings that the object's properties of beauty, ugliness, intelligence, etc. induced in the subject. Predicates involving morality (*wrong, immoral*) would be appropriate under *find* to the extent that they are also coerced into being about experiencing a feeling of wrongness, rightness or outrage. Predicates involving modality would be appropriate under *find* insofar as the modals are

associated with feelings and sensations.

So the question we began with: "Why do subjective predicates get embedded under *find*?" gets the answer, "Because people feel differently about different things." People differ in what it is like for them to find something likely. A counterstance is simply a way of feeling differently. And *find* has the discourse effects of inarguability because people cannot override another person's senses.

4.6 Conclusion

In this chapter, I have looked at different explanations of why *find* (or the equivalent attitude verb in another language) selects for predicates that also occur in propositions inspiring faultless disagreement. After considering several options, I conclude that Stephenson (2007a) was right all along: *find* is simply an attitude verb with the evidential basis of direct experience. Most of the criticisms of her view arose from two gaps: confusing Subjective-Find with Discovery-Find; and not having a solid account of what "direct experience" means.

This chapter makes significant progress in bridging both gaps. By showing that Subjective-Find and Discovery-Find differ in their Aktionsart, I am able to show that Bouchard's (2012) criticisms of *find* being experiential are misguided; most of his apparent counterexamples are actually using Discovery-Find, and his own account does not explain legitimate modal uses of Subjective-Find. I show that an experiential account explains everything that a counterstance-based or argument/index-based account does, and also explains how *find* selects for modal adjectives, something that no previous account of *find* has addressed to my knowledge.

In addition, except for a few examples, I have stayed with English *find*, not extending the claims to German, Norwegian, French, Russian or other languages' equivalents. Probing at the differences between languages in what experiential basis is required will doubtless uncover interesting and unexpected areas to explore. For now, this chapter's merit is in paving the road to exploring what it is like to find something likely.

4.7 Appendix: some observations on *consider*

Although most of this chapter is about *find*, this appendix section discusses how my theory views *consider,* starting with the following puzzle:

(101) On a game show[9], a contestant is asked to choose between three doors. Behind one of the doors is a car, a valuable prize; behind the other two doors are useless booby prizes. The contestant has no conclusive evidence as to which door conceals the prize; she has to depend entirely on intuition and luck. She gives the matter some thought, and then says,

a. ??The car is behind door 2.

b. ??I find the car (to be) behind door 2.

c. ??I consider the car to be behind door 2.

d. I believe the car to be behind door 2.

Why are assertions (a)-(c) infelicitous? We can explain that *find* is infelicitous because the contestant has no direct experience with the space behind the door. We can explain that the bare assertion is infelicitous because the speaker has no real grounds for asserting this: she does not know for sure that the car is behind number 2, and is setting herself up to say something false. But why is *consider* wrong here, given that we know the speaker is claiming that only she believes this, and does not require others to believe this?

What previous work on *consider* has not taken into account is that there are actually two verbs [10] *consider NP XP*,[11] which differ in Aktionsart and Moorean effects. I will call them $consider_A$ (for Activity) and $consider_S$ (for Stative). I will demonstrate the difference

9. This is a setup known in probability research as the Monty Hall problem, after the host of the game show *Let's Make A Deal*.

10. This observation is due to Patrick Munoz, and I also thank him for much of the following discussion, although he has a different position on how to interpret *consider*.

11. Not to speak of the verb *consider NP*, e.g. *We are considering your application*, which is syntactically different still.

using the following examples:

(102) a. *consider*$_A$: We consider the ropes frictionless for the purposes of this physics problem.

b. *consider*$_S$: The Ancient Chinese considered themselves the center of the universe.

They pass all the Aktionsart tests in (Dowty, 1979) for statives versus activities, showing that *consider*$_A$ is an activity, and *consider*$_S$ is a stative:

(103) Only non-statives may occur in the progressive:

a. We are considering the ropes frictionless for this physics problem.

b. ??The Ancient Chinese were considering themselves the center of the universe.

(104) Only non-statives may occur in the complement of *force* or *persuade*:

a. We persuaded the professor to consider the ropes frictionless for this physics problem.

b. I forced/persuaded my parents to consider me an adult.

c. ??Their viewpoint forced/persuaded the Ancient Chinese to consider themselves the center of the universe.

(105) Only non-statives may be modified by *deliberately* or *carefully*:

a. We deliberately considered the ropes frictionless in this physics problem (so as to simplify the solution).

b. ??The Ancient Chinese deliberately considered themselves the center of the universe.

(106) Only non-statives may be used in the imperative:

a. Consider the ropes frictionless in this physics problem.

b. ??Consider yourself the center of the universe!

(107) Only non-statives occur in pseudo-cleft constructions:

a. What we did (to solve the problem) was consider the ropes frictionless.

b. ??What the Chinese did was consider themselves the center of the universe.

Moreover, they have different Moorean effects: the subject of $consider_A$ p does not have to believe p, while the subject of $consider_S$ p does have to believe p:

(108) a. We consider ropes frictionless for the purposes of this physics problem, even though we know they are not.

b. #The Chinese considered themselves the center of the universe, even though they knew they were not.

I conclude from this (as an initial step in the analysis) that $consider_A$ is synonymous with "treat as":

(109) We are treating the ropes as frictionless for this physics problem.

On the other hand, $consider_S$ requires at least "treat as" + "believe." We can test this by using a setup where one may believe some relationship between a noun phrase and a predicate without actually being able to *treat* the referent of the noun phrase accordingly. For example, this would be the case with a mythological creature, which subjects may have beliefs about, but be unable to interact with.

(110) a. The Ancient Egyptians believed the god Khnum to be the creator of the world on his potter's wheel.

b. ??The Ancient Egyptians treated the god Khnum as the creator of the world.

c. ?The Ancient Egyptians considered the god Khnum the creator of the world.

Thus we get the entailment that $consider_S$ must entail $consider_A$, but not vice versa.

This explains the puzzle of the Monty Hall problem:

(111) The contestant looks at the three doors and says:

a. #I consider the car to be behind Door 2.

We see that $consider_A$ is not licensed here: she does not have grounds to treat Door 2 as if there were a car behind it, since there might not be. Thus, $consider_S$ also fails, because even if she believes that there is a car behind Door 2, she cannot treat Door 2 as if there was.

Once we have the distinction between $consider_A$ and $consider_S$, with "treating as" being a part of the definitions of both, the effects of counterstances observed by Kennedy & Willer (2016) fall out of the usage of "treat as." Talking about treating X as Y is only felicitous when there is a choice about whether to treat X as Y or as Z. When there is no choice, *treat as* is infelicitous, so are both *consider* verbs, and counterstances are ruled out:

(112) a. ??I treat oxygen as breathable.

b. ??I consider oxygen breathable.

c. Let's treat ropes as frictionless for the purposes of this physics assignment.

d. ??Let's treat ropes as frictionless for our mountain climbing expedition.

Treating X as Y is certainly subject to coordination by stipulation:

(113) a. Let's agree to consider Lee vegetarian.

b. Let's agree to treat Lee as vegetarian.

We would even predict that this should be true of $consider_A$ cases but not $consider_S$ cases, as beliefs should not be so flexible as to be subject to stipulation. Stipulation would mean overruling a belief.

Consider has often been mentioned in the subjectivity literature as a difficult verb to define. This appendix, peripheral to my main focus on *find*, showed that perhaps part of the difficulty is that there are two different verbs. Splitting the problem should make it easier for further research.

CHAPTER 5

WHAT IT ALL MEANS

"But your *class [on aesthetics] — your class is a cult classic. I* love *your class. Your class is all about never ever saying* I like the tomato. *That's why so few people take it — I mean, no offence, it's a compliment. They can't handle the rigor of never saying* I like the tomato. *Because that's the worst thing you could ever do in your class, right? Because the tomato's not there to be* liked."

—Zadie Smith, *On Beauty* (2005)

5.1 Introduction

Recall Touchstone's story with which we began this dissertation:

(1) ROSALIND: Where learned you that oath, fool?

TOUCHSTONE: Of a certain knight that swore by his honor they were good pancakes, and swore by his honor the mustard was naught. Now, I'll stand to it, the pancakes were naught and the mustard was good, and yet was not the knight forsworn.

— William Shakespeare, *As You Like It*, Act I, Scene 2, lines 51-52.

Our goal in this dissertation was to establish what was it about *good* and *naught* in describing pancakes and mustard that allowed Touchstone and the knight to disagree on how it was appropriate. We saw that although there have been many intriguing attempts to model *how* it is that truth and falsity applies to such words, most of them took the distinction for granted. In Chapter 3 of this dissertation, I argued that the difference must lie in whether people are allowed to have differing perceptions and differing categorizations that apply to a particular word. That is, the knight's perceptions of whether the pancakes were *naught* meaning 'not tasty' were not the same as Touchstone's and could not be, while Touchstone and the knight should have a consensus on what it meant for the pancakes to be *naught*

meaning 'rotten.' In Chapter 4, I showed how this drives the experiential requirement of *find*.

This chapter allows us to sum up the contribution of the rest of the dissertation. I am not proposing a new theory to compete with expressivism, relativism, contextualism and its variants, metalinguistic negotiation, and other models. Indeed, in this chapter I will show that many of them can consistently incorporate my conclusions, and I can incorporate theirs.

This chapter is structured as follows. In the next section, I give an overview of the main conclusions we have reached in the previous chapters concerning the sources of subjectivity. Then I dedicate Section 5.3 to addressing a previous link from subjectivity to evidentiality (Korotkova, 2016) and showing how my proposal differs from it. Section 5.4 goes through the models I discussed in 1.4, and discusses how the incorporation of my proposal benefits (or hinders) them. In Section 5.5, I conclude the chapter, and the dissertation itself.

5.2 A review of our conclusions

In the previous two chapters, we looked at multiple characteristics that distinguish subjective propositions from objective ones. Here I provide a brief review of the features my model needs to account for. First, there are the well-known features of subjective propositions that I discussed in Chapter 1:

(2) Faultless disagreement:

Sue: Avocados are tasty.

Zoe: No, they're not! (neither is viewed as saying something false)

(3) Faultless disagreement in the comparative:

Sue: Avocados are tastier than carrots.

Zoe: No, carrots are tastier than avocados.

(4) Embedding under *find*:

Sue: I find avocados tasty/?vegetarian.

(5) Modification by *to X/for X* propositional phrases:

Sue: Avocados are tasty to me.

In Chapter 2, I added a number of pragmatic features of subjective propositions that have been less documented, which I review below:

(6) Disagreement as rejection from the common ground rather than as contradiction in truth value (Khoo & Knobe (Forthcoming 2016); see also Gutzmann (2016)):

Sue: Avocados are tasty, so let's have some with our toast.

Zoe: No, they're not! Let's not have any!

(7) Faulty disagreement patterns when one speaker has much weaker evidence than the other (Gunlogson & Carlson, 2016):

a. Sue: Avocados are tasty to me.

Zoe: #No, they're not.

b. Sue: I find avocados tasty.

Zoe: #No, you don't.

(8) The Acquaintance Inference (not a presupposition!) (Ninan, 2014):

Sue: Avocados are tasty, #though I've never tried one.

(9) The uptake requirement:

a. Sue: Avocados are tasty.

Zoe: [silence]

Sue: I take it you don't agree. / #Good, we're on the same page.

b. Sue: Avocados are tasty.

Chloe: Avocados are tasty; #I've never had one, but Sue said so.

We saw in Chapter 2 that the pragmatic features of subjective propositions boil down to the following:

(10) a. A subjective proposition requires a certain form of relevant direct knowledge (which form varies from proposition to proposition but is mostly experiential) for its assertion conditions.

b. If other speakers have the form of relevant direct knowledge that contradicts it, they are justified in rejecting the subjective proposition from updating the common ground.

c. If other speakers lack the direct knowledge, they cannot accept or reject the subjective proposition's update into the common ground.

d. That the speaker of the subjective proposition p believes p is automatically added to the common ground.

e. If the subjective proposition has a subject other than the speaker — *Avocados are tasty to Bob,* as uttered by Sue — then it is treated just like an objective proposition and the knowledge base for it does not have to be direct.

In Chapter 2, I noted that all of these look exactly like the behavior of objective predicates and their assertion conditions, except for one factor — that other speakers can have relevant direct knowledge that contradicts the asserted proposition, which knowledge is most often experiential, and which I am calling *evidence*.

(11) The requirements for faultless disagreement:

a. Speaker A felicitously asserts p_A from evidence e_A, and Speaker B felicitously asserts p_B from evidence e_B.

b. p_A and p_B both have the same references.

c. p_A and p_B are contradictory/not truth-compatible/cannot both update the common ground.

d. e_A and e_B are reliable: e_A should for all reasonable people lead to the belief p_A and e_B should for all reasonable people lead to the belief p_B.

e. e_A and e_B are considered equally reasonable.

As I argued in Chapter 2, I follow a model of (c) in which p_A and p_B cannot both update the common ground to interpret disagreement; I relied on the unprovability of the equality of e_A and e_B to determine faultlessness; even though I presented some ways in which e_A and e_B are not equal, I showed that the pragmatic requirements of a common ground require us to treat them as if they are. We ended up with the following principles in Chapter 3:

(12) **Principle 1:** The more perceptions differ, the more subjective the predicates based on those perceptions are.

with the corollary that predicates dealing with smell and taste, where genetic variability makes perceptions differ quite a bit, would thus be more subjective than predicates dealing with visual or auditory evidence.

(13) **Principle 2:** The more categorizations differ, the more subjective the predicates based on those categorizations are.

(14) **Principle 3:** The more categorizations are explicitly the same, the less subjective predicates based on those categorizations are. (Or: categorizations can be made the same if you name them right.)

Therefore, when perceptions converge and categorizations are shared, disagreements shouldn't be faultless.

A question may arise: are differing perceptions and/or categorizations a necessary condition for subjectivity? Suppose that humans invented some food, let's call it 'SuperChocolate,' that absolutely everyone unanimously takes pleasure from; it directly impacts the pleasure

center of the brain. Would we be justified in asserting that SuperChocolate is tasty? Would *tasty* still be subjective in those cases? Similarly, if everyone who had dissimilar tastes died in some epidemic, and the humans left on earth all had very similar tastes, would *tasty* still have subjective semantics?[1] I would expect that *tasty* would become an objective trait. However, it could have possibly dissimilar subjective readings, like *sweet* or *out of tune*, both facts about the food/sound that are verifiable using computer sensors but which provide different responses in people (to one person, a slightly out-of-tune note may not be that troubling, while to another, it would be a very unpleasant experience).

Knowing what we now know about the requirements for differing perceptions and differing categorizations, we can group predicates in order along a spectrum of subjectivity:

(15)

Predicate Type	**Differs in Perception**	**Differs in Categorization**
Objective	No	No
Dimensional	No	Yes
Modal	No	Yes
Moral	No or Partly	Yes
Evaluative	Partly	Yes
Personal Taste	Yes	Yes

Taking a closer look at *find* in Chapter 4, we saw that *find* is tightly linked to experience. In fact, even for modal expressions, which do not seem to have an experience component, the acceptability of *find* shows that this is true for some modals and not others, and it's precisely the modals that can rely on "gut feeling," such as epistemics and bouletics, that take *find*, and not deontic modals. In the case of subjective propositions embedded under *find*, we confirmed that experience is presupposed (and faultless disagreement is infelicitous). However, for bare subjective expressions, the link to experience, that is, the Acquaintance Inference, is more complex, and I will discuss it in the following section, as I address the

1. I am grateful to Itamar Francez and Patrick Munoz for discussion that led to this point.

relationship of subjectivity to evidentiality.

5.3 The relationship to evidentiality

In the previous section I have linked my model of the sources of subjectivity to the major subjectivity theories on the market. In this section, I want to look farther afield, and connect the discussion to the theories of evidentiality, since perception and categorization are categories of evidence.

A major dissertation in evidentiality that has already attempted to link it to subjectivity is that of Korotkova (2016), and here I will present her arguments that subjectivity drives evidentiality, and explain in detail how my model shows that the situation is actually the other way around.

Evidentials are morphemes or words that show the source of the speaker's knowledge for the proposition she asserts, be it direct, visual, inferred, hearsay, etc. In about 25% of the world's languages they are nearly obligatory (though never completely obligatory, McCready & Ogata (2007) affirm). They may appear as verb tenses (Bulgarian), particles (Japanese), adverbs (English non-obligatory evidentials like *apparently, reportedly*), or attached morphemes (Turkish, Quechua).

During the period between the work of Faller (2002) and the work of Korotkova (2016), evidentials in human languages were generally assumed to come in two different types:

(16) *Epistemic evidentials*

a. Occur in St'at'imcets (Matthewson *et al.* , 2007), Japanese (McCready & Ogata, 2007), Bulgarian (Izvorski, 1997)...

b. Behave like epistemic modals

c. Act at the propositional level

d. Do not apply to a proposition p if p is known to be true or false (compare epistemic modal *#The keys must be in my pocket, but they are not*)

e. Are embeddable

(17) *Illocutionary evidentials*

a. Occur in Cuzco Quechua (Faller, 2002), Cheyenne (Murray, 2011)

b. Behave like English parentheticals such as *they say, I find, I gather...* (Murray, 2011)

c. Act at the speech-act level

d. Do *not* carry a commitment to the truth of p (Faller has a Quechua example: *They left me a lot of money REPORTATIVE, but as you can see, VISUAL they did not leave me a cent*)

e. Appear to not be embeddable, although Mbya seems to be a counterexample

f. May appear in questions and imperatives, at least in Cheyenne

The main upshot of Korotkova's 2016 dissertation is that she shows across languages that apparent restrictions on evidential behavior that led to this pattern were not actually based on the semantics at all, as scholars such as Faller and Murray had assumed. She shows that the reason the evidentials in Quechua or Cheyenne were not embeddable were due to syntactic restrictions. She thus shows that evidentials across human languages are all semantically of one type, although she refrains from committing to whether that type is epistemic modal or modifying speech acts. This is an important and commendable result.

The topic of subjectivity, however, relates to the second half of her dissertation, where she argues that evidential content actually reflects subjective content. Previously, Murray (2011) had used data from contradictions and arguments to analyze evidentials in Cheyenne as having not-at-issue content. That is, just like appositives (Potts, 2002b,a) and the emotional content of epithets like *that bastard* (Potts, 2007), the evidential content of a statement cannot be disagreed with and is difficult to question. The Cuzco Quechua evidentials studied by Faller (2002) follow a similar pattern of non-challengeability, as illustrated by Korotkova (2016) based on (Faller, 2002):

(18) a. Ines-qa qaynunchay ñaña-ta=**n** watuku-rqa-n
Ines-TOP yesterday sister-3-ACC=DIR visit-PST-n
"Ines visited her sister yesterday, I saw."

b. Mana=n chiqaq-chu
not=DIR true-NEG
"That's not true."

(i) $=\neg$ [Ines visited her sister] p

(ii) $\neq \neg$[You saw that Ines visited her sister] *evidential requirement*

– Korotkova (2016) ex. 137, based on Faller (2002), ex. 116-117b.

That is, telling the speaker *That's not true* can challenge the proposition itself, but not the speaker's evidence for it.

Korotkova adds, "The pattern illustrated in (137) is observed in many other geographically unrelated languages, e.g. in Bulgarian (South Slavic; Izvorski 1997), Cheyenne (Algonquian; Murray 2014), Georgian (South Caucasian; Korotkova 2012), German (Germanic; Faller 2007) and St'át'imcets (Salish; Matthewson et al. 2007). Based on the data from available studies of evidentiality within formal semantics, the non-challengeability of the ER is a *universal* property of grammatical evidentials [italics original]. It does not depend on the morphosyntactic category of evidentials in a given language (unlike e.g. syntactic embeddability; Chapter 3) or on the type of information source (e.g. hearsay evidentials are non-challengeable just as direct evidentials are)." (p. 96)

Korotkova argues that the evidential requirement is not not-at-issue. Appositives, accepted to be not-at-issue, cannot be directly challenged with *That's not true*, as in the following example:

(19) a. Ortcutt, a spy, lost his passport.

b. That's not true.

(i) $=\neg$ [Ortcutt lost his passport] *main sentence*

(ii) $\neq \neg$[Ortcutt is a spy] *appositive*

However, evidentials follow a stronger pattern. Not-at-issue content like the projected content of appositives or famous presupposition projectors like *too* or *again* can be contradicted with *You are mistaken* even if they cannot be contradicted with *That's not true*. I give an example in English that Korotkova also confirms works for Turkish and Bulgarian:

(20) a. Alice: California, the largest state, legalized marijuana.

b. Bob: That's not true.

(i) = ¬ [California didn't legalize marijuana] *main sentence*

(ii) ≠ ¬[California is not the largest state] *appositive*

c. Bob: You're mistaken.

(i) = ¬ [California didn't legalize marijuana] *main sentence*

(ii) = ¬[California is not the largest state] *appositive*

(if there is a continuation such as *Alaska is the largest state*)

Korotkova argues that instead, evidential content patterns with first-person pain reports and hope ascriptions: to A's statement *I have an awful headache* or *I hope that Trump will win*[2], neither *You are mistaken* nor *That's not true* are felicitous as a response from B. Both responses sound odd, assuming that B knows better than A does herself how her body feels or what her hopes are. A has exclusive access to her pain receptors and to her hopes. Similarly, says Korotkova, "I argue that *all* evidentials denote experiences to which individuals have exclusive access, regardless of the source" (p. 109, italics original).

Korotkova groups the category of Subjective Content as including first-person (A) attitude reports (*I hope*), (B) taste ascriptions (*It tastes good to me*), (C) psych verbs (*I am excited*) and (D) statements about pain (*It hurts*). Evidentials, she argues, are in the same group due to being non-challengeable under any circumstances.

Readers of Chapter 2 would recognize that the non-challengeability of subjective content is treated as *faulty disagreement* by Gunlogson & Carlson (2016). Both Korotkova and

2. Korotkova's dissertation was written and defended prior to the 2016 U.S. election.

Gunlogson and Carlson analyze this as access to private information: in the case of predicates of personal taste, only the speaker has access to her taste. In the case of evidentials, only the speaker has access to her evidence.

However, to fit into my theory, I must argue that in uniting subjectivity with evidentiality, Korotkova is mistaken about which is the source and which is the symptom when looking at privileged access to information versus difference in the information available.

First, the taste ascriptions that Korotkova compares to evidential content are those that have *to me* — e.g, *It tastes good to me*. As we saw in previous chapters, these behave quite differently from bare taste assertions like *It is tasty*, in two major ways: (1) the experience involved is presupposed for predicates with *to me* (and *find*), while as Ninan (2014) shows, for bare predicates the experience involved is not presupposed; and (2) predicates with *to me* (and *find*) are indeed non-challengeable, while bare predicates are challengeable and that is what faultless disagreement is.

Thus evidentials and subjective predicates modified with *to me* and *find* do share features of non-challengeability, but that does not mean that evidentials should fall under subjective content.

A crucial argument against uniting predicates of personal taste and evidentials is that one's position on predicates of personal taste can change. Although wine or beer may at first taste sour and unpleasant, a wine-tasting course or detailed guidance from a beer expert can lead one to change one's position on it. As we remember, when tastes change (our perceptions or categorizations change), we feel that we should retract our earlier statements, as illustrated in MacFarlane's Fish Sticks example I will reproduce once more:

(21) Fish Sticks Example:

When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I've changed my mind about the tastiness of fish sticks. So, if someone said, "But

> you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I didn't mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane (2014) pp. 13-14)

To MacFarlane, even predicates modified by *to me* still require a retraction. But predicates modified by *to me* still may not allow a retraction.

On the other hand, evidential content cannot be retracted: once one acquires hearsay evidence, it will always be hearsay evidence. It can be overridden by stronger evidence, or both could be acknowledged, as in the Pedro's Hen example used by Faller (2002): Pedro, a farmer, goes to his henhouse to see one of his hens missing and a trail of blood and feathers. Pedro knows that foxes sometimes steal hens, so he conjectures that a fox took his hen. Then he encounters his neighbour who tells him that she saw a puma (a much rarer predator) take his hen.

(22) a. Atuq-chá wallpa-y-ta apa-rqa-n
Fox-CNJ hen-1-ACC take-PST1-3
'I concluded that a fox took my hen.'

b. Ichaqa wasi masi-y riku-sqa puma-s apa-n-man ka-rqa
But house friend-1 see-SQA puma-REP take-PST.3-IRR be-PST.3
'But my neighbour (lit. house friend) saw [it] and [she says] a puma took it.'

Thus, Pedro is judging by context as to which evidence should be considered more reliable — in this case, the inferential evidence and the hearsay evidence seem equally good (on the one hand, the neighbour is trustworthy; on the other, pumas are rare in this region and given the evidence, a fox is more likely). Thus we see that evidence *source* is distinct from evidence *reliability*, except that direct evidence generally takes priority (unless the speaker is unable to trust her own senses due to illness or impairment).

On the other hand, although Pedro may decide between inferential and hearsay evidence, it seems very odd for him to later say, "I said a puma took my hen, I heard, but I was wrong, I didn't actually hear that" or "I said a fox took my hen, I concluded, but I didn't actually conclude that." If he utters this, he would be taken as having lied earlier. This does not parallel with retraction of taste claims, where it does not at all seem like the child MacFarlane was lying if he later retracts his statement.

Korotkova argues that evidential content is subjective content. But why does one need to mark subjective content as such? Not just because one has private access to it, but because therefore *it is different for different people* (and for the same person at different times in the case of inferential evidence or tastes). One has just as much private access to direct evidence, because visual evidence still goes into the brain. However, people generally agree on what they see at the same time and location. Thus we mark the evidential content that may vary from person to person, including hearsay. Therefore, the factor of differing direct evidence seems to lead to subjective content, rather than the factor of subjective content leading to differing evidence.

Korotkova's argument leaves what subjective content *is* undefined, other than that it is private. However, that doesn't distinguish it from other sensory information. My analysis of it as content where the evidence can be different for different people clarifies the distinction between what it is and what it is not.

5.4 The effect on previous models

This section echoes the structure of Section 1.4. There, I went through a summary of the prevailing theories about how to model subjective propositions, and their ideas on what causes subjectivity. Here, I go through the theories in the same order, and explain how my evidential model of subjectivity would fit into each of them.

It is not the goal of this section, or of the entire dissertation, to argue which model would therefore be 'best' or 'right.' In Chapter 3, I have shown human tendencies to differ in

perception that, in my opinion, must be incorporated into models of language. That such factors as specific anosmias exist, causing human perceptions to differ, is a fact. That people tend to analyze the same situation differently depending on their experiences and priorities is also a fact. If models of language incorporate this, they are more realistic and plausible. That said, I do show that some of the models on the current market would need to make a lot more adjustments than others do.

5.4.1 Expressivism

As I said in the section on expressivism in Chapter 1 (1.4.2), I am giving modern expressivism in the vein of Gutzmann (2016) and Silk (2016) a lot of attention because their method of bypassing the Frege-Geach problem allows their models to be used to explain the results of Khoo & Knobe (Forthcoming 2016). Essentially, to recap, traditional expressivism held that subjective propositions are a different type of speech act than standard truth-conditional propositions. In the view of Gutzmann, they express emotional content. In the view of Khoo and Knobe, disagreement with subjective propositions (and indeed, with objective propositions) has the pragmatic function of rejecting update to the common ground rather than having opposing truth values. Gutzmann shows that disagreement and modus ponens can function independently of truth conditions — that is, one can disagree or construct modus ponens reasoning with speech acts that a priori do not have truth conditions, such as promises, compliments, and commands — and thus disagreement about subjective propositions is not a valid reason to assume that they are necessarily truth-conditional.

However, if we look at the principles above, we see that the differences between subjective and objective propositions are gradable, and are predicted to be so. Perceptions may differ radically or slightly between agents: e.g. one agent could completely lack a sense of smell, or just have trouble perceiving vanillin, compared to another. Categorizations, similarly, should and do differ in a gradual way: an agent may completely not notice a factor, or notice it but interpret it differently. Thus subjectivity should be a spectrum.

As far as I understand the speech act theory of Austin (1962) and Searle (1962), speech act categories are not predicted to be gradable; a statement should not be half compliment and half assertion, or a bit of a promise but mostly a threat. Judging by this, the expressivist model, which assumes that a statement is either subjective or non-subjective, should contradict the results of incorporating differing perceptions.

However, the pragmatic and illocutionary layer of a statement functions as a separate layer from the semantics. A statement could be a bare declarative or question, and depending on the context, may simultaneously be an assertion, compliment, promise or threat. For example, I can illustrate this with the same declarative statement, which the reader is free to imagine said in different tones of voice to make the pragmatic force clearer:

(23) a. Batman will return. (declarative assertion of a future fact)
b. Batman will return. (a promise to Batman's fans)
c. Batman will return. (a threat to Batman's enemies)
d. (in a mocking voice) Batman will return. (an insult to Batman's fans)

(24) a. The velociraptor is hungry. (a declarative observation linking the observed behavior of the velociraptor to its state)
b. The velociraptor is hungry. (a request to feed it)
c. The velociraptor is hungry. (a warning to the addressee)
d. The velociraptor is hungry. (a threat to throw the addressee to it)

All of these in each group are semantically the same statement, using the same words composed in the same way, and with the same reference (assuming a unique Batman or a single salient velociraptor). But they would have very different pragmatic functions.

Similarly, we can accept the expressivist model of Gutzmann (2016) by assuming that it works solely on the pragmatic layer, which varies independently of the semantic layer. Thus, subjective predicates, by expressing differences in perception and categorization, do have a

separate role in speech acts than objective predicates do. Instead of expressing *The world is this way*, they express *My perceptions and judgements are this way, are yours?*.

Thus we get the uptake requirement noted by Beltrama (2016) — that subjective expressions carry an inquisitive character. We can use the tools of expressivism to account for their pragmatic function. As I argued in Chapter 2, we can adapt Beltrama's model in order to make the common-ground update of an expressive proposition be the join of the views of all speakers. Uttering the negation of that proposition blocks common-ground update.

5.4.2 Relativism

What does a judge index actually entail? How is one relativized to a speaker or other person (or cat)? And how do words acquire one?

The evidence model allows us to answer those questions. If I assume a relativist model, I would assume that every predicate has a default trivial judge index. But certain predicates have judge indices that differ from person to person. These would be the predicates that depend on differing perceptions or differing categorizations. Other predicates — the conventionally objective ones such as *wooden* or *prime number* — would have a judge index, but as it is the same for everyone, it is trivial. And still others — such as color terms, for which human perceptions do vary though not as much as for taste — would have judge indices that resolve to values that are *mostly* the same for most speakers, or *converge* to be the same.

How do we explain MacFarlane's (2014) context of utterance and context of assessment indices? He argues for these based on the need we feel for retracting our statements in cases similar to his Fish Sticks Example and Coffee Shop Example, which I reproduce again below.

(25) Fish Sticks Example:

When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I've changed my mind about the tastiness of fish sticks. So, if someone said, "But

you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I didn't mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane (2014) pp. 13-14)

(26) Coffee Shop Example:

Suppose you are standing in a coffee line, and you overhear Sally and George discussing a mutual acquaintance, Joe:

SALLY: Joe might be in China. I didnt see him today.

GEORGE: No, he can't be in China. He doesn't have his visa yet.

SALLY: Oh, really? Then I guess I was wrong.

SALLY: Oh, really? #Still, I was right when I said "Joe might be in China," and I stand by my claim. (MacFarlane (2014), p.240)

MacFarlane incorporated the contexts of utterance and assessment in order to account for differences in taste or knowledge state that require retraction even for the same person, once their knowledge state changes (such as learning about more quality seafood than fish sticks, or ruling out the possible world where Joe is in China as impossible).

What has changed between utterance and assessment in MacFarlane's examples, I argue, is the agent's categorization (unless the child MacFarlane's taste buds are physiologically different from the adult MacFarlane's). Thus the judge indices diverge from each other, because the categorizations of the judge have diverged.

If we allow a judge index for all things, however, relativism retains the problem pointed out by Pearson (2012): having more indices than world and time breaks the tidy semantic model of Kratzerian semantics. As I quoted Pearson (2012) in Chapter 1, "Such a move complicates the semantics by requiring that each predicate takes an additional individual

argument as well as a world (and time) argument, with implications across the inventory of semantic types: the intension of a one-place predicate like *table* or *run* is now of type $< s, < e, < e, t >>>$, that of a two-place predicate like *love* is of type $< s, < e, < e, < e, t >>>>$ and so on." (Pearson (2012), p. 2)

However, Lasersohn (2017) has already considered this and adapted the semantic model accordingly. He argues that there must be a hidden variable in order for *to X* or *for X* to attach correctly, although he does not explain why, if there is such a hidden variable, subjectivity can still shift.

I can conclude from this that if we incorporate into relativism a predictive method for when judge indices converge and when they diverge (while allowing that all predicates have judge indices, but some would converge for all judges), there is no reason why relativism is not an appropriate framework for modelling subjective predicates. This dissertation's results, incorporated into relativism, would solve a pressing problem that Lasersohn acknowledges — that he takes the difference between predicates of personal taste and other predicates as a given.

5.4.3 Relativism and modality

A further question remains: can we explain the problem that Stephenson (2007a,b) wrestles with, that modal expressions and predicates of personal taste seem to both be context-sensitive and have similar pragmatic values, but only the latter and not the former provides exocentric readings? I reproduce Stephenson's examples that show this:

(27) A: How's that new cat food you bought?
B: It must be tasty, because the cat ate almost all of it. (Exocentric reading where the judge of taste is the cat)

(28) Mary: Wow, the dog really likes the dog food you're feeding him.
Sam: (#)Yeah, I think it might be table scraps. (No exocentric reading where the

judge of possibility is the dog)

However, Stephenson may have judged too hastily that modal expressions do not provide exocentric readings. Consider Egan's (2007) James Bond example:

(29) James Bond has just returned to London after a long day of infiltrating SPECTRE's secret base in the Swiss Alps, planting a bug in the main conference room, and slipping out by night after leaving persuasive but misleading evidence of his presence in Zurich. Sipping martinis in MI6 headquarters while monitoring the newly placed bug, Bond and his CIA colleague Felix Leiter overhear a conversation between Blofeld and his second in command, Number 2. Some facts to bear in mind about everyone's epistemic situation: Both Bond and Felix, obviously, know that Bond is in London. It is compatible with everything that Blofeld and Number 2 know that Bond is in Zurich. In the course of a discussion of the pros and cons of various nefarious plans, Number 2 says to Blofeld,

a. Number 2: Bond might be in Zurich.

(i) Blofeld (to Number 2): That's true.

(ii) Leiter (to Bond): ??That's true.

Egan uses this case to argue for a relativist model for the truth of modals. Note, however, that Leiter could ironically answer:

(30) Number 2: Bond might be in Zurich.

a. Leiter (to Bond): Yeah, you might be.

It requires a specific intonation, but is an acceptable conversational move. We take others' perspective as a matter of course, involving fiction, drama, irony and sarcasm and other reasons. This does not even require special grammatical moves.

Similarly, von Fintel & Gillies (2011) have the Mastermind example, involving a popular game where one player creates an arrangement of colored pegs behind a screen, and the other player tries to guess the arrangement.

(31) Pascal and Mordecai are (still) playing Mastermind. After some rounds where Mordecai gives Pascal hints about the solution, Pascal asks whether there might be two reds. Mordecai answers:

a. That's right. There might be.

As von Fintel and Gillies put it, "He can answer this way even if he knows there aren't two reds. As far as the norms of assertion go, it's as if he had uttered an explicit claim about Pascal's evidence." Von Fintel and Gillies argue further that there is no need for relativism (they refer to relativists as "CIA agents" for "contexts, indices and agents") in the analysis of epistemic modals.

We can return to Stephenson's example, which I will reproduce:

(32) Mary: Wow, the dog really likes the dog food you're feeding him.
Sam: (#)Yeah, I think it might be table scraps. (No exocentric reading where the judge of possibility is the dog)

If Stephenson's example had left the modal unembedded, then exocentric readings are available. Sam *can* say ironically, assuming the dog's perspective:

(33) Mary: Wow, the dog really likes the dog food you're feeding him.
Sam: Yeah, there *might* be table scraps in there.

Sam ironically assumes the dog's perspective, even though he is aware that there are no table scraps in the dog food, but the dog generally acts enthusiastically in the hope of getting some. This is a perfectly natural example of an exocentric modal. Thus Stephenson's error was

in inserting *I think* into the example, which of course anchors it to Sam's thoughts and removes the possibility of the dog's perspective. As Korotkova (2016) argues, all embedded modals under attitude verbs shift to the perspective of the embedder, as demonstrated in the following example (her ex. 218, p. 143):

(34) Scylla thinks [that Odysseus' ship *might* pass Charybdis]...

a. NON-SHIFTED, SPEAKER-ORIENTED: ...#but Scylla is sure it would pass.

b. SHIFTED, SUBJECT-ORIENTED: ...but I am sure it would pass.

In this case, the judge of *might* has to be the thinker, Scylla, and not the speaker. Similarly, we see that in the earlier, *Yeah, I think it might be table scraps*, the judge of *might* has to be the subject of the embedded clause, *I*.

5.4.4 Outlook-based semantics

Coppock's (2015) model of outlook-based semantics is essentially a notational modification of both the simple relativism as proposed by Kölbel (2002) and the judge-index relativism proposed by Lasersohn (2005). She simplifies the model by proposing *outlooks* instead of worlds and judges. As I explained in Chapter 1, outlooks have most of the same properties of possible worlds: they determine the extensions of predicates and relations, including both objective and subjective (what Coppock calls *discretionary*) predicates. Coppock still includes sets of worlds in her overall semantics, but only for objective predicates (e.g. *red, wooden, prime number*). An objective proposition describes a set of worlds, as well as a set of outlooks, while a subjective proposition describes only a set of outlooks.

However, as I also pointed out in Chapter 1, Coppock still does not draw a firm distinction between what exactly makes a subjective or discretionary predicate. My proposal from Chapter 3 fills that gap. If we adopt the view of outlook-based semantics, outlooks are determined by situations where different people would have different perceptions forming

their evidence for the extension of the predicate, or different — indeed, discretionary — ways of categorizing the world accordingly. Worlds would be established when there is an objective consensus, either based on quantitative measurements or on general consensus like color. Outlooks would be left to most other things.

Thus, although outlooks provide an elegant simplification to the semantic notation, and solve some of the concerns that Pearson (2013) has concerning the introduction of judges, I do not see that they provide strong advantages concerning the notion of truth.

5.4.5 Contextualism

As we recall from Chapter 1, contextualist theories argue that subjective predicates have a judge argument in their Kaplanian content, which is what determines the variability in truth. Stojanovic (2007) argues that many forms of contextualism end up being notational variants on relativism. She argues that this is because there is no such thing as faultless disagreement: disagreements are either genuine disagreements or they are misunderstandings. From this, she presents a function translating relativist semantics to contextualist semantics as follows:

(35) To prove the equivalence between CS and RS, I define a bi-directional translation procedure T between the two formal languages for which the following holds. Let S_c and S_r be respectively sentences in the languages of CS and of RS, let f_1, f_2 be assignments of values to free variables, and let w be a world of evaluation and u a judge. Then:

a. S_r is true with respect to f_1, w and u iff $T(S_r)$ is true with respect to f_1, T and w, where assignment f_1^T is defined in terms of f_1 and u.

b. S_r is true with respect to f_2 and w iff $T(S_c)$ is true with respect to f_2, w and u^T, where u^T is a judge value obtained directly from f_2.

(Stojanovic (2007), p. 700)

In the formal language that Stojanovic uses, there is a variable x_T that is used for the implicit judge argument associated with taste predicates.

However, contextualism is the major theory that runs into the most trouble with the view of subjectivity that I have argued for in the last few chapters. This is because in my view of things, subjectivity is a flexible notion: whether people perceive and categorize notions differently depends on the people involved (and on whether we accept that their categorizing is valid, or assume that they have a mental illness or sensory disturbance or are insufficiently educated). Thus if we assume a fixed argument associated with predicates of taste (and aesthetic predicates, modal predicates, vague predicates, etc.), we must explain how that argument grows more or less salient as the perceptions of the agents involved converge or diverge.

A second problem with the contextualist reasoning is that there is no other consistent syntactic evidence for the judge argument. There is no syntactic difference between *Obama won the debate (at a presidential election debate)* and *Sharon won the debate (at the World Universities Debating Championships)*, yet one is subjective and one is not. Lasersohn may argue that *to X* attachments provide such evidence, but *?Obama won the debate to me* sounds very strange, despite the faultlessness of the disagreement.

Because of the variability in how subjective concepts are, and the lack of syntactic difference between subjective and non-subjective predicates, we are faced with two options: (a) all predicates have a judge argument and for the objective ones, it is trivial: assigning it to any speaker for predicates such as *wooden* or *prime* would not change the truth conditions of the proposition; (b) none of them do, and the judge argument does not actually exist.

In the previous section, I suggested that relativism adopt the *all* scenario and have every predicate have a judge index, but where perceptions and categorizations do not significantly vary, it is trivially saturated. I would suggest the same for the contextualist model as expounded by Stojanovic and her fellow contextualists.

If we accept that this is the case, then we have the question of whether we identify the

contextualist argument with the Experiencer argument. Although Bylinina (2013, 2016) has argued that this is the case, she has also shown that aesthetic and evaluative predicates, as well as modals, do not have Experiencer arguments in languages that make them explicit in the syntax. Thus the judge = Experiencer reasoning becomes tenuous; at least, it cannot apply to all categories of subjective predicates.

5.4.6 Judge-free contextualism

If Moltmann (2010) and Pearson (2013) wish to argue that certain predicates have a first-person generic that is filled by the speaker and those he or she empathizes with, I provide an analysis for what the empathize-with relation would actually entail. By my reasoning, to empathize with an agent would mean that one assumes that the agent shares the same perceptions and the same categorizations as oneself, or at least, would react in a similar way to similar stimuli, as in the case of cat food.

Grinsell (2012, 2017) provides a firm mathematical and economic foundation using social choice theory as to why vagueness exists. In his 2012 paper and especially in his 2017 dissertation, he ties vague adjective standards and multidimensional adjectives to the problems of indecision as to whether a standard applies or what is optimal. He draws a rigorous analogy between proposition models and adjective spaces, and the behavior of voters as analyzed by the Marquis de Condorcet and Kenneth Arrow.

His analysis of collective choice shows that as long as preferences differ (and satisfy the conditions of unanimity, anonymity and continuity, which he explains) there will not be a fair aggregation function. Taking for granted that voters/judges/speakers will be diverse and have different preferences, Grinsell shows that vagueness effects *have to* arise, mathematically. For a full proof, I refer readers to (Grinsell, 2017), chapter 4.

An important feature of Grinsell's reasoning, both in the voting paradigms that he bases his analysis on and in his own analysis of vague and multidimensional adjectives, is *anonymity*. That is the notion that no one's preferences stand out or are more important

than any other's (or in the case of a voting paradigm, that all votes count equally). We can see this as applying to faultless disagreement. The reason that the disagreement is seen as faultless is precisely because of anonymity. Yet Arrow's Theorem that Grinsell cites provides a way out of cycling.

My proposal supplements Grinsell because it explains *why* preferences would differ. Thus it provides a more complete picture from which Grinsell's reasoning naturally follows.

5.4.7 Metalinguistic negotiation

As we have seen earlier, metalinguistic negotiation (Barker, 2002, 2012; Sundell, 2010; Plunkett & Sundell, 2013) is essentially a negotiation about categorizations. Since metalinguistic negotiation is about using language to establish norms, it is about establishing what words cover what categories. For example, Ludlow's (2008) famous faultless disagreement over whether the racehorse Secretariat would count as an athlete (whether the category of *athlete* includes the requirement *human*) is a predictable type of disagreement.

We would thus predict that metalinguistic negotiation would be used for those words and propositions where categorizations would differ — and crucially, we would predict that it would *fail* for words and propositions where perceptions themselves would differ. Differing perceptions cannot be negotiated with. There is no point in Zoe persuading Sue to believe that vanilla smells wonderful if Sue congenitally can't smell vanilla, or in Zoe persuading the color-blind Bob that the tulips are a gorgeous shade of red when to him there is very little difference between the flowers and their leaves. If Bob or Sue concede to Zoe's norms, they are being insincere.

However, just because metalinguistic negotiation would not succeed in cases of perceptual differences does not mean that following the protocols of metalinguistic negotiation would not serve a purpose there anyway. As Beltrama (2016) argues, we can use subjective propositions to elicit information, including differences in perception:

(36) a. Greg: The movie was awesome.

Fred: No, I didn't like it.

Greg: Ahah, that's hilarious! I knew you were the kind of person that doesn't like these movies.

b. Greg: The movie started at 8 pm.

Fred: No, it didn't. It started at 9 pm!

Greg: # Oh, that's interesting! We have different perspectives on this.

Beltrama's characters Fred and Greg can use metalinguistic negotiation on the meaning of *awesome* to try to alter each other's categorization. However, if Sue and Zoe argue about the taste of avocado, they are unlikely to change each other's minds (although, as with the experience of a wine-tasting course, they might). Their negotiation would serve a different purpose, such as in the following argument:

(37) Sue: Avocados are delicious.

Zoe: No, they're disgusting.

Sue: Oh, I didn't know you felt that way. Okay, I was planning to serve Mexican tonight, but I'll do your burrito without guacamole.

The negotiation over the taste of avocado, or another proposition that hinges on a perceptual difference, would fail as negotiation — the participants could not adopt an agreement that would satisfy them both. But it would allow the participants to know that their perceptions differ, and adjust their behavior accordingly. By learning what other people *like*, we can learn what they *are like*.

5.4.8 Counterstance contingency

The essence of the counterstance model (Kennedy & Willer, 2016) is that *consider* and *find* are used in contexts where there are doxastic alternatives in the discourse, which match on

factual information about the world but do not match on decisions about how to resolve indeterminacy of linguistic meaning; those alternatives are called *counterstances*. In a sense, we have the intuition that it is 'up to the speaker' to decide whether a certain word or phrase applies in the given context, and there is an equally felicitous option where it does not.

Kennedy and Willer define the difference between subjective and objective propositions as follows:

(38) Instead, we propose that [the distinction between 'subjective' and 'objective' predicates] corresponds to a pragmatic distinction between those aspects of the common ground that discourse participants take to be grounded in objective facts of the world and those they take to be (to a certain extent) arbitrary matters of linguistic practice. To model this distinction, we embellish our context model with a function k that tracks the contingency of the stipulations involved in achieving an information state. k takes an information carrier s and derives a set κ(s) of s's COUNTERSTANCES: each such counterstance agrees with s on its factual information but disagrees on contextually salient decisions about linguistic practice...Relatedly, we can say that a proposition is subjective just in case we can find some information state s that accepts it — the proposition is true at all possible worlds in $s$but that fails to be accepted by one of s's counterstances. We label this COUNTERSTANCE CONTINGENCY. (Kennedy & Willer (2016) p. 5-6)

But how can we predict whether a certain proposition is part of the facts about the world, or part of the counterstances? This is where the evidential analysis can add to the counterstance model — and find a limitation for it.

Differing categorizations seem to map precisely to what Kennedy and Willer call contextually salient decisions about linguistic practice in their model of counterstance contingency. For an illustrative example, we have the gymnastics judge Chris Grabowecky's recollection of judging Alexei Nemov at the 2004 Olympics that I quoted in Chapter 3: "He [Nemov]," said

Grabowecky, "had errors in his routine that the public doesn't recognize as errors because they're not judges. But he had those beautiful release skills and they go, 'Wow, that guy should win.' But they don't see that he's catching his release close, bending his knees here, taking a step on the landing." (Myers (2016), p. 86)

An expert gymnastics judge would make a contextually salient decision as to whether *excellent gymnastics* applies to the performance he has just seen. There are other doxastic alternatives available — perhaps those where he would have judged the gymnast by the height of his releases rather than by the technique with which he catches them or the straightness of his legs. The judge, after training, experience, and familiarity with the code drawn up by his fellow judges, would choose the counterstance where *excellent gymnastics* applies here. The same would be true of similar evaluative predicates like *beautiful, expert, good music*, etc. These reflect categorizations, which themselves reflect the viewer's intelligence and knowledge, and may be explicitly codified.

Thus, we can map differing categorizations to differing counterstances. The facts about the world, such as what Nemov did, are the same. But how can we incorporate the different perceptions? As we saw in Chapter 3, different perceptions are not subject to persuasion; they are congenitally what they are. They are difficult to explicitly codify because we do not have access to others' brains and interiority.

Consider, according to Kennedy and Willer, presupposes mere counterstance contingency, whereas *find* requires a more stringent condition, which they call *radical counterstance contingency*. With radical counterstance contingency, the prejacent is counterstance contingent relative to every partition introduced by κ_c, i.e. no matter how we resolve uncertainty of meaning based on parameters that support coordination by stipulation, it remains uncertain.

At first glance, this seems to be very similar to differing perceptions. Intuitively, these cannot be coordinated by stipulation – one cannot come to an agreement and stipulate that we now perceive or do not perceive vanillin or the color green. Thus the counterstance model

can work just fine with the differing perceptions and categorizations as an under-layer — if we have differing perceptions of something, we would predict that this something would be radically counterstance contingent. If we have only differing categorizations, while our perceptions seem to be pretty much similar, then we can have coordination by stipulation on the lexical items involved, and thus have ordinary counterstance contingency.

5.5 Conclusion

This dissertation has looked into a problem that the subjectivity literature rarely looks into: what makes subjective predicates differ from objective predicates and how subjectivity varies. It concluded that subjectivity would naturally ensue from people being different; they have both different ways of looking at the world and different ways of classifying it.

In this chapter, I took the conclusions we reached from Chapters 3 and 4, and discussed how they applied to the landscape of subjectivity research. I related the conclusions to the research on evidentiality and assertion conditions. I then went through the same models I had discussed in Chapter 1, and explained how my conclusions help them.

There are still many avenues of further research in the interaction between subjectivity and evidentiality. For example, Krawczyk (2012) and Taranto (2003) argue that *obviously* and *clearly* require a shared evidential context between speaker and interlocutor. This is a proposal that still needs to be tested for subjective predicates, and indeed, the literature on the use of subjective predicates in languages with more compulsory evidentials is quite sparse.

Another section of subjective predicates that I have not yet explored is how they relate to sensory expressions like *This dress looks blue and black* or *It sounds off-key.* The exploration of *find* in Chapter 4 may support further avenues of that research.

A topic that I have currently left for future research is explaining the restriction on the assertion conditions of subjective predicates particularly, which manifests as the Acquaintance Inference. Although my conclusions are a step forward in predicting what predicates

would be subjective, and provide a simple explanation for the use of *find*, they do not explain the Acquaintance Inference not being a presupposition and not being a condition on belief, which is one of the most puzzling aspects of subjectivity.

REFERENCES

Austin, J. L. 1962. *How To Do Things With Words.* William James Lectures 1955. Cambridge: Cambridge University Press.

Ayer, Alfred Jules. 1936. *Language, Truth and Logic.* London: Gollancz.

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy*, **25**, 1–36.

Barker, Chris. 2012. *Negotiating taste.* Manuscript draft.

Barker, Chris. 2013. Negotiating taste. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(2-3).

Beltrama, Andrea. 2016. *Bridging the Gap: Intensifiers Between Semantic and Social Meaning.* Ph.D. thesis, University of Chicago.

Bierwisch, Manfred. 1989. The semantics of gradation. *Pages 71–262 of:* Bierwisch, Manfred, & Lang, Ewald (eds), *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation.* Berlin: Springer-Verlag.

Borkin, Ann. 1973. *To be* and not *to be. Pages 44–56 of: Papers from the Ninth Regional Meeting of the Chicago Linguistic Society.*

Bouchard, David-Etienne. 2012. *Long-Distance Degree Quantification and the Grammar of Subjectivity.* Ph.D. thesis, McGill University.

Bourget, David, & Mendelovici, Angela. 2016. Phonomenal Intentionality. *Pages https://plato.stanford.edu/archives/fall2016/entries/phenomenal–intentionality/ of:* Zalta, Edward N. (ed), *Stanford Encylopedia of Philosophy*, fall 2016 edn. Metaphysics Research Lab, Stanford University.

Bylinina, Lisa. 2013. *The Grammar of Standards: Judge-dependence, Purpose-relativity, and Comparison Classes in Degree Constructions.* Ph.D. thesis, University of Utrecht.

Bylinina, Lisa. 2016. Judge dependence in degree constructions. *Journal of Semantics*, October, 1–41.

Cameron, Leslie. 2014. Pregnancy and olfaction: a review. *Frontiers of Psychology*, **5**(67), Online.

Comesaña, Juan. 2006. A Well-Founded Solution to the Generality Problem. *Philosophical Studies*, **129**(1), 27–47.

Coppock, Elisabeth. 2015 (January 19). *Outlook-based Semantics*. Manuscript draft.

Croy, Ilona, Olgun, Selda, Mueller, Laura, Schmidt, Anna, Muench, Marcus, Hummel, Cornelia, Gisselmann, Guenter, Hatt, Hanns, & Hummel, Thomas. 2015. Research report: Peripheral adaptive filtering in human olfaction? Three studies on prevalence and effects of olfactory training in specific anosmia in more than 1600 participants. *Cortex*, **73**, 180 – 187.

Dechaine, Rose-Marie, Cook, Clare, Meuhlbauer, Jeffrey, & Waldie, Ryan. 2014 (May). *(De)constructing Evidentiality*. Manuscript in submission to special issue of Lingua on evidentiality.

Dowty, David. 1979. *Word Meaning and Montague Grammar*. Netherlands: Springer. Chap. The Semantics of Aspectual Classes of Verbs in English, pages 37–132.

Ducrot, Oswald. 1975. Je trouve que. *Semantikos*, **1**(1), 63–88.

Egan, Andy. 2007. Epistemic modals, relativism and assertion. *Philosophical Studies*, **133**, 1–22.

Egan, Andy, Hawthorne, J., & Weatherson, B. 2005. Epistemic modals in context. *Pages 131–170 of:* Preyer, G., & Peter, G. (eds), *Contextualism in philosophy: Knowledge, meaning and truth*. Oxford University Press.

Faller, Martina. 2002. *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. Ph.D. thesis, Stanford University.

Farkas, Donka, & Bruce, Kim. 2010. On reaction to assertions and polar questions. *Journal of Semantics*, **27**(1), 81–118.

Fleisher, Nicholas. 2013. The dynamics of subjectivity. *Pages 276–294 of: Proceedings of SALT 23*.

Frege, Gottlob. 1956. *Translations from the Philosophical Writings of Gottlob Frege*. Blackwell Publishing. Chap. Negation, pages 117–135.

Geach, Peter T. 1965. Assertion. *Philosophical Review*, **74**(4), 449–465.

Giannakidou, Anastasia, & Mari, Alda. Forthcoming 2016. Mixed (non)veridicality and mood choice in complement clauses: starting with emotive verbs. *Pages – of: Veridicality in Grammar*. University of Chicago Press.

Gibbard, Allan. 1992. *Wise Choices, Apt Feelings*. Harvard University Press.

Gladwell, Malcolm. 2004. The Ketchup Conundrum. *The New Yorker*, September 6, –.

Grinsell, Tim. 2012. Social choice theory and linguistic vaguencess. *In: Proceedings of SALT 22*.

Grinsell, Tim. 2017. *Semantic Indecision*. Ph.D. thesis, University of Chicago.

Gunlogson, Christine, & Carlson, Gregory. 2016. Predicates of experience. *Pages 169–200 of:* Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds), *Subjective Meaning: Alternatives to Relativism*. Linguistische Arbeiten, vol. 559. De Gruyter.

Gutzmann, Daniel. 2016. If expressivism is fun, go for it! *Chap. 2, pages 21–46 of:* Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds), *Subjective Meaning: Alternatives to Relativism*, vol. 559. Berlin: De Gruyter.

Hirvonen, Sanna. 2014. *Predicates of Personal Taste and Perspective Dependence*. Ph.D. thesis, University College London.

Izvorski, Roumyana. 1997. The Present Perfect as an Epistemic Modal. *Pages 222–239 of:* Lawson, Aaron (ed), *Proceedings of SALT 7*.

Kamp, J. A. W. 1975. Two theories about adjectives. *Pages 123–155 of:* Keenan, Edward L. (ed), *Formal Semantics of Natural Language*. Cambridge University Press.

kant1790 Kant, Immanuel. 1892 [1790]. *Critique of Judgement*. http://oll.libertyfund.org/titles/kant-the-critique-of-judgement: Online Library of Liberty.

Kaplan, David. 1989. Demonstratives. *Pages 481–563 of:* Almog, J., Perry, J., & Wettstein, H. (eds), *Themes From Kaplan*. Oxford: Oxford University Press.

Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry*, **4**, 169–193.

Kennedy, Christopher. 2007. Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, **30**, 1–45.

Kennedy, Christopher. 2011. Ambiguity and Vagueness: An Overview. *Chap. 23, pages – of:* Maienborn, C., P. Portner, & von Heusinger, K. (eds), *Semantics: An International Handbook of Natural Language Meaning*. The Hague: Mouton de Gruyter.

Kennedy, Christopher. 2013 (January). *Two Sources of Subjectivity: Qualitative Assessment and Dimensional Uncertainty*. Draft.

Kennedy, Christopher. 2016. Two kinds of subjectivity. *Pages 105–126 of:* Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds), *Subjective Meaning: Alternatives to Relativism*, vol. 559. De Gruyter.

Kennedy, Christopher, & McNally, Louise. 2005. Scale Structure, Degree Modification, and the Semantic Typology of Gradable Predicates. *Language*, **81**, 345–381.

Kennedy, Christopher, & McNally, Louise. 2010. Color, context and compositionality. *Synthese*, **174**(1), 79–98.

Kennedy, Christopher, & Willer, Malte. 2016. Subjective attitudes and counterstance contingency. *Pages – of: Proceedings of SALT 26*, vol. 26.

Khoo, Justin, & Knobe, Joshua. Forthcoming 2016. Moral disagreement and moral semantics. *Nous*, —.

Klecha, Peter. 2012. Diagnosing modality in predictive expressions. *Journal of Semantics*.

Klecha, Peter. 2014. *Bridging the Divide: Scalarity and Modality*. Ph.D. thesis, University of Chicago.

Klein, Ewan. 1980. A Semantics for positive and comparative adjectives. *Linguistics and Philosophy*, **4**, 1–45.

Knaapila, Antti, Hwang, Liang-Dar, & Reed, Danielle R. 2012. Genetic Analysis of Chemosensory Traits in Human Twins. *Chem Sense*, **37**(9), 869–881.

Kölbel, Max. 2002. Faultless Disagreement. *Pages 55–73 of: Proceedings of the Aristotelian Society*, vol. 104.

Korotkova, Natasha. 2016. *Heterogeneity and uniformity in the evidential domain*. Ph.D. thesis, University of California at Los Angeles.

Krawczyk, Elizabeth. 2012. *Inferred Propositions and the Expression of the Evidence Relation in Natural Language Evidentiality in Central Alaskan Yup'ik Eskimo and English*. Ph.D. thesis, Georgetown University.

Lasersohn, Peter. 2005. Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy*, **28**, 643–686.

Lasersohn, Peter. 2009. Relative truth, speaker commitment, and control of implict arguments. *Synthese*, **166**, 359–374.

Lasersohn, Peter. 2017. *Subjectivity and Perspective in Truth-Theoretic Semantics*. Oxford Studies in Semantics and Pragmatics, vol. 8. Oxford University Press.

Lassiter, Daniel. 2011. *Measurement and Modality: The Scalar Basis of Modal Semantics*. Ph.D. thesis, New York University.

Ludlow, Peter. 2008. Cheap Contextualism. *Philosophical Issues*, **18**.

MacFarlane, John. 2007. Relativism and disagreement. *Philosophical Studies*, **132**, 17–31.

MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press.

Majid, Asifaa, & Burenhult, Niclas. 2013. Odors are expressible in language, as long as you speak the right language. *Cognition*, **130**(2014), 266–270.

Malamud, Sophia, & Stephenson, Tamina. 2015. Three Ways To Avoid Commitments: Declarative Force Modifiers in the Conversational Scoreboard. *Journal of Semantics*, **32**, 275–311.

Markt, Sarah C, Nuttall, Elizabeth, Turman, Constance, Sinnott, Jennifer, Rimm, Eric B, Ecsedy, Ethan, Unger, Robert H, Fall, Katja, Finn, Stephen, Jensen, Majken K, Rider, Jennifer R, Kraft, Peter, & Mucci, Lorelei A. 2016. Sniffing out significant "Pee values": genome wide association study of asparagus anosmia. *BMJ*, **355**.

Matthewson, Lisa, Rullmann, Hotze, & Davis, Henry. 2007. Evidentials as epistemic modals: evidence from St'at'imcets. *In: The Linguistic Variation Yearbook*, vol. 7. John Benjamins.

McAuley, Julian, & Leskovec, Jure. 2013. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. *Pages – of: Proceedings of the ACM International Conference on World Wide Web*.

McCready, Eric. 2015. *Reliability in Pragmatics*. Oxford Studies in Semantics and Pragmatics, vol. 4. Oxford University Press.

McCready, Eric, & Ogata, Norry. 2007. Evidentiality, modality and probability. *Linguistics and Philosophy*, **30**, 147–206.

McNally, Louise, & Stojanovic, Isidora. 2015. Aesthetic Adjectives. *In:* Young, James (ed), *(To appear in) The Semantics of Aesthetic Adjectives*. Oxford University Press.

Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds). 2016. *Subjective Meaning: Alternatives to Realism*. Vol. 559. Berlin: De Gruyter.

Menashe, Idan, Abaffy, Tatjana, Hasin, Yehudit, Goshen, Sivan, Yahalom, Vered, Luetje, Charles, & Lancet, Doron. 2007. Genetic Elucidation of Human Hyperosmia to Isovaleric Acid. *PLoS Biol*, **5**(11), e284.

Moltmann, Friederike. 2010. Relative truth and the first person. *Philosophical Studies*, **150**, 187–220.

Murray, Sarah. 2011. *Evidentiality and the Structure of Speech Acts*. Ph.D. thesis, Rutgers University.

Myers, Dvora. 2016. *The End of the Perfect 10*. New York: Touchstone.

Ninan, Dilip. 2014. Taste predicates and the acquaintance inference. *Pages 290–309 of: Proceedings of SALT 24*, vol. 24. SALT.

Nouwen, Rick. 2007. *Predicates of (im)personal taste*. Unpublished manuscript.

Pearson, Hazel. 2012. A Judge-Free Semantics for Predicates of Personal Taste. *Journal of Semantics*, **30**, 1–25.

Pearson, Hazel. 2013. *The Sense of Self: Topics in the Semantics of De Se Expressions*. Ph.D. thesis, Harvard University.

Pedersen, Walter. 2012. *Faultless disagreement and semantic relativity*. Manuscript at http://people.linguistics.mcgill.ca/ walter.pedersen/papers/pedersen-FD.pdf.

Pitt, David. 2004. The Phenomenology of Cognition Or What Is It Like To Think That P? *Philosophy and Phenomenological Research*, **LXIX**(1), –.

Plunkett, David, & Sundell, Tim. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, **13**(23), 1–37.

Potts, Chris. 2002a. The lexical semantics of parenthetical as and appositive which. *Syntax*, **5**, 55–88.

Potts, Chris. 2002b. The Syntax and Semantics of As-Parentheticals. *Natural Language and Linguistic Theory*, **20**, 623–689.

Potts, Christopher. 2007. The Expressive Dimension. *Theoretical Linguistics*, **33**(2), 165–198.

Potts, Christopher. 2014. *Handbook of Contemporary Semantics*. 2 edn. Wiley-Blackwell Publishing. Chap. Presupposition and implicature, pages –.

Richard, Mark. 2004. Contextualism and Relativism. *Philosophical Studies*, **119**(May), 215–242.

Rotstein, C., & Winter, Y. 2004. Total adjectives vs. partial adjectives: Scale structure and higher order modifiers. *Natural Language Semantics*, **12**, 259–288.

Saebø, Kjell Johan. 2009. Judgment Ascriptions. *Linguistics and Philosophy*, **32**, 327–352.

Sassoon, Galit. 2012. A typology of multidimensional adjectives. *Journal of Semantics*, Online.

Schroeder, Mark. 2008a. *Being For: Evaluating the Semantic Problem of Expressivism*. Oxford University Press.

Schroeder, Mark. 2008b. What is the Frege-Geach Problem? *Philosophy Compass*, 703–720.

Searle, John. 1962. Meaning and speech acts. *Philosophical Review*, **71**(4), 423–432.

Silk, Alex. 2016. How To Be An Ethical Expressivist. *Philosophy and Phenomenological Research*, **91**(1), 47–81.

Simons, Mandy, Tonhauser, Judith, Beaver, David, & Roberts, Craige. 2010. What Projects and Why. *In: Proceedings of SALT 20*.

Stalnaker, Robert. 1978. Assertion. *Pages 315–332 of:* Cole, Peter (ed), *Syntax and Semantics, Volume 9: Pragmatics*. Academic Press, NY.

Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy*, **25**(5/6), 701–721.

Stephenson, Tamina. 2007a. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, **30**, 487–525.

Stephenson, Tamina. 2007b. *Towards a Theory of Subjective Meaning*. Ph.D. thesis, MIT.

Stevenson, Charles. 1937. The emotive meaning of ethical terms. *Mind*, **46**(181), 14–31.

Stojanovic, Isidora. 2007. Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy*, **30**, 691–706.

Sundell, Tim. 2010. Disagreements about taste. *Philosophical Studies*, June, Online.

Svenonius, Peter. 1994. *Dependent Nexus: Subordinate Predication Structures in English and the Scandinavian languages*. Ph.D. thesis, University of California at Santa Cruz.

Taranto, Gina. 2003. *Discourse Adjectives*. Ph.D. thesis, University of California at San Diego.

Umbach, Carla. 2016. Evaluative propositions and subjective judgments. *Pages 127–168 of:* Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds), *Subjective Meaning: Alternatives to Relativism*. Linguistische Arbeiten, vol. 559. De Gruyter.

Valentin, Dominique, Chollet, Sylvie, Beal, Severin, & Patris, Bruno. 2007. Expertise and memory for beers and beer olfactory compounds. *Food Quality and Preference*, **18**(5), 776 – 785.

van Rooij, Robert. 2010. Measurement and interadjective comparisons. *Journal of Semantics*, –.

van Wijnbergen-Huitink, Janneke. 2016. Subjective meaning: an introduction. *Chap. 1, pages 1–20 of:* Meier, Cecile, & van Wijnbergen-Huitink, Janneke (eds), *Subjective Meaning: Alternatives to Relativism*, vol. 559. Berlin: De Gruyter.

Vanderbilt, Tom. 2016. *You May Also Like: Taste In an Age of Endless Choice*. Knopf.

Vardomskaya, Tamara. 2014. Subjective Adjectives as a Cooperation Problem. *Pages 191–204 of:* Kohlberger, Martin, Bellamy, Kate, & Dutton, Eleanor (eds), *ConSOLE XXII: Proceedings of the 22nd Conference of the Student Organization of Linguistics in Europe (8-10 January 2014, Lisbon)*. University of Leiden.

von Fintel, Kai, & Gillies, Anthony. 2011. 'Might' made right. *Pages 108–130 of:* Egan, Andy, & Weatherson, Brian (eds), *Epistemic Modality*. Oxford University Press.

von Fintel, Kai, & Gillies, Thony. 2010. Must...stay...strong! *Natural Language Semantics*, **18**, 351–383.

Wolf, Lavi. 2015. *Degrees of Assertion*. Ph.D. thesis, Ben-Gurion University of the Negev.

Wright, Crispin. 1994. *Truth and Objectivity*. Harvard University Press.