

Familiarity Inferences, Subjective Attitudes and Counterstance Contingency*

Towards a Pragmatic Theory of Subjective Meaning

Christopher Kennedy[†] and Malte Willer[‡]

[†]Department of Linguistics, University of Chicago, 1115 E. 58th Street,
Chicago, IL 60637, ck@uchicago.edu

[‡]Department of Philosophy, University of Chicago, 1115 E. 58th Street,
Chicago, IL 60637, willer@uchicago.edu

April 11, 2022

Abstract

Subjective predicates have two interpretive and distributional characteristics that have resisted a comprehensive analysis. First, the use of a subjective predicate to describe an object is in general felicitous only when the speaker has a particular kind of familiarity with relevant features of the object; characterizing an object as *tasty*, for example, implies that the speaker has experience of its taste. Second, subjective predicates differ from objective predicates in their distribution under certain types of propositional attitude verbs. The goal of this paper is to argue that these features can be explained in a uniform way and within a broadly truth-conditional approach to semantic content, given a view of subjective language as an essentially pragmatic, context-sensitive phenomenon. Specifically, we propose that what renders an issue subjective in discourse is speakers' awareness of *counterstance contingency*: contingency relative to information states that represent alternative pragmatic stances.

1 The Plot

Predicates of personal taste such as *tasty* or *fun* come with an inference of direct experience: when a speaker uses one of these predicates to describe an object or event, she

*For valuable discussion and feedback on earlier versions of this work, we are grateful to Chris Barker, Peter Laserson, John MacFarlane, Patrick Muñoz, Rachel Rudolph, Kjell Johan Sæbø, Tamara Vardomskaya, Anubav Vasudevan, and audiences at SALT 26, PhLiP 4, the 2017 Annual Meeting of the Linguistic Society of Japan, the Alternative Logics and Natural Language Workshop at the Hebrew University of Jerusalem Logic, Language and Cognition Center, the 2018 World Congress of Philosophy, Northwestern University, New York University, Yale University, the University of Chicago, the University of Illinois, Urbana-Champaign, the University of Maryland, Nagoya University, and Nanjing University. Special thanks to three anonymous L&P reviewers and to the handling editor Patrick Grosz for extensive helpful comments. We also gratefully acknowledge the Neubauer Collegium for Culture and Society for their generous support of this research.

typically presents herself as having first-hand knowledge of that object or event, in the sense that she has direct experience of it in the relevant way (MacFarlane 2014; Ninan 2014; Pearson 2013; Franzén 2018; Anand and Korotkova 2018; Muñoz 2019; Ninan 2020). Following an utterance of (1) or (2) with (1a) or (2a), for instance, would likely cause raised eyebrows; if the speaker lacks the relevant experience of the taste of sea urchin or the activity of downhill skiing, it would be much more appropriate for her to choose hedged variants such as (1b) or (2b).¹

- (1) I have never tried sea urchin.
 - a. # It's (not) tasty.
 - b. I hear it's (not) tasty.
- (2) I have never gone downhill skiing.
 - a. # It's (not) fun.
 - b. It must (can't) be fun.

This is noteworthy since many other predicates lack a similar direct experience component: a straight assertion that sea urchin is orange does not at all require that the speaker has seen sea urchin; it may express a judgment formed on the basis of testimony. A central goal of this paper is to make progress toward understanding why this is so.

Part of the explanandum here is that direct experience inferences show a complex pattern of projection. Negation does not cancel the implication that the speaker has first-hand experience with the item under consideration, as shown by (1a) and (2a), but hedging does, as shown by (1b) and (2b). Similarly, to state that Kim was recently in Tokyo and enjoyed a tasty dish at her favorite izakaya is not for the *speaker* to suggest that she actually tasted Kim's meal. Such "exocentric" uses of taste predicates are tied to tastes and sensibilities other than the speaker's and thus differ from "autocentric" uses in which the item under consideration is evaluated based on the speaker's own tastes and sensibilities (Lasersohn 2005). Here we are especially interested in the fact that embedding a taste ascription under a SUBJECTIVE ATTITUDE VERB such as English *find* comes with a direct experience inference, but embedding under *believe* does not:

- (3) I have never tried sea urchin, but ...
 - a. ... I believe it to be tasty.
 - b. # ... I find it tasty.
- (4) Kim has never tried sea urchin, but ...
 - a. ... she believes it to be tasty.
 - b. # ... she finds it tasty.

There is nothing strange about believing something to be tasty in the absence of first-hand knowledge of its taste — observations of others' reactions or reports from reliable sources could justify such a judgment. But one cannot find something tasty without actually having tasted it.

¹A note on notation: we use # throughout to indicate that an utterance of the relevant sentence is unacceptable (either in general or in the specified context, as appropriate), and indicate our assessment of the source of the unacceptability in the descriptive and analytical portions of the paper.

Existing approaches in the literature all leave something to be desired, and to get the problem into proper view it is helpful to see why. [Ninan \(2014\)](#) considers an EPISTEMIC account of direct experience. Start with the knowledge norm of assertion defended at length by, for instance, [Williamson \(1996, 2000\)](#): one must assert a sentence ϕ in some context c only if one knows that ϕ is true as used in c . And add to this the following acquaintance principle: whenever a taste predicate is used autocentrically, knowing that x is tasty (or that it is not tasty) requires first-hand knowledge of x 's taste. It then follows immediately that an utterance of (1a) is problematic: to be in a position to assert that sea urchin is tasty, one must know that sea urchin is tasty, which in turn requires having sampled it — exactly what the first sentence denies. Since the acquaintance principle does not impose any constraints on hedged autocentric uses of predicates of personal taste — or of their epistemically modalized uses, for that matter — it is also not surprising that utterances of (1b) or of (2b) get a pass.

One important problem with this story is that the acquaintance principle is not obviously correct, since knowledge claims about (for instance) taste that are based on indirect evidence are in general felicitous ([Muñoz 2019](#)).

- (5) I know that the licorice is tasty ...
 - a. ... because Alfonse made it.
 - b. ... because it's Finnish.

(5a) and (5b) easily roll of the tongue, and this is unexpected if the acquaintance principle were in fact true.²

Second, even if correct, the acquaintance principle does not immediately generalize to cases where direct experience inferences arise in embedded contexts, including under *find*. It states that what it takes to know that x is tasty in autocentric contexts, but does not say what it takes to know that Kim (or oneself) finds x tasty. It is, of course, not unreasonable to claim that one cannot know that S finds x tasty without knowing that S has actually sampled x , and then to add that knowing that S believes x to be tasty comes with no such requirement. But this fact calls as much for an explanation as the one articulated by the original acquaintance principle.

Treating direct experience inferences as a PRESUPPOSITIONAL affair is the major alternative to an epistemic account in the existing literature ([Pearson 2013](#); [Anand and Korotkova 2018](#); [Ninan 2020](#); [Willer and Kennedy 2020](#)). Here the main challenge is to explain why these inferences behave in some but not all respects like presuppositions: why, for instance, they project from under the scope of negation (as presuppositions do) but disappear under the scope of epistemic modals (unlike ordinary presuppositions). While progress has been made on this front, the projection facts surrounding English *find* and *believe* still remain a bit of a puzzler.

- (6) Mary is an only child, but ...
 - a. # ... I find her brother attractive.
 - b. # ... I believe her brother to be attractive.
- (7) I've never seen Mary's siblings, but...

²See [Muñoz \(2019, Section 4.1.1\)](#) for additional critical discussion of the acquaintance principle. [Ninan \(2014\)](#), we should add, acknowledges that his explanatory strategy leaves some critical questions unresolved, though his concerns are different from the ones voiced here.

- a. # ... I find her brother attractive.
- b. ... I believe her brother to be attractive.

In (6), the presupposition that Mary has siblings (triggered by the reference to her brother) has a tendency to project from under *believe* as well as from under *find*, but the inference of direct experience projects only from the latter, as shown by the contrast in (7). No account of the direct experience requirement — presuppositional or otherwise — can be complete without a discussion of what makes subjective attitude verbs such as English *find* special.³

Our goal in this paper is to argue that direct experience inferences in both matrix assertions and under subjective attitude verbs are a special case of what we will call FAMILIARITY INFERENCES — familiarity with a set of facts that are relevant for determining whether a predicate applies — that underwrite taste judgments specifically and “subjective” judgments more generally. Our general strategy here will be to let an explanation of why subjective attitude attributions impose distinct familiarity constraints on the attributee inform our explanation of why subjective assertions impose those very same constraints on the speaker. In brief, attitude ascriptions involving *find* (or its interesting cousin *consider*) require their complement to be subjective in a particular way — a notion we propose to analyze as a sensitivity to distinct discourse alternatives that arise pragmatically from language users’ sophisticated awareness that (what they take to be) matters of fact underdetermine how people speak and what people think. We will show how the semantics of subjective attitude attributions thus understood triggers familiarity inferences, and then move on to show how the same explanation can be extended to account for the specific inferences in assertions involving taste and other subjective predicates.

The paper is structured as follows. Section 2 presents the key data involving subjective attitude verbs, and shows why existing accounts are inadequate. In Section 3 we first explain and motivate our own analysis of the selectional properties of subjective attitude verbs in informal terms and then provide a formal analysis. Section 4 shows how the framework provides the basis for an explanation of familiarity inferences for both subjective attitude verbs and for unembedded uses of subjective predicates. Section 5 picks up some loose ends and Section 6 concludes.

2 Subjective attitude ascriptions

We begin by explaining what an adequate semantics for subjective attitude verbs should look like, focusing on English *find* and *consider*. Section 2.1 describes the unique selectional properties of these verbs, discusses two previous attempts to explain these patterns, and argues that they are insufficiently general. Section 2.2 substantiates our claim above that direct experience inferences are a special case of a more general category of inferences requiring familiarity with a particular set of facts and shows how these inferences track the meaning of the predicate in the complement of a subjective attitude verb.

³Ninan (2020) proposes that intentional operators cancel the direct experience inference and then suggests that *believe* but not *find* belong to this category, thus effectively adopting a proposal from Sæbø (2009). Part of our goal here is to show that Sæbø’s proposal is neither necessary nor sufficient to account for the data surrounding subjective attitude verbs.

2.1 Complement selection

What makes subjective attitude verbs notable, and distinct from doxastic attitude verbs like *think* and *believe*, is that they imply that it is somehow a “matter of discretion” whether the predicate in their complement truthfully applies to its argument, and they correspondingly reject complements in which this relation can only be construed as a matter of contingent fact (see, e.g., Bouchard 2012; Bylinina 2017, Fleisher 2013; Kennedy 2013; Hirvonen 2014; Reis 2013; Sæbø 2009; Stephenson 2007; Umbach 2016; Vardomsakaya 2018; and Muñoz 2019). Consider, for example, a context in which Kim has been handed a glass of some unknown beverage to sample. She can report her judgment about the quality of the taste of the drink using either *find* or *consider* as in (8a–b), because there is an intuitive sense in which it is a matter of her discretion whether the predicate *delicious* truthfully applies; the doxastic attitude verb *think* is also acceptable here.

- (8) a. I find this drink delicious.
b. I consider this drink delicious.
c. I think this drink is delicious.

But she cannot felicitously use *find* or *consider* to report her judgment that the drink is fermented — though she can use *think* — because, intuitively, whether it is fermented is not a matter of discretion, but rather a matter of fact:

- (9) a. # I find this drink fermented.
b. # I consider this drink fermented.
c. I think this drink is fermented.

Find and *consider* are not fully interchangeable, however: not all predicates that have a subjective flavor are acceptable under the former. For example, there is a certain amount of discretion in the use of the predicate *empty*: does it hold only of containers that are completely devoid of contents, or can it be used when only a small amount of stuff remains inside? And so if Kim picks up the bottle from which her drink was poured and sees only a few drops of liquid left inside, she could describe the situation using *consider* (or *think*), but not using *find*:

- (10) a. # I find this bottle empty.
b. I consider this bottle empty.
c. I think this bottle is empty.

Furthermore, even when the same predicate can embed under both *find* and *consider*, as with *delicious* in (8a–b), there is a difference in meaning between the variants (cf. McNally and Stojanovic 2017). The *find*-sentence expresses a judgment that sounds, in some sense, more “sensual” or qualitative, while the *consider*-sentence expresses a judgment that sounds more “intellectual” or quantitative. The pairs in (11) bring out this difference quite clearly. The variants in (11a) have a more intimate feel than those (11b), and so would sound more natural in a conversation between lovers, while the sentences in (11b) would be more appropriate than those (11a) in context in which the speaker is trying to maintain professional distance, such as a conversation between an acting coach and an actor.

- (11) a. I find you fascinating/irresistible/sexy.
 b. I consider you fascinating/irresistible/sexy.

One of our goals in this paper is to develop a semantics for subjective attitude verbs that captures the intuition that they differ from plain doxastic attitude verbs in requiring it to be in some way a matter of discretion and not a matter of fact whether the preja-cent is true, that uses this to derive their selectional properties, and that, at the same time, allows us to encode the finer-grained differences between *find* and *consider* which explain contrasts like those in (10) and (11). We will present our account momentarily; in the remainder of this section, we review two alternative approaches to the selectional properties of subjective attitude verbs and argue that neither checks all of the analytical and empirical boxes.

The first approach, articulated in Sæbø 2009, provides a type-theoretic account of the contrast between (8a) and (9a)/(10a). In this analysis, predicates like *delicious* are distinguished from predicates like *fermented* and *empty* in having an implicit judge argument:

- (12) a. $\llbracket \textit{delicious} \rrbracket^{c,w} = \lambda x \lambda y. x$ tastes very good to y at w
 b. $\llbracket \textit{fermented} \rrbracket^{c,w} = \lambda x. x$ has undergone fermentation at w
 c. $\llbracket \textit{empty} \rrbracket^{c,w} = \lambda x. x$ is devoid of contents at w

And the function of a *find*-type subjective attitude verb is to saturate the judge argument of its complement with the denotation of its subject; *find* does not introduce any content of its own:

- (13) $\llbracket \textit{find } \phi \rrbracket^{c,w} = \lambda x. \llbracket \phi \rrbracket^{c,w}(x)$

The problem with (9a) and (10a), then, is that the embedded small clauses lack unsaturated judge arguments, and so fail to serve as legitimate arguments for *find*. Although Sæbø does not address familiarity inferences, the meaning that this analysis derives for (8a) is equivalent to the meaning of an autocentric interpretation of the bare form as uttered by Kim, which as we have seen, has a direct experience inference of its own. This opens the door to an account of the former in terms of the latter.

However, a drawback of this kind of analysis is that it does not easily generalize to address the distinctions between *find* and *consider* that we saw above. As a further illustration of this challenge, consider the use of the adjective *long* to describe temporal duration, as in (14).

- (14) The 8-hour flight from Chicago to Frankfurt is longer than the 13-hour flight from Chicago to Tokyo.

(14) can be heard in two ways: as a quantitative statement about temporal duration of the flights, or as a qualitative statement about the experience of the temporal duration of the flights (Kennedy 2013). Assuming the temporal modifiers are accurate, (14) will normally be heard as false on the quantitative reading, but it could very naturally be heard as true on the qualitative reading, if, say, the person making the claim typically flies economy to Frankfurt and first class to Tokyo. And in line with what we saw above, embedding (14) under *find* vs. *consider* disambiguates it: (15a) has only a qualitative reading, while (15b) has only a quantitative reading, and so suggests that Kim is taking something else into account in making her judgment, other than flight time alone, such as the time it takes to get to the airport, or the time spent in immigration lines.

- (15) a. Kim finds the 8-hour flight from Chicago to Frankfurt longer than the 13-hour flight from Chicago to Tokyo.
 b. Kim considers the 8-hour flight from Chicago to Frankfurt longer than the 13-hour flight from Chicago to Tokyo.

At the same time, when only a quantitative interpretation is plausible, embedding under *find* is bad. This is the case in (16a), since Kim cannot have a qualitative experience of the passage of prehistoric time. (16b) is fine, however, as long as it is taken to be a matter of some discretion as to where the boundaries between geologic eras are drawn.

- (16) a. # Kim finds the Triassic Period longer than the Permian period.
 b. Kim considers the Triassic Period longer than the Permian period.

While there is clearly some kind of meaning distinction between the quantitative and qualitative senses of an adjective like *long*, there is no obvious type-theoretic reflection of this difference: both senses are gradable, for example, and both have the same basic syntactic distribution, with the one exception being embeddability under *find*. So while *find* and *consider* are evidently sensitive to different ways that a predicate can be “subjective,” there is no independent evidence that this difference corresponds to a difference in semantic type.

A second type of approach effectively takes direct experience inferences as the analytical starting point, and accounts for the selectional properties of *find* in terms of lexically encoded evidential constraints. Stephenson (2007) proposes that *find* is a doxastic attitude verb that differs from *believe* and *think* (and *consider*) in imposing an additional requirement that the prejacent denote a proposition that the attitude holder can have direct experience of. Responding to shortcomings of this account noted by Sæbø (2009), Muñoz (2019) adds that this direct evidence is impossible for any agent other than the subject of the attitude verb to have. So roughly, the evidence that some drink tastes very good, given some individual anchor x , is that it tastes very good *to* x ; only x can tell directly whether this is so, and hence *delicious* embeds under *find*. In contrast, the evidence that some drink has undergone fermentation, or that some container is devoid of contents, given some individual anchor x , is evidence of the sort that anyone in principle can have directly, so neither *fermented* nor *empty* embed under *find*. The upshot of this kind of approach is that “the infelicity of *find*-reports [...] tends to hold in virtue of the lexical semantics of predicates, tracking whether they contain some component specially sensitive to direct evidence, in such a way that there can in principle be direct evidence for the relevant hyperintension that only one agent can have” (Muñoz 2019, p. 274).

The problem with linking embeddability under *find* to some kind of direct evidence constraint is that it is too restrictive. It is, for sure, not unreasonable to think that experiential predicates such as *delicious*, *tasty*, and *fun* encode a distinct sensitivity to evidence that is only directly accessible to one individual but not another, but not all predicates that embed under *find* are of this particular kind. The following naturally occurring examples demonstrate that, in addition to experiential predicates, *find* can embed: character trait predicates (*brave*, *irresponsible*, *naive*, *evil*, *heroic*, *stupid*, *arrogant*, *petty*, *mean-spirited*); aesthetic predicates (*kitsch*, *gaudy*, *over the top*, *dynamic*, *profound*, *flexible*, *elegant*); and moral predicates (*wrong*, *right*, *unacceptable*, *permissible*, *impermissible*) (cf. Vardomsakaya 2018).

- (17) a. How do you feel about Timothy Treadwell? Do you find him brave and interesting or irresponsible and naive?
- b. I don't find him evil à la Moriarty but I don't find him heroic either.
- c. Men find him cowardly and women find him disturbing
- d. I find him stupid, arrogant, petty, and decidedly mean-spirited.
- (18) a. Some may find [the temple] kitsch, some may find it gaudy, some may find it over the top, but the level of devotion and respect shown by the multitude of Buddhist pilgrims on the day we visited suggests that they view it as religiously significant.
- b. I always look at the drawing ... if I find it dynamic, profound and flexible, it touches a soft spot of mine.
- c. This is also a valid solution, and a practical one in some languages, but few people will find it elegant.
- (19) a. I find cheating wrong, mostly because it says you can't be trusted.
- b. I don't find it right that people who make little pay more taxes towards the community than big business.
- c. Like many in the game I find it unacceptable that so little has been done to reform Fifa.
- d. In this case, Unger thinks that we will find it permissible to push the one person on roller skates in front of the trolley, even though if it were the only alternative to letting the trolley kill the six, we would find it impermissible. (Kamm 2007, p. 197)

It is unclear what aspect of the lexical semantics of *brave*, *arrogant*, *over the top*, *dynamic*, *wrong*, *impermissible* and so forth could be sensitive to direct evidence in some distinguished form. There is, to be sure, an evaluative component to the meanings of these predicates, and it makes sense to say that speakers assign extensions to them in ways that vary according to their own affective state. But it does not obviously follow that direct evidence for such evaluations can only be had by the evaluator, any more than direct evidence for an assessment of what counts as empty can only be had by the assessor. We take this as reason to reject an analysis of the selectional properties of *find* in terms of a requirement of direct experience specifically.

2.2 Familiarity Inferences

We saw in Section 1 that *find*-reports involving predicates of personal taste and other experiential predicates trigger the inference that the attitude holder has direct experience relevant for assessing the truth of the prejacent. This is not so for *find*-attributions with non-experiential predicates — as we just noted — but these do give rise to something similar, as shown by the following contrasts between *find* and *believe*:

- (20) The only thing Kim knows about Timothy Treadwell is that Werner Herzog made a documentary about him. And because of that ...
 - a. # ... she finds him brave.
 - b. ... she believes that he's brave.

- (21) Lee doesn't know anything about Smith's solution, but he does know Smith's style of analysis, so ...
- a. # ... he finds the solution elegant.
 - b. ... he believes that the solution is elegant.
- (22) They don't know what Jones said to Smith, but Jones has a history of making insensitive remarks, so ...
- a. # ... they find the remark unacceptable.
 - b. ... they believe that the remark was unacceptable.

One can believe that someone or something is brave, elegant or unacceptable based on rather sparse information, but to *find* someone or something brave, elegant or unacceptable, one must evidently know — or at least believe that one knows — a particular set of facts: how a person responds to danger, how many stipulations a solution must make in order to work, the content of, or intentions behind, a specific utterance or action, and so forth. Sometimes this kind of connection to the relevant set of facts entails acquaintance in the strict sense of direct perceptual experience, e.g. for aesthetic predicates that characterize visual properties of an object, and, of course, for predicates with experiential semantics:

- (23) The only thing Kim knows about sea urchin is that many people enjoy eating it. And because of that ...
- a. # ... she finds it tasty.
 - b. ... she believes that it's tasty.

But the full pattern suggests that direct experience is just a special case of a more general FAMILIARITY requirement: for a *find*-attribution to be felicitous, the attitude at play must constitute a genuine assessment of, and so require familiarity with, the facts that matter for whether or not the attitude is true. And what those facts are is a function of the lexical semantics of the embedded predicate: when the predicate has an experiential semantics, familiarity requires experiencing things a certain way, and manifests itself as a distinct direct experience requirement, à la Stephenson (2007) and Muñoz (2019). But when the predicate describes a character trait, or a moral property, familiarity manifests itself in other ways.

Familiarity inferences are, moreover, a feature of subjective attitude verbs generally, and not just a feature of *find* specifically. Replacing *find* with *consider* in (20a)–(23a) does not improve acceptability, and we see a similar contrast between *consider* and *believe* if we look at examples involving predicates that embed under *consider* but not *find*, such as *heavy* on its quantitative use:

- (24) The only thing Kim knows about this suitcase is that it's big, so ...
- a. # ... she considers it heavy.
 - b. ... she believes that it's heavy.

And an example like (25) demonstrates quite clearly that what is required is familiarity with the *relevant* facts, in this case facts about ingredients and preparation. When it is made explicit that the attitude holder lacks such information, a *consider*-report is unacceptable, even in the presence of direct experience:

- (25) Kim doesn't know the ingredients that went into this cake, but based on its taste...
- a. # ... she doesn't consider it gluten-free.
 - b. ... she doesn't believe that it's gluten-free.

At the same time, a *consider*-report is also fine in the absence of direct experience, as long as the attitude holder is familiar with the relevant facts:

- (26) Kim hasn't tried the cake, but she knows that it was made using rice flour from a mill that also produces wheat flour, so ...
- a. ... she doesn't consider it gluten-free.
 - b. ... she doesn't believe that it's gluten-free.

Finally, for both *find* and *consider*, familiarity inferences show the projection pattern of presuppositions. They project out of negation:

- (27) a. # Kim has never tried sea urchin because she doesn't find it tasty.
 b. # The only thing Kim knows about this suitcase is that it's small, so she doesn't consider it heavy.

And out of questions: as shown by (28)–(29), it is infelicitous to ask whether a subjective attitude holds when it is manifest in the context that the relevant familiarity conditions are not satisfied.

- (28) A: Kim has never eaten sea urchin, but wants to try it.
 B: Why? Does she #find/ think it tasty?
- (29) A: Kim has no information about the weight of this suitcase, but she doesn't want to carry it.
 B: Why? Does she #consider/ think it heavy?

At first glance, it looks like familiarity inferences need not project out of modals, much as we saw with unembedded uses of taste predicates:

- (30) a. Kim has never eaten sea urchin, but she might find it tasty.
 b. Kim has no information about the weight of this suitcase, but she might consider it heavy.

However, the second clauses in examples like (30a–b) are most naturally understood as implicit conditionals, whose antecedents introduce content that triggers assessment of the modal claim at an index where familiarity is satisfied (*...but if she eats it/acquires the information...*). If we construct examples in which the modal claim is unambiguously assessed at an index where familiarity does not hold, the sentences are infelicitous:

- (31) a. # Kim has never eaten sea urchin, but it's possible that she finds it tasty.
 b. # Kim has no information about the weight of this suitcase, but it's possible that she considers it heavy.

Subjective attitude verbs, then, presuppose that the attitude holder is familiar with a set of facts relevant for assessing the truth of the prejacent, and what these facts are depends on the meaning of the embedded predicate. As a consequence, it will not do to tie the selectional properties of *find* to direct evidence. First, this would fail to account for the more generalized familiarity inferences that we see when *find* embeds non-experiential predicates, and second, it would tell us nothing about why *consider* also comes with distinct familiarity inferences. Instead, we propose to analyze the selectional criteria of both types of subjective attitude verbs in terms of a feature that all of the predicates they embed share: subjectivity, of a sort that we will make precise in the next section. We will then show that the familiarity implications of subjective attitudes flow naturally from our analysis of subjectivity, and we will conclude by showing how the account can be generalized to capture acquaintance inferences of unembedded uses of subjective predicates.

3 Analysis

3.1 Counterstance contingency

Our analysis of subjective attitude verbs, first articulated in Kennedy and Willer (2016) and further refined here, is closest in spirit to proposals developed by Coppock (2018) and Bouchard (2012). Coppock replaces possible worlds with “outlooks,” which are refinements of worlds that settle not only matters of fact but also matters of opinion, and then lets all predicates — including predicates of personal taste — have ordinary extensions relative to these refined points of evaluation. Since opinions differ, a world will allow for different refinements and thus correspond to multiple outlooks, and it makes good sense to call a predicate “discretionary” if its extension varies across the outlooks corresponding to a single world. (Objective predicates, in contrast, will at most vary in their extensions across worlds.) A proposition is discretionary just in case its truth-value varies across the outlooks corresponding to a single world. Coppock’s proposal for Swedish *tycka* — which patterns with English *find* in many ways — is that it presupposes that its complement is discretionary. A similar proposal for *find* is articulated by Bouchard (2012) when he suggests that it carries a “subjective contingency presupposition:” keeping all the non-subjective facts constant, it must be possible to judge the complement clause true, and it must be possible to judge it false.

Like Coppock and Bouchard, we suggest that what makes subjective attitudes special is that they carry a contingency presupposition. However, our proposal is specifically designed to account for the fine-grained differences between *find*- and *consider*-type subjective attitude verbs. There is, for instance, a distinct sense in which *empty* is a discretionary predicate — indeed, it lives happily under the scope of *consider* — but as we have seen, it is not discretionary “enough” to embed felicitously under *find*. We will develop an account of the notion of contingency at play here that makes sense of exactly this observation.

Another distinguishing feature of our analysis is that it does not try to locate the distinction between subjective and objective predicates exclusively in the lexicon. Take, for instance, the difference between (32a) and (32b):

- (32) a. Kim considers Crimea part of Russia.

- b. # Kim considers Siberia part of Russia.

The intuitive explanation of the contrast is that the sovereignty over Crimea is (at the time of writing this paper) disputed, hence the use of *consider* in (32a) seems appropriate, while Siberia being part of Russia would count as an objective fact, hence the use of *consider* in (32b) is odd. But this is not a matter of semantics; it is simply a matter of what background information the discourse context provides or can be accommodated. We thus propose that the kind of contingency that licenses embeddability under subjective attitude verbs, and gives rise to “subjective meaning” more generally, is fundamentally pragmatic.

Specifically, our key proposal is that the contingency that licenses embeddability under subjective attitude verbs arises from language users’ sophisticated awareness that (what they take to be) matters of fact only partly determine what we say and think. To get the guiding intuition into view, observe that in the following two examples, replacing *believe* with *consider* signals that the formation of the attitude under consideration must have involved a “leap from the facts:”

- (33) a. Kim believes the gas tank to be full.
b. Kim considers the gas tank full.
- (34) a. Kim believes herself to be a Chicagoan.
b. Kim considers herself a Chicagoan.

For example, (33b) signals that Kim’s commitment to the gas tank being full is based on something more than just knowledge of how much fuel it contains, such as how much she is willing to spend for gas. And (34b) would be appropriate in a context in which Kim in fact lives in one of the Chicago suburbs, but has formed a self-identity based around her love of deep-dish pizza and the Chicago Cubs. Plain belief attributions, to be clear, do not exclude that adopting the commitment involves a leap from the facts; but the use of *consider* explicitly signals the attitude to be subjective in this specific way.

The more concrete proposal is that a subjective attitude ascription asserts belief in the proposition expressed by the complement clause, and presupposes the contingency of this belief across a set of contextually provided alternatives to the attitude holder’s doxastic state, all of which agree on the salient facts of the matter but disagree on judgments about those facts. We label these alternatives COUNTERSTANCES and the contingency across them COUNTERSTANCE CONTINGENCY. Each of these contextually generated alternatives constitutes a distinct “pragmatic stance” in the sense that the choice of one rather than another is a practical affair, reflecting (perhaps unarticulated) practical decisions, intentions, and plans about, centrally but not exclusively, language use. We will comment on the full range of practices that are of relevance here in Section 5.4, but the paradigmatic cases are those in which aspects of underdetermined meaning are resolved in one way or another: those in which an agent’s doxastic state constitutes a pragmatic stance on language use by categorizing, say, race horses as athletes or 1.75 meter gymnasts as tall.⁴ The resulting beliefs are counterstance contingent in the sense that it is possible to agree on the facts (the strength, speed and stamina of racehorses, the heights of gymnasts, and so on) and still disagree on the issue, simply in virtue of adopting

⁴See MacFarlane 2016, 2020 for a thoroughly pragmatic spin on the familiar labeling of vagueness as indecision that is congenial to the general perspective proposed here.

different stances on language use. It is in this sense that subjective attitude ascriptions imply that the truth of the prejacent is a “matter of discretion,” and it is for this reason that statements like (33b) and (34b) can imply that the attitude holder’s commitment involves a “leap from the facts” — though not necessarily that the attitude holder herself sees things this way. As far as she is concerned, her position may be one that is entirely dependent on factual considerations. We will say more about this in Section 3.4.

Counterstances are thus pragmatic enrichments of a set of shared background assumptions that are treated as matter of fact in context. It should be uncontroversial that the attitudes ascribed in (33) and (34) are counterstance contingent in context: one may just as well think that one’s fuel budget does not matter for whether some gas tank is full, or that one does not qualify as Chicagoan unless one actually lives (or at least has lived) in the city of Chicago. It likewise does not take much imagination to see how the belief that Crimea is part of Russia (32a) could be counterstance contingent in 2022 when this manuscript went to press, eight years after the annexation of the former by the latter and during the Russian invasion of Ukraine. Is a referendum sufficient for some region to become part of another country, or does the referendum need to be recognized by other states? If so, how many, and does the referendum also have to be in accordance with the local constitution? Nothing about the meaning of *part of Russia* settles these questions, but how one answers them obviously decides whether Crimea counts as a part of Russia or not. In contrast, it is much less easy to see how one’s stance on these or other controversial issues about statehood could matter for the political status one is inclined to assign to Siberia, at least not without additional stage setting or accommodation.⁵ In the absence of such information, the proposition that Siberia is part of Russia fails to be counterstance-contingent, and fails to embed under *consider*.

⁵Accommodation can help us understand why the negation of (32b) is fine:

- (i) Kim doesn’t consider Siberia part of Russia.

First, note that negations of both *consider* and *find* are most naturally interpreted with “Neg-raising,” as though the negation is in the embedded clause. (This is a feature they share with plain doxastic attitude verbs like *think* and *believe*.) The proposition that Siberia is not part of Russia is of course no more counterstance contingent than the proposition that it is, in typical contexts, but characterizing Kim as committed to a proposition that is obviously false invites us to accommodate a basis for making sense of her judgment, and one way to do this is to suppose that there is some legitimate, alternative stance which would render it true. There is, in contrast, no corresponding pressure to find a way to make sense of a judgment that is obviously true, as in (32b).

Here the difference in meaning between *consider* and *believe* is informative: intuitively, (ii) characterizes Kim as “clueless” in a way that (i) does not.

- (ii) Kim doesn’t believe that Siberia is part of Russia.

This is because to say (ii), in a context in which it is otherwise taken for granted that Siberia is part of Russia, is just to say that Kim is mistaken about the facts. To say (i), on the other hand, invites accommodation of some issue that render it unsettled whether Siberia counts as part of Russia — say, a question about the political legitimacy of boundaries drawn as a result of Russian conquest of the Khanate of Sibir — and communicates that Kim’s “error” involves a disagreement on how to settle this question, in light of this issue. To the extent that we are willing to accommodate such issues, cases like (i) will make sense; when this kind of accommodation is implausible, negation does not help:

- (iii) # Kim considers/doesn’t consider the sum of two and two equal to four.

3.2 Radical counterstance contingency

Our basic proposal for explaining the contrast between subjective attitude verbs and plain doxastic attitude verbs, then, is that the former, but not the latter, presuppose the counterstance contingency of the opinion at play, and that speakers are inclined to treat an opinion as counterstance contingent if (but, as we will discuss in Section 5.4, not only if) they can see it as sensitive to the contingencies of linguistic practice. It remains to explain the more fine-grained differences between *consider*-type and *find*-type subjective attitude verbs. Our key proposal is that the latter presuppose a distinguished kind of subjectivity that we label RADICAL COUNTERSTANCE CONTINGENCY, which flows from a distinguished kind of pragmatic underdetermination that we now attempt to explain.

The guiding idea here is that not all counterstance contingency is created equal. Sometimes it makes sense for speakers to propose to coordinate a stance by stipulation. This is what we see, for example, in (35a–c), where “*for present purposes*” should be heard as referring to some salient task, action or goal whose execution somehow requires categorization of objects according to whether they satisfy the predicate. For example, in (35a–c), this might be: choosing which actors will play Elves and which will play Dwarves in a stage production of *The Lord of the Rings*; deciding what kinds of meals to serve the guests at a party; formulating tax policy; and deciding which bottles to include in a review of “inexpensive wines.”

- (35) For present purposes,
- a. let’s count any actor over 2 meters as tall and any actor under 1.75 meters as short.
 - b. let’s count anyone who eats shellfish but no other animals as vegetarian.
 - c. let’s count any family with annual income greater than \$250K/year as rich.
 - d. let’s count any bottle of wine that costs under \$20 as inexpensive.

Of course, an interlocutor is free to reject any of these stipulations; the point is that it can be a natural discourse move to propose them, in the right context, in order to explicitly fix a basis for linguistic categorization in the service of some practical purpose (cf. MacFarlane 2020).

This kind of move is not felicitous with all predicates, however. (36a–d), for example, sound decidedly odd.

- (36) For present purposes,
- a. # let’s count any linguistic puzzle that resists explanation as fascinating.
 - b. # let’s count anyone willing to camp in grizzly bear territory as brave.
 - c. # let’s count any painting that uses more than five colors as dynamic.
 - d. # let’s count any action that harms no one as permissible.

Note that it is not the case that we simply cannot talk about what counts as fascinating, brave, dynamic or permissible: we may suppose that it is part of the sense of *brave* that an act only counts as such if it is right, and we certainly may ask which acts count as permissible and which do not. What we cannot do, in a natural way, is stipulate a specific criterion as the basis for categorizing objects according to these predicates. We may, of course, always try to make others see things just the way we do, but unlike what we saw

with the predicates in (35), it does not make sense for the predicates in (36) to simply stipulate some arbitrary criterion as the basis for establishing a conversational convention on how to use them.⁶

Given some set of counterstances, then, we may ask which ones agree on just those factors that allow for COORDINATION BY STIPULATION of the sort that we see in (35). Partitioning that set along these lines delivers a set of equivalence classes of counterstances, such that each cell agrees on decisions that support coordination by stipulation, but not on those that do not. We say that an opinion is RADICALLY COUNTERSTANCE CONTINGENT just in case it is contingent with respect to each of these equivalence classes. Our proposal for *find*-type subjective attitude verbs is that they require the opinion at play to be radically counterstance contingent.

Any proposition that is radically counterstance contingent is also (merely) counterstance contingent, but not the other way around. Our analysis thus captures the fact that *find* is more selective than *consider*. And clearly, radical counterstance contingency runs deeper than mere counterstance contingency because it persists even if there is agreement *both* on the facts *and* on issues that can be stipulated away. This feature of the analysis supports an account of why it is that predicates that embed under *find* are felt to be “more subjective” than those that embed only under *consider*, and also forms the basis for an account of the difference in meaning between *find* and *consider* attributions involving predicates that embed under both, as we will discuss in Section 5.3.

But in virtue of what does a particular expression call for a pragmatic enrichment that does or does not naturally allow for coordination by stipulation? What is it about the predicates in (35) that leads to the possibility of coordination by stipulation, and what is it about the predicates in (36) that makes such a move unnatural? It is tempting to say that those aspects of meaning that cannot be stipulated are precisely those that relate to individual experience (cf. Muñoz 2019, p. 274, fn. 44), but we have seen that not all expressions that embed under *find* are experiential — at least not in a way that does not strip that notion of any theoretical value beyond its utility for accounting for subjective attitude verbs.⁷

Our answer to these questions takes inspiration from Lasersohn (2005, p. 669), who suggests that discourse contexts leave the truth-conditions of certain natural language constructions underdetermined in a distinct way. Adopting the standard view from Kaplan (1989), Lasersohn takes predicate extensions — and so truth — to be sensitive to various parameters, complete valuations of which correspond to different contexts qua formal objects. While concrete utterance situations typically determine a unique value for many of these parameters (e.g., the speaker), they do not do so for all of them. In

⁶Not surprisingly, given the observations in Section 2, a single predicate may show different behavior depending on how it is used. For example, the adjective *heavy* patterns with the predicates in (35) when it is used quantitatively, and with those in (36) when it is used qualitatively:

- (i) For present purposes,
 - a. let’s count any suitcase that weighs more than 50kg as heavy.
 - b. # let’s count any dessert that contains more than 4tbs of butter as heavy.

⁷One reason to avoid this outcome is that, as Bylinina (2017) shows, there are good reasons to think that “having experiential semantics” defines a natural class of expressions, relative to various grammatical phenomena, which includes predicates of personal taste but excludes other predicates that embed under *find*, such as evaluative and aesthetic predicates.

particular, Lasersohn claims that concrete situations of utterance fail to single out anyone in particular for being the judge, a parameter that is relevant for fixing the extensions of (among other expressions) predicates of personal taste like *tasty*. In fact *any* individual, regardless of whether or not they are participating in the conversation, counts a potential candidate for filling that role, from which it follows that no utterance situation uniquely determines the extension of a judge-dependent predicate: there are as many possible ways of answering whether something counts as *tasty* in a concrete utterance situation as there are possible judges. Lasersohn famously gives this idea a relativist gloss, but the point is perfectly general: some aspects of (formal) context that matter for truth or falsity are *essentially* underdetermined, in the sense that their underdetermination persists *no matter* the particulars of the utterance situation.

Other aspects of context, in contrast, carry their underdetermination only *incidentally*. For example, assuming the standard view that whether an individual counts as *tall* depends on whether their height exceeds a contextual threshold, we may conclude from the vast literature on vagueness that the relevant facts about particular situations of utterance — the salient comparison class, the interests or expectations of the discourse participants, etc. — indeed fail to identify a unique height for the threshold. But such facts, together with general principles of informativity, nevertheless constrain the threshold in systematic and predictable ways (see e.g. [Lassiter and Goodman 2014](#); [Qing and Franke 2014](#); [Lewis 2020](#)). So it may be that a particular situation of utterance fails to settle whether someone with a height of 1.75 meters is tall or not, but those same facts may very well settle whether someone with a height of 2 meters is tall and whether someone with a height of 1.5 meters is not tall. At the same time, we can easily think of alternate contexts, involving, say, different comparison classes, in which it is settled whether an individual with a height of 1.75 meters is tall or not.

If all of this is correct, then it is easy to see why some instances of underdetermination would resist a proposal to coordinate by stipulation. A discourse move of this kind is essentially an invitation to transition to a context in which a single determinate stance definitively resolves the issue under consideration in one way or another; but if the underdetermination at play is an essential one, alternative stances are salient no matter the context, and thus the discourse move is bound to fail. In contrast, transitioning to a context with a single determinate stance on the matter is at least a possibility if the underdetermination at play is an incidental affair. In brief, the proposal is that propositions can, in context, be radically counterstance contingent insofar as their truth or falsity is sensitive to a criterion that is not merely incidentally but essentially underdetermined, and whenever this is so we expect that underdetermination to resist coordination by stipulation.

One may still want to know why natural languages should manifest the suggested distinction between incidental and essential underdetermination. Here we start from the observation that the resolution of a particular instance of underdetermination in one way rather than another, in context, involves privileging a particular stance over alternatives. But not all stances may be afforded such privilege. In particular, we would like to suggest that those criteria which are essentially underdetermined are precisely those which language users resist privileging in this way. The intuition is that it is fairly “cheap” to privilege a stance in which someone with a height of 1.75 meters or greater is tall or one in which someone who eats mollusks but no other animal is vegetarian. But it is rather more costly to privilege a stance in which a certain type of linguistic

puzzle stimulates an experience of fascination, or one in which a particular set of features supports a positive aesthetic judgment about a painting, because doing so may very well mean ascribing a quality to an object that is incompatible with one’s own attitudes, dispositions or experiences. Doing so is not impossible — indeed, this is plausibly what is going on in exocentric uses of predicates of personal taste — but it is marked, and requires a certain amount of accommodation. One could imagine cashing out the relevant notion of cost here in either cognitive or social terms (or both) — something we do not have the space for here — but the core idea is that essential underdetermination can be viewed as a sort of “leveling of the linguistic playing field,” which ensures that, for certain kinds of meaning-determining criteria, no perspective can take priority over others. (See [Barker 2013](#), [Lasersohn 2017](#), ch. 11, and [Grinsell 2017](#) for similar ideas.)

3.3 Formal implementation

This section provides a formal characterization of the counterstance space, counterstance contingency, and the semantics of subjective attitude verbs with atomic complements. Semantic values are assigned relative to possible worlds and in light of some contextual parameter (w_c is the world of the context of utterance). An information state $i \subseteq W$ is a set of possible worlds and we write “ $i \models p$ ” whenever i is *committed* to the proposition p , i.e. for all $w \in i$, $w(p) = 1$ (p is true at all possible worlds in i). As usual, W is the set of possible worlds and we let I be the set of all information states (i.e. the powerset of W , which is also the set of all propositions). An *issue* is a partitioning of W . Given some proposition p , \bar{p} is the negation (i.e. complement) of p . The issue as to whether p is true is *resolved* by some i , $i \triangleright p$, just in case $i \models p$ or $i \models \bar{p}$.

Each possible world maps, among other things, each constant of our target language to its world-invariant denotation $d \in D$ (here and throughout we will not mark the difference between a constant and its denotation to simplify the notation). We assume that our semantic models provide some doxastic accessibility relation. Fix some subset $D_a \subseteq D$ as the set of *doxastic agents*. Then $\text{Dox}: (D_a \times W) \mapsto I$ assigns to each doxastic agent a and possible world w the set of possible worlds compatible with what a believes at w . We define the standard truth-conditions for belief attributions on that basis (here and throughout we assume that a denotes a doxastic agent):

$$(37) \quad \llbracket a \text{ believes } \phi \rrbracket^{c,w} = 1 \text{ iff } \text{Dox}(a, w) \subseteq \llbracket \phi \rrbracket^c$$

So far, so familiar. The key additional component is a contextually determined COUNTERSTANCE SELECTION FUNCTION which maps an individual and a proposition to a set of pairwise disjoint sets of doxastic alternatives. Each such alternative — each counterstance — represents a distinct pragmatic stance, in the sense described in Section 3.1: it agrees with all other alternatives about the salient facts of the matter, but differs in the judgments it renders about those facts, in ways that, in particular, represent legitimate decisions about linguistic practice in the presence of semantic and pragmatic underdetermination. And each *set* of alternatives contains all those counterstances that agree on decisions about incidental underdetermination — those aspects of pragmatic stances allow for coordination by stipulation — and disagree on decisions about essential underdetermination, for the reasons articulated in Section 3.2.

The formal details of the counterstance selection function are spelled out in (38).

(38) **Definition: Counterstance selection function**

A contextually determined *counterstance selection function* $\kappa_c: (D_a \times I) \mapsto 2^{\mathcal{P}(I)}$ maps each doxastic agent and proposition to a set of sets of information states (a COUNTERSTANCE SPACE) such that for all $\pi, \pi' \in \kappa_c(a, p)$, $\pi \cap \pi' = \emptyset$. We call each $\pi \in \kappa_c(a, p)$ a *cell* in $\kappa_c(a, p)$. The set of counterstances to some individual a 's doxastic state in c is defined as $\mathcal{C}_c(a) = \{i \in I: i \in \bigcup \kappa_c(a, p) \text{ for some } p \in I\}$. We will say that a counterstance is *live* in context c just in case it is a counterstance to some individual's doxastic state in c . We drop the subscripts in expressions such as " κ_c " whenever this is harmless.

We will say more about how the arguments of a counterstance selection function constrain the objects in a counterstance space in the next section. But here and throughout we shall assume that context highlights a set of issues as matters of fact and that all counterstances to some individual a 's doxastic state will share a 's commitment (or lack thereof) to these issues:

(39) **Constraint: Preservation of Matters of Fact**

Let M be an issue that counts as a matter of fact in context c and let $\mathcal{C}_c(a)$ be the set of counterstances to a 's doxastic state: then for every $p \in M$ and $i \in \mathcal{C}_c(a)$: $\text{Dox}(a, w) \models p$ iff $i \models p$.

We will not attempt to offer a theory of what makes an issue a matter of fact in discourse, not least because this is a context-sensitive affair. For our purposes, it suffices to assume that language users have no trouble identifying a set of issues that are taken as non-negotiable in a concrete discourse situation, for instance: how tall someone is, the physical description of some painting, the events leading up to the 2014 Crimean status referendum, the history of Siberia, and so on. Whatever these issues are, all counterstances to some individual's doxastic state must agree on them in the sense that they resolve — or fail to resolve — these issues just as the individual's beliefs do.⁸

The two notions of counterstance contingency that we introduced informally in Sections 3.1 and 3.2 can then be stated more precisely as follows.

(40) **Definition: (Radical) counterstance contingency** Take any proposition p and counterstance space K :

- i. p is *counterstance contingent* with respect to K iff $i \models p$ and $j \models \bar{p}$ for some $i, j \in \bigcup K$.
- ii. p is *radically counterstance contingent* with respect to K iff for all $\pi \in K$: $i \models p$ and $j \models \bar{p}$ for some $i, j \in \pi$.

We will say that a doxastic agent a 's commitment to p is (radically) counterstance contingent in context c , or that a has a (radically) counterstance commitment to p in c , just in case p is (radically) counterstance contingent with respect to $\kappa_c(a, p)$ and a believes that p .

⁸Note that in fixing a set of issues that count as matters of fact, contexts constrain but do not determine the set of counterstance contingent propositions. Given (39), issues that are accepted as matters of fact do not allow for counterstance contingent resolutions. However, some issues may only admit of a clear answer, and so fail to be counterstance contingent — such as whether an individual with a height of 2 meters is tall — and yet not count as a matter of fact in the relevant sense. This fact will play a role in part of our derivation of familiarity inferences in Section 4.1 below.

Let us illustrate the proposal using a very simple model consisting of four counterstances (ovals) containing two indices apiece (circles), grouped together in two equivalence classes (round-cornered rectangles). A counterstance space need not include the relevant agent’s information state — more on this momentarily — but may do so and here it is located in the solid-edged oval in the top left corner.

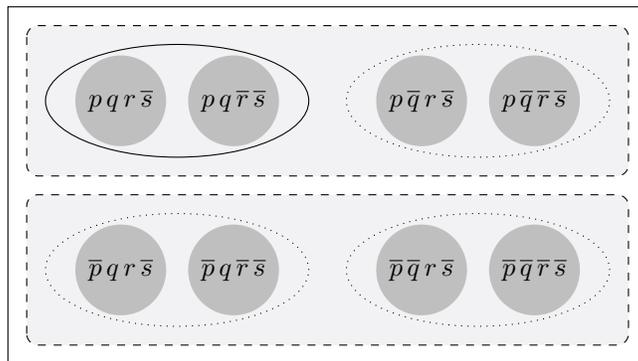


Figure 1: A counterstance space

The proposition r as well as its negation fail to be counterstance contingent since no counterstance is committed to them: all counterstances are agnostic about r . The proposition \bar{s} fails to be counterstance contingent as well, but for a different reason: all stances reject s . The commitment to \bar{s} thus differs from the commitment to p , which fails to be preserved across counterstances; at the same time, there is no variation in commitment to p across counterstances within a single cell: in the top cell all counterstances accept p , and in the bottom cell all reject p . So p is counterstance contingent, but not radically so. Finally, q is radically counterstance contingent since in every cell — and thus despite agreement on matters that allow for coordination by stipulation — there are counterstances that vary in their commitment to q .

The basic proposal then is that *consider* and *find* are in their core at issue content just like *believe* — they express doxastic attitudes — but differ in their presuppositions. The former presupposes that its complement is counterstance contingent, while the latter presupposes radical counterstance contingency. The following clauses articulate these claims more precisely:⁹

⁹As noted in section 2, Muñoz (2019) argues that only *consider* has doxastic at-issue content, not *find*. This claim is based on the putative contrast in (i).

- (i) Alphonse doesn’t find/#consider licorice tasty, but he thinks that it is.

According to Muñoz, the *find* variant is acceptable in a context in which Alphonse is an autocentric evaluator who has forgotten what licorice tastes like, and so holds a false belief, but the *consider* variant is not. We agree that the *find* sentence can have this interpretation, but we do not see a strong difference between *find* and *consider* here, especially if we spell out the details of the context a bit more:

- (ii) a. Alphonse doesn’t find/consider licorice tasty. But he’s forgotten what it tastes like, and now he mistakenly thinks/believes that it *is* tasty.
- b. Alphonse thinks/believes that licorice is tasty. But that’s because he’s forgotten what it tastes like. In fact, he doesn’t find/consider it tasty.

- (41) a. $\llbracket a \text{ considers } \phi \rrbracket^{c,w}$ is defined only if $\llbracket \phi \rrbracket^c$ is counterstance contingent with respect to $\kappa_c(a, \llbracket \phi \rrbracket^c)$.
 b. If defined, then $\llbracket a \text{ considers } \phi \rrbracket^{c,w} = 1$ iff $\text{Dox}(a, w) \subseteq \llbracket \phi \rrbracket^c$
- (42) a. $\llbracket a \text{ finds } \phi \rrbracket^{c,w}$ is defined only if $\llbracket \phi \rrbracket^c$ is radically counterstance contingent with respect to $\kappa_c(a, \llbracket \phi \rrbracket^c)$.
 b. If defined, then $\llbracket a \text{ finds } \phi \rrbracket^{c,w} = 1$ iff $\text{Dox}(a, w) \subseteq \llbracket \phi \rrbracket^c$

In ordinary contexts c , an agent’s doxastic commitment to $\llbracket \text{Lee is vegetarian} \rrbracket^c$ is counterstance contingent but not radically counterstance contingent, while a doxastic commitment to $\llbracket \text{Lee is fascinating} \rrbracket^c$ would be radically counterstance contingent. This is why simple predications of *vegetarian* are acceptable under *consider* but not under *find*, and why simple predications involving *fascinating* are acceptable under both types of subjective attitude verbs. Relatedly, in ordinary contexts c , an agent’s doxastic commitment to $\llbracket \text{Siberia is part of Russia} \rrbracket^c$ fails to be counterstance contingent, whereas in contrast a doxastic commitment to $\llbracket \text{Crimea is part of Russia} \rrbracket^c$ would be counterstance contingent (but not radically so), which explains why the latter, but not the former, is acceptable under the scope of *consider* (but not *find*).

3.4 Summary

Summarizing, we have proposed that subjective attitude verbs presuppose that their complements are subjective in a specific way: the truth of their complements is guaranteed to vary across a set of alternatives to the attitude holder’s doxastic state, all of which agree on the salient facts of the matter but disagree on judgments about those facts. One appealing feature of this analysis is that it situates the explanation of the selectional properties of subjective attitude squarely within a broader set of analyses of the selectional properties of modals and other attitude verbs, which are based on the idea that such expressions introduce different kinds of contingency or “non-settledness” conditions on their complements. Such analyses include von Fintel and Gillies’ (2010) evidential analysis of epistemic *must*, Condoravdi’s (2002) analysis of the distribution of future-oriented interpretations of possibility modals, Giannakidou and Mari’s (2017) analysis of the future as an epistemic modal, and Giannakidou and Mari’s (2015) analysis of the distribution of indicative vs. subjunctive mood in the complements of emotive attitude predicates.

Stepping back, on the story told here, subjective attitude verbs have their place in language because ordinary speakers know that the conventions of (in particular) linguistic practice often fail to conclusively settle predicate extensions: even if the facts are settled, opinions may differ because underdetermination allows for the possibility that doxastic

On our account, these facts do not indicate that *find* and *consider* fail to entail belief in the prejacent but rather that one may come to believe a proposition based on hearsay but fail to believe that proposition — or even come to believe its negation — based on direct acquaintance with the item under consideration. Braun (1998) argues that there is nothing irrational about this in his defense of the Russelian analysis of propositional attitude reports.

That said, our analysis is compatible with the at-issue content of *find* and *consider* being stated in terms of something other than vanilla doxastic attitudes, and our account of the selectional properties of subjective attitude verbs does not rely on any assumptions about their at-issue content at all. However, it is crucial to our analysis of their familiarity inferences that their at-issue meanings at least include a doxastic component, so for simplicity, we will continue to characterize this part of their meaning as identical to that of *believe*.

agents make judgments in incompatible ways. Whether this possibility is real may depend on what the facts in question are: we take an individual’s belief that Crimea is part of Russia to be counterstance contingent, for instance, precisely because we know that the situation on the ground is complicated. Subjectivity, as we have characterized it here, is a fundamentally pragmatic phenomenon, since language users generate counterstances and thus counterstance contingency in context using world knowledge. And finally, since the presence or absence of counterstances is a pragmatically determined feature of the discourse context, so too is the distinction between “objective” and “subjective” predicates, at least insofar as this distinction answers to the embeddability under subjective attitude verbs.

Let us reiterate here as well that in attributing a counterstance contingent belief to a doxastic agent, one does not thereby attribute to that agent the recognition that the truth of their belief hinges on the kinds of factors that determine a pragmatic stance, rather than on the facts. To characterize an agent a ’s commitment to p as counterstance contingent is to say that p is counterstance contingent with respect to a distinct counterstance space $\kappa_c(a, p)$. What space that is, and thus whether p is counterstance contingent, depends on what the attributee a believes, for the simple reason that the counterstances within that space are alternatives to *that individual’s* doxastic state. But of course it also depends on the contextually determined counterstance selection function κ_c , and we have argued that κ_c is not an *internal* affair, in the sense of being determined by what the *attributee* takes to be a matter of discretion, but is rather an *external* matter, depending on, most prominently, what the *attributor* takes to be a matter of discretion. As we noted in Section 3.1, in saying that Kim considers herself a Chicagoan, for instance, the attributor represents Kim’s belief as involving a leap from the facts, *regardless* of whether Kim herself sees it that way. The technical apparatus developed in Section 3.3 is set up to allow for this interpretation: κ_c is not fixed by the attributee a in any way, and so p may be counterstance contingent with respect $\kappa_c(a, p)$ even if a does not recognize the matter as such.

Subjective attitude ascriptions thus have an at-issue content that is fixed by the attributee’s doxastic state, and a not-at-issue content that depends on matters external to the attributee’s doxastic state. The former requires the attributee to have a doxastic commitment towards the prejacent; the latter requires that the “factual content” of the attributee’s belief state, as the *attributor* sees it, fails to ensure this commitment. So even if one thinks that doxastic commitments must in some sense be transparent to the attributee, that a commitment is counterstance contingent need not be, at least as long as the attributee is not a speech act participant.¹⁰

¹⁰Our analysis thus leads us to expect a difference between examples involving third person subjects of subjective attitude attributions and those in which a subjective attitude is *self-ascribed*. To see that this is so, consider an argument over balls and strikes in baseball. Although the rules of the game provide a definition of the strike zone as the area over home plate from the midpoint between a batter’s shoulders and the top of the uniform pants and a point just below the kneecap, when the batter is in their stance and prepared to swing at a pitched ball, it is recognized that the umpire has some discretion in deciding how sharp these boundaries are, and so how to categorize a pitch as a strike or a ball. It is also explicitly stipulated in Rule 8.02a that “any umpire’s decision which involves judgment, such as, but not limited to, whether a batted ball is fair or foul, whether a pitch is a strike or a ball, or whether a runner is safe or out, is final. No player, manager, coach or substitute shall object to any such judgment decisions.” Against this background, compare the reactions in (ia) and (ib) of a batter to a called third strike:

- (i) a. I think that pitch was a ball!

Nothing we have done so far shows that the key ideas of our proposal for subjective attitude verbs cannot be elaborated in other ways. But it is not trivial to get the details right, especially if we care for a framework that gets not only the facts about *find* but also those about *consider* straight, and does so in a principled way. Coppock (2018), for instance, suggests that the distinction between counterstance contingency and radical counterstance contingency can be captured in an outlook based framework as follows. Recall that a proposition is objective if its truth-value does not vary across outlooks, and that it is discretionary just in case it fails to be objective. We may then say that a proposition is counterstance contingent if it is discretionary, and we can say that it is radically counterstance contingent just in case it is strongly discretionary in the following sense: for every possible world, there is an outlook at which the proposition is true and an outlook at which it is false (Coppock 2018, pp. 133-4).

Coppock’s distinction may track something important but it does not capture counterstance contingency the way it is used here to explain the full set of embedding facts about *consider*. Recall in particular the contrast between (32a) and (32b), repeated below.

- (32) a. ✓ Kim considers Crimea part of Russia.
 b. # Kim considers Siberia part of Russia.

The outlook based approach can model this data by stipulating that the proposition that Siberia is part of Russia varies only across worlds, while the proposition that Crimea is part of Russia varies within worlds, across outlooks. But as Coppock herself acknowledges (2018, p. 134), the formal framework has nothing to say about why this is so — for this or any other case. The counterstance approach, in contrast, provides the answer: the truth of the proposition that Crimea is part of Russia varies with different outcomes of unresolved socio-political decisions about sovereignty; the truth of the proposition that Siberia is part of Russia does not.

Furthermore, as we will argue in the next section, the counterstance approach provides a basis for answering the question we began with: how to explain the familiarity inferences associated with certain uses of subjective predicates. We begin with our account of the familiarity inferences associated with uses of subjective attitude verbs, and then turn to the case assertions of unembedded taste predicates.

-
- b. I consider that pitch a ball!

(ia) might get the batter ejected from the game, because it comes across as a real objection (albeit a weaker one than unembedded “*That pitch was a ball!*”, which would certainly result in ejection). (ib), on the other hand, just sounds weird, because it amounts to an acknowledgement that the issue is a matter of discretion, and yet it is one that the batter has no authority over. In contrast, a spectator observing a batter’s surprised reaction to a called third strike use either (iia) or (iib) without suggesting that the batter is confused about who decides about balls and strikes.

- (ii) a. The batter thinks that pitch was a ball.
 b. The batter considers that pitch a ball.

For related reasons, we also expect there to be noteworthy differences between third person and second person subjective attitude attributions, but we leave spelling out the details to another day.

4 Familiarity inferences explained

In this section, we show how a pragmatic theory of subjective meaning based on counterstance contingency provides the basis for an account of the familiarity presuppositions of subjective attitude verbs and, with the addition of some additional assumptions about the expressive force of speech acts, for an account of acquaintance inferences in matrix assertions of predicates of personal taste. The central idea, which we flesh out in what follows, is that the counterstance selection function is not unconstrained, but is subject to two contextual constraints, which we call STRONG OPINIONATEDNESS and AFFECTIVE GROUNDING, and which together derive familiarity inferences. The former, we show below, follows directly from the fact that counterstances agree on matters of fact, and is responsible for the fact that subjective judgments imply that the attitude holder must have rather specific beliefs about the facts on which the truth of their judgment turns: the height of an individual that one considers tall, the ingredients of a dish that one considers vegetarian, and so forth. The latter accounts for the additional fact that, depending on the predicate, subjective judgments further imply that the attitude holder’s beliefs about the facts have come about in a particular way: that one’s beliefs about the taste of sea urchin are formed by tasting it, rather than being told about it, for example. Affective Grounding does not follow directly from the axioms of counterstance contingency, but, we claim, neither is it *ad hoc*. Instead, it reflects the interaction between lexical semantics and context in a pragmatic theory of subjective meaning such as ours, in which counterstances model doxastic agents’ sensitivity to alternative information states each of which constitutes a legitimate pragmatic stance on (among other things) the use of particular predicates.

4.1 Strong Opinionatedness

To say that a doxastic agent a resolves the issue as to whether q is just to say that a ’s doxastic state entails q or entails \bar{q} . We claim that there is a formal connection between contextual constraints on the set of live counterstances and constraints on how opinionated a doxastic agent must be for one of their beliefs to qualify as counterstance contingent. Let us begin by defining what it takes for a proposition to settle a question in context:

(43) **Definition: Settling the Question**

Let p and q be propositions: p settles the question as to whether q is true in context c iff for all $a \in D_a$ and for all $i \in \mathcal{C}_c(a)$, $i \cap p \models q$, or for all $a \in D_a$ and for all $i \in \mathcal{C}_c(a)$, $i \cap p \models \bar{q}$.

The intuition here is that in some cases, the truth of p renders q a “clear case.” This happens in particular whenever all counterstances that are live in context give the same answer to the question as to whether q is true, given p . For example, in a context in which Kim is a professional basketball player, the truth of the proposition that *Kim’s height is 215 centimeters* settles the question as to whether *Kim is tall* is true, because there are no live counterstances in which basketball players with a height of 215 centimeters fail to count as tall: such individuals are clear cases of tall basketball players.¹¹ Likewise the truth of *Kim’s height is 175 centimeters* also settles this question, because there

¹¹In the 2019-20 season, the average National Basketball Association player was 198.8 cm tall.

are no live counterstances in which such individuals count as tall: they are clearly not tall basketball players. But the truth of *Kim’s height is 200 centimeters* does not settle the question as to whether *Kim is tall* is true, because there are live counterstances in which such individuals are tall and ones in which they are not: they are “borderline tall” basketball players.

We can now derive the following fact:

(44) **Fact: Strong Opinionatedness**

If some issue M counts as a matter of fact in context, and if one resolution p of M settles the question as to whether q is true, then an agent a ’s commitment to q cannot be counterstance contingent unless a ’s doxastic state rules out p .

In other words, an agent’s commitment to q cannot be represented as counterstance contingent unless their doxastic state is understood to rule out any matter of fact that would render q a clear case. To see why, suppose that $\text{Dox}(a, w)$ is compatible with p . Then each counterstance to $\text{Dox}(a, w)$ must be so as well, given Preservation of Matters of Fact (39). And if p settles the question as to whether q is true, all counterstances must either (i) support q when strengthened with p or (ii) reject q when strengthened with p . If (i) holds, no counterstance fully rejects q ; if (ii) holds, no counterstance fully accepts q . In either case, there are no two counterstances i and j such that $i \models q$ and $j \models \bar{q}$, and so q cannot be counterstance contingent.

This is just to say that in representing some opinion as counterstance contingent, one is imposing distinct descriptive constraints on the attributee’s doxastic state: as far as the attributor is concerned, *the facts on the ground, as they are represented by the attributee, must not already settle the question*. As a consequence, attributions of counterstance contingent beliefs characterize the attributee as opinionated in a way that is distinct from, and indeed stronger than attributions of plain beliefs, in a way that captures a good part of what we have been calling familiarity inferences.

Let us demonstrate this by working through a concrete example:

- (45) a. Yara believes Zach to be tall.
 b. Yara considers Zach tall.

Given standard assumptions about the truth conditions of gradable predicates, both (45a) and (45b) assert that Yara believes Zach’s height h_z to meet or exceed the contextual standard for tallness θ_c : for all $w \in \text{Dox}(y, w_c)$, $h_z(w) \geq \theta_c$. (45b), in addition, presupposes the counterstance contingency of this belief: there must be at least one counterstance committed to Zach being tall, and another committed to Zach not being tall. Assuming that Zach’s height counts as a matter of fact in context, all counterstances must share Yara’s beliefs about what this is. And assuming that the value of the threshold is constrained by the context but (incidentally) underdetermined, counterstances will vary in this value. It is this variation in thresholds across counterstances that allows for Yara’s attitude to be potentially counterstance contingent.

For instance, suppose that Yara has actually seen Zach, and has formed the belief that his height is between 175 cm and 177 cm; suppose also that the contextually viable range of threshold values, determined by a salient comparison class, is between 175 cm

and 180 cm.¹² Such a context will generate (among others) counterstances corresponding to different valuations of θ_c , but given Preservation of Matters of Fact, each of these counterstances must agree on Yara’s beliefs about Zach’s height, which is just to say that each contains only worlds in which Zach’s height is in the 175–177 cm range. In the table in (46), each column represents a counterstance corresponding to values of $\theta_c \in \{\theta_{175}, \theta_{176}, \theta_{177}, \theta_{178}, \theta_{179}, \theta_{180}\}$, and the rows indicate the truth of the proposition that Zach is tall in worlds that agree with Yara’s representation of the facts $\{w_{175}, w_{176}, w_{177}\}$.

(46)	$\llbracket h_z(w) \geq \theta_c \rrbracket = ?$	θ_{175}	θ_{176}	θ_{177}	θ_{178}	θ_{179}	θ_{180}
	w_{175}	1	0	0	0	0	0
	w_{176}	1	1	0	0	0	0
	w_{177}	1	1	1	0	0	0

In this context, Yara’s belief counts as counterstance contingent because the counterstance based on θ_{175} guarantees its truth and any of the ones based on θ_{178} , θ_{179} or θ_{180} guarantee its falsity. (The ones based on θ_{176} and θ_{177} are agnostic).

Let us now consider a context in which the comparison class and range of contextually viable thresholds remain the same, but it is made explicit that Yara does not have specific information about Zach’s height (cf. the examples we used in Section 2.2 to first illustrate familiarity inferences with *consider*):

- (47) The only thing Yara knows about Zach is that he comes from a family of tall people, so ...
- a. # ... she considers him tall.
 - b. ... she believes that he’s tall.

Here we have no basis for treating Yara as having the strong opinion about Zach’s height that she had above. At most, we can assume her opinion to be based on general statistical knowledge about height distributions in tall families, which means that we cannot exclude worlds from Yara’s doxastic state in which Zach has a height of, say, 190 cm. But this means that even the most demanding counterstance — the one based on θ_{180} — is now open to the possibility that Zach is tall, as shown in (48).

(48)	$\llbracket h_z(w) \geq \theta \rrbracket = ?$	θ_{175}	θ_{176}	θ_{177}	θ_{178}	θ_{179}	θ_{180}
	w_{190}	1	1	1	1	1	1

Yara’s belief fails to be counterstance contingent, and so use of *consider* in (48b) is correctly predicted to be unacceptable.¹³

¹²Here we treat beliefs about heights and the range of potential thresholds as categorical for simplicity, but in actual discourse situations, both values are probabilistic (see e.g., Schmidt, Goodman, Barner, and Tenenbaum 2009; Lassiter and Goodman 2014, 2015; Qing and Franke 2014; Goodman and Frank 2016; Qing 2020; Xiang, Kennedy, Xu, and Leffel forthcoming; Xiang, Kramer, and Kennedy 2021). This does not change the general point we are making here.

¹³An anonymous reviewer worries that Strong Opinionatedness as we have formulated it here is too strong, pointing to cases like (i), which they judge to be more acceptable than (47), particularly if there is background information such that Yara assigns high probability to Zach’s height being between 176cm and 180cm.

- (i) Yara doesn’t know Zach’s exact height. But she knows he’s taller than Xena, who she knows to be 175cm tall, so she considers him tall.

We agree with this judgment, but claim that it is precisely such background information about Yara’s beliefs that allow Strong Opinionatedness to be satisfied. We will say more about this below.

This example involves the gradable predicate *tall*, a paradigmatically vague term, but the account generalizes to all those expressions that we characterized as “incidentally” underdetermined expressions in Section 3.2, i.e. expressions that have conventionalized clear cases (in context) but also afford a borderline area of application in which the conventions for use of the expression do not return a verdict. These include in addition to gradable adjectives, expressions such as *athlete*, *planet*, and *gluten-free*:

- (49) a. Lee considers Secretariat an athlete.
- b. Lee believes that Secretariat is an athlete.
- (50) a. Kim considers Pluto a planet.
- b. Kim believes that Pluto is a planet.
- (51) a. Pat considers the soup gluten-free.
- b. Pat believes that the soup is gluten-free.

In all of these examples, the (a) sentences involving *consider* come with familiarity inferences that are lacking in the corresponding (b) examples involving *believe*: in the former, the attitude holder’s beliefs about the relevant objects should be detailed enough to render them borderline cases relative to the corresponding predicates; in the latter, this need not be the case. Simply being told that Secretariat is an athlete may be enough to justify the belief that this is so, for example, but it does not by itself provide an appreciation of the facts that make the case a matter of controversy and licenses use of *consider*. Instead, one should also have learned e.g. that Secretariat, an exceptional runner, is not human (unlike Olympic sprinters, clear athletes), but participates in organized sporting events (unlike cheetahs, clear non-athletes).

In sum, a belief that an object x has the property denoted by a predicate P counts as counterstance contingent only when the attributee’s beliefs about the facts on the ground are specific enough to put x squarely in (what the attributor takes to be) the contextually determined borderline region for application of P . This constraint, which we refer to as Strong Opinionatedness, flows directly from Preservation of Matters of Fact, the condition that all counterstances agree on matters of fact: for, if the facts (as represented by the attitude holder) *were* compatible with x being a clear (non-) case of P , then the proposition that x is (not) P would be compatible with all of the attitude holder’s counterstances, and so would fail to be counterstance contingent. This, we submit, accounts for the general empirical observation that attitude ascriptions involving *consider* impose stronger familiarity constraints on the attributee than do the corresponding plain belief attributions.¹⁴ That said, Strong Opinionatedness on its own does not derive the full set

¹⁴In the cases we have been considering, in which a singular term is used to attribute a belief about a particular object, familiarity inferences stemming from Strong Opinionatedness manifest as beliefs about those objects (Zach’s height, Secretariat’s athleticism, etc.). But this need not be the case. Consider (i).

- (i) Kim considers every full professor old.

This example has an interpretation in which *every full professor* takes scope over *consider*, attributing to Kim singular beliefs about every individual in the (contextually restricted) domain of quantification and implying that Kim has specific beliefs about each individual’s age. But it also has an interpretation in which *every full professor* takes scope under *consider*, attributing to Kim a general belief about full professors. On this reading, it does not require that Kim have specific beliefs about any individual full professors, but it *does* require that she have specific beliefs about the distribution of ages of these individuals, and from the attributor’s perspective, this distribution should be such that it is borderline

of facts that we have thus far treated as instances of familiarity inferences. In particular, it does not capture the fact that, when predicates of personal taste and other “evaluative” expressions are embedded under *find* and *consider*, the attributee must also have formed their strong opinion in a certain way. We address this issue in the next section, but two additional remarks are in order before moving to this discussion.

First, it is important to stress once more that for an attribution of a counterstance contingent attitude to be felicitous, it is the *attributor* who must take the attributee’s belief to be based on information that leaves the case underdetermined; the *attributee*, in contrast, may very well think that the case is clear. An utterance of (50a) commits the speaker to the position that Kim’s beliefs about Pluto leave open its status as a planet, for example depending on whether one accepts the International Astronomical Union’s 2006 re-definition of the term, but Kim may be entirely unaware of such considerations, and consider Pluto a clear case. And furthermore, to imply that the attributee is strongly opinionated about the underlying facts is not to endorse the relevant opinions: as far as the attributor is concerned, the attributee’s judgment may have been formed based on faulty reasoning or misinformation. To say (51a), for example, is to say that Pat has specific opinions about what is in the soup and/or how it was made, but it is not to say that these opinions are correct.

Second, one may wonder how Strong Opinionatedness, as we have derived it here, squares with cases such as the following.

(52) Jane is at least 18, so I consider her an adult.

Here one might worry that (52) is assertable even though the speaker does not explicitly rule out that, as far as they know, Jane could be considerably older than 18, and thus a clear case of an adult in any intuitive sense of the word. So, is Strong Opinionatedness too strong after all?

We have two responses to this concern. First, if the context is one in which the only factor relevant for categorizing someone as an adult is their age — say, a decision about who to admit to an “adults only” performance — then examples like (52) are acceptable only to the extent that one can accommodate some additional constraints on the attributee’s beliefs to satisfy Strong Opinionatedness (cf. note 13). When such supplementary assumptions are explicitly ruled out, the use of *consider* is odd, as shown by (53).

(53) # Jane is at least 18, and possibly as old as 30, so I consider her an adult.

And second, the most natural understanding of examples such as (52) is actually as an articulation of the opinion that being at least 18 is sufficient for Jane to pass as an adult, *independent of other considerations*, i.e. as a claim made in a context in which the truth of the proposition hinges on factors other than those having to do with age. This reading is particularly clear in (54), which makes explicit that the fact that Jane has the property of being at least 18 is enough to count her as an adult, even in the presence of behavioral facts that would otherwise dictate a different judgment.

(54) Jane is at least 18, so I consider her an adult despite her childish behavior.

whether every full professor is old. Similar considerations apply to other examples involving quantifiers, but a full exploration of this issue would take us well beyond the scope of this paper. (See [Ninan 2020](#), [Willer and Kennedy 2022](#) and Section 5.1 for additional discussion.)

This example highlights once again that the use of *consider* requires familiarity with the *relevant* facts, and which facts count as relevant is a context-dependent affair, dependent on what kinds of features determine what counts as a borderline vs. clear case of a predicate on different occasions of use.

4.2 Affective Grounding

In the preceding section, we observed that an agent’s judgment cannot count as counterstance contingent unless their doxastic state rules out any matters of fact that settle the question about that judgment, which in turn means that the individual must be familiar enough with the facts on which the truth of the judgment turns to ensure that it counts as a borderline case in context. This constraint, which we called Strong Opinionatedness, derives the familiarity inferences that arise from embedding incidentally underdetermined expressions under *consider*, but it does not account for the full variety of familiarity inferences we observed earlier, for two reasons.

First, when we turn to essentially underdetermined predicates — those that embed under *find* — we see that Strong Opinionatedness does no work for us. Consider an agent who loves the taste of cilantro, and believes that anything that contains cilantro is bound to taste good. One might think that such things then count as clear cases of *tasty*, but this is not the case. As we noted in Section 3.2, to say that *tasty* is essentially underdetermined is to say that there are as many ways of answering the question of whether something is tasty as there are possible judges: it is part of the conventions of the use of this expression that pretty much anything may be deemed tasty, which is just to say that all cases are borderline cases, in the sense that matters for Strong Opinionatedness. The upshot is that Strong Opinionatedness imposes little to no constraints on the doxastic state of the attributee when it comes to essentially underdetermined predicates, given that such expressions afford little to no clear cases in context.

Second, even if it were possible to impose some constraints of this sort, they would not be enough. As we have seen, embedding predicates of personal taste and other evaluative expressions under *consider* and *find* suggests not only that the attributee is strongly opinionated, but also that the opinion was formed in a distinct way, which depends on the predicate. One can read an exhaustively detailed description of the taste of cilantro, sufficient to form strong opinions about its taste — indeed, sufficient to *believe* that it is tasty (or not) — but one can neither *find* cilantro tasty (or not) nor *consider* it tasty (or not) without actual experience of its taste.

Our key suggestion in response to this issue is that essentially underdetermined predicates introduce distinct constraints on the set of counterstances available in context, which are rooted in their evaluative lexical semantics and which reflect the difference between *taking a stance* on the issues they address and merely forming an opinion about them. We begin by observing that our model not only allows us to talk about an agent’s beliefs (at a world) simpliciter but also to say which of those beliefs are *grounded* in one’s experiences or moral sentiments, adapting ideas in [Willer and Kennedy \(2020\)](#). Taking inspiration from [Franzén \(2018, 2020\)](#), we will generalize this idea to identify a broad category of AFFECTIVELY GROUNDED beliefs — those beliefs that are grounded in an affective state with the corresponding content, may it be an experience, moral sentiment, emotion, or other cognitive episode that results from “being struck” by an object or event in a certain way. What all of these affective states have in common is that their very

existence requires awareness on behalf of the subject: one cannot be struck by some object or event in some way without one’s being sensitive to at least some of its features. For a proof to strike one as elegant, for instance, one must be aware of certain properties of that proof. Our account of the particular familiarity inferences of evaluative language then boils down to the proposal that an evaluative attitude may be counterstance contingent only if the attitude is affectively grounded, and since affective grounding requires awareness on behalf of the subject, it follows that attributions of counterstance contingent attitudes in general require the attributee’s familiarity with some distinct feature of the object or activity under consideration.

Before offering our explanation of why counterstances are so constrained, let us make the proposal more precise, starting with the case of experiential grounding and the direct experience inferences that result from embedding experiential predicates like *tasty* and *fun* under a subjective attitude verb. Start by saying what it takes for a proposition to put an object in the (anti-)extension of some predicate (in context):

- (55) A proposition $p \subseteq W$ places an object $x \in D$ in the (*anti-*)extension of a predicate β in context c iff (i) for all $w \in p$, $x \in \llbracket \beta \rrbracket^{w,c}$ or (ii) for all $w \in p$, $x \notin \llbracket \beta \rrbracket^{w,c}$.

A proposition p places an object in the extension or anti-extension of a predicate just in case each world in that proposition does.

We may now state which of a subject’s beliefs are EXPERIENTIALLY GROUNDED:

- (56) **Definition: Experientially Grounded Beliefs**

Consider some p , a , and w : p is *grounded in a ’s experience* at w iff p is entailed by the content of one of a ’s (past or present) experiences at w . The function $\text{Exp}: (D_a \times W) \mapsto 2^I$ assigns to each doxastic agent a and possible world w a set of propositions $\text{Exp}(a, w) = \{p: p \text{ is grounded in } a\text{’s experience at } w\}$. If $\text{Dox}(a, w) \models p$, then a ’s belief that p is *experientially grounded* at w iff $p \in \text{Exp}(a, w)$.

Here we assume that experiences have content and as such rule out certain ways the world could be.¹⁵ For instance, experiencing some object as red is to be in a state with some representational content: it rules out certain ways the world could be, specifically those in which the object is not red. An agent’s experientially grounded beliefs are then those that are “backed up” by past or present experiential episodes with the relevant content. A belief that some object is red is experientially grounded, for instance, in case the subject is experiencing (or has experienced) the object as red.

Next we impose a formal constraint on the counterstance selection function: whenever the proposition under consideration concerns the application of an experiential predicate, a counterstance must be agnostic on the issue unless the agent’s experiences resolve it in one way or another. More precisely (recall that $i \triangleright p$ just in case $i \models p$ or $i \models \bar{p}$):

- (57) **Constraint: Experiential Grounding**

Let β be an experiential predicate and suppose that p places some object d in the (anti-)extension of β in context c , then for all a and i in $\mathcal{C}_c(a)$: $i \triangleright p$ only if $p \in \text{Exp}(a, w_c)$ or $\bar{p} \in \text{Exp}(a, w_c)$.

¹⁵The thesis that experience is fundamentally a matter of representing the world as being a certain way traces back at least to Kant (1781) — on one prominent reading anyway — and figures prominently in the work of thinkers as different as Evans (1982), Peacocke (1983), McDowell (1994), Dretske (1995), Chalmers (1996), and Byrne (2001).

Experiential familiarity inferences follow directly. According to (57), if a counterstance to an agent’s doxastic state resolves p , then the agent must have (or have had) an experience with p or with \bar{p} as its content. For instance, it is only if the agent has experienced downhill skiing as fun (or not) that his or her experiences speak to the issue as to whether downhill skiing is fun, and it is only in that case that one of the agent’s counterstances may agree and another disagree with the opinion. It then follows by common sense that embedding an experiential predicate under *find* or *consider* gives rise to a direct experience inference: doing so is to attribute an experientially grounded belief. Subjective attitude verbs presuppose that the opinion at play is counterstance contingent, and according to (57), this can only be so if the opinion is experientially (and not just doxastically) grounded. But one cannot — and here is where common sense comes in — have experience of some (kind of) object as tasty (or not) without tasting it, and one cannot have experience of some (kind of) activity as fun (or not) without participating in it. Since *believe*-attributions do not presuppose counterstance contingency, they do not trigger this kind of inference.

The familiarity inferences of moral predicates can be derived in a similar way. While the nature of moral thought is obviously a matter of debate, it seems uncontroversial that many of our moral opinions flow from (or are at least accompanied by) MORAL SENTIMENTS — sentiments such as moral approval and disapproval. The obvious suggestion would then be that moral predicates require an opinion to be grounded in moral sentiment for counterstances to resolve the issue:¹⁶

(58) **Definition: Morally Grounded Beliefs**

Consider some p , a , and w : p is grounded in a ’s moral sentiment at w iff p is entailed by the content of one of a ’s (past or present) moral sentiments at w . The function $\text{Mor}: (D_a \times W) \mapsto 2^I$ assigns to each doxastic agent a and possible world w a set of propositions $\text{Mor}(a, w) = \{p: p \text{ is grounded in } a\text{’s moral sentiments at } w\}$. If $\text{Dox}(a, w) \models p$, then a ’s belief that p is *morally grounded* at w iff $p \in \text{Mor}(a, w)$.

(59) **Constraint: Moral Grounding**

Let β be a moral predicate and suppose that p places some object d in the (anti-)extension of β in context c , then for all a and i in $\mathcal{C}_c(a)$: $i \triangleright p$ only if $p \in \text{Mor}(a, w_c)$ or $\bar{p} \in \text{Mor}(a, w_c)$.

We will not dive into the details of moral psychology here but again common sense seems sufficient to derive familiarity inferences. Clearly, moral sentiments do not require direct experience: one can disapprove of war without having experienced it. But insofar as sentiments are essentially responsive attitudes — they are reactions to some feature of an object or activity — having any kind of moral sentiment toward an object or situation requires awareness of some of its features; moreover, that feature must be morally relevant in the sense that awareness of it actually triggers a sentiment of (for instance) approval or disapproval in the subject. This is why, we suggest, embedding a moral predicate under a subjective attitude verb implies that the attributee is distinctly familiar with

¹⁶This is not to endorse a fully noncognitivist position and identify moral beliefs with moral sentiments, any more than our previous proposal for experiential adjectives identified beliefs about taste with taste experiences. In particular, nothing we have said rules out the existence of *bona fide* moral beliefs that lack moral grounding — the claim is that such opinions do not qualify as counterstance contingent and hence can only be ascribed using *believe* and *think*, not *find* or *consider*. See [Schroeder 2010](#) for an overview of the variety of noncognitivist positions in ethics.

the subject matter: counterstance contingency entails moral grounding of the attitude ascribed, which in turn entails that the attributee must possess awareness of a morally relevant feature of the object or activity under consideration.

The key maneuver here, then, is to suggest that an agent’s counterstances can resolve some issue only if the agent’s opinion is suitably grounded. Following [Willer and Kennedy \(2020\)](#), let us now add the claim that what suitable grounding amounts to is a lexical semantic matter: it depends on the kind of predication(s) involved.¹⁷ When the predicate is experiential, counterstance contingency requires experiential grounding, and requires familiarity with experientially relevant features of the predicate’s argument; when the predicate is moral, counterstance contingency requires moral grounding, and requires familiarity with morally relevant features of the predicate’s argument. But the strategy is general enough to account for the variety of familiarity inferences that we highlighted earlier, such that for any particular evaluative predicate, counterstance contingency requires grounding in the relevant affective state, and requires familiarity with those features of the predicate’s argument that are involved in triggering the relevant affective state.¹⁸

But a key question remains: *why* do evaluative predicates impose distinct grounding constraints on counterstance contingent attitudes? We do not see a way to offer a proof from first principles here; with [Franzén \(2018, 2020\)](#), we take the proposal to be closer to an inference to the best explanation. But let us offer some reason to believe that, in the broader context of our analysis of subjective meaning as sensitivity to counterstance contingency, that the proposal is not *ad hoc*. First, counterstances represent particular kinds of resolutions of issues, distinct from ordinary doxastic alternatives: resolutions whose differences track differences in pragmatic stance, in the sense that the issue is not what the facts are, but how to interpret them. Second, counterstances are, by definition, counterstances to some agent’s doxastic state. As such, we claim, they are in a position to take a stance on an issue only if the relevant agent is. And crucially, taking a pragmatic stance on an evaluative issue involves more than simply adopting an opinion. This territory has, indeed, been well covered by the literature on moral and aesthetic testimony: why does it seem strange to form a moral or aesthetic opinion — say that the death penalty is wrong or that the Eiffel Tower is beautiful — just based on another’s say-so? Everything we say here is compatible with the position that such opinions may be justified if the circumstances are right. But in line with what has been suggested by critics and defenders of moral and aesthetic testimony alike, we submit that there seems to be something lacking from opinions thus formed, in that they do not constitute responses to some relevant moral or aesthetic feature of the object or event under consideration. Testimonial reasons, in brief, are not moral or aesthetic reasons. And we add here — again in line with the existing literature — that what tells these two types of reasons is that the latter, but not the former, are distinctly emotional achievements (or failures).¹⁹

¹⁷In [Willer and Kennedy 2020](#), grounding is taken to be a lexically specified use condition on predicates; our proposal here, in contrast, treats grounding as a lexical constraint on counterstances. We will say more about the implications of this difference below.

¹⁸In fact, the strategy can be made fully general, if we follow [Willer and Kennedy \(2020\)](#) and treat doxastic grounding as the unmarked case, required by non-evaluative predicates, since counterstances, *qua* doxastic alternatives, already satisfy doxastic grounding.

¹⁹Thus [Enoch \(2014\)](#) suggests that moral deference is in some cases the only sensible response to moral uncertainty, but also notes that “forming a moral judgment by deference and then acting on it is much less of a moral achievement than forming the true judgment without deference, because it does not constitute the appropriate response to the morally relevant features of the case” (p. 255). For relevant discussion of this point see also, among many others, [Hopkins \(2007\)](#), [Hills \(2009\)](#), and [Callahan \(2018\)](#).

The case generalizes to judgments of taste and to opinions formed on the basis of inference, when the reasoning involved does not connect the reasoner to the relevant experiential features of the object or situation under consideration.²⁰ What underwrites our intuitions in these cases, we suggest, is that an evaluative judgment lacks a distinct kind of “stability” if it is not suitably grounded in the speaker’s affective attitudes, and thus does not constitute a genuine stance on the issue. And this is because, intuitively, the judgment remains susceptible to how the object or situation would actually strike the agent in case he or she came into contact with its relevant features. For instance, one may have excellent testimonial or inferential reasons for believing that cilantro is tasty — and perhaps these will even affect one’s experiences — but it would be plain weird to uphold that belief if one ends up tasting but not liking it. In contrast, if one does not like cilantro, others’ say-so or arguments to the contrary have no purchase, though of course they may give one reason to try it again. When it comes to matters of taste, individual experiences trump all other considerations.

Similarly for all evaluative judgments: even if an agent has (testimonial or inferential) reasons for holding them, they appear irrational or at least disingenuous if they do not accord with the relevant aspects of the agent’s affective state.²¹ It is in this sense, then, at least as far as the folk theory of ordinary speakers is concerned, that there is a primacy of affective grounding when it comes to evaluative attitudes, and that a legitimate stance on evaluative matters requires affective grounding. Assuming again that a judgment is counterstance contingent only if there is a stance to be taken in the first place, it is not surprising that evaluative judgments can be counterstance contingent only if suitably grounded.

Our proposal is in crucial respects similar to other recent proposals for the acquaintance inference. [Franzén \(2018\)](#) and [Willer and Kennedy \(2020\)](#) suggest that assertions express states of mind and that evaluative assertions express distinct evaluative attitudes that require familiarity with the subject matter. [Ninan \(2020\)](#) appeals to supervenient evaluations and maintains that an assertion about taste can only evaluate to super-true if they are grounded in the speaker’s experience. In our framework, we put assumptions about what is required to take an evaluative stance at the foundation and then move on to derive the familiarity inferences of evaluative assertions on that basis (see the next subsection). All of these proposals require some bedrock assumptions about evaluative language and evaluative states of mind to get off the ground. What we suggest here is to articulate these assumptions against a comprehensive background framework

[Hills \(forthcoming\)](#) articulates a view about aesthetic testimony congenial to what we have suggested about aesthetic stances; see also e.g. [Wollheim \(1980\)](#), [Meskin \(2004\)](#), and [Hopkins \(2011\)](#) for relevant discussion. [Budd \(2003\)](#) and [Livingston \(2003\)](#) are skeptical of the thesis that all aesthetic judgments require direct experience of the object under consideration, but they also suggest that such experience is required to genuinely appreciate a work of art.

²⁰As we noted earlier, one can form genuine moral opinions based on testimony as long it conveys knowledge of morally relevant features. One may come to condemn (or approve of) the United States Invasion of Panama, for instance, based on what one has read about it. Presumably, one can also come to know of such features via inference. Genuine judgments of beauty, in contrast, seem to require direct experience, while it is possible to end up finding a proof elegant after being told about it. We will not try to dig even deeper here and answer the question of why some features that allow one to form an evaluative judgment are relatable via testimony and inference while others do not.

²¹It is precisely this pressure to accord with the agent’s affective state, we suspect, that leads to essential underdetermination and the impossibility of coordination by stipulation, in line with our discussion at the end of Section 3.2: affective states, after all, cannot be stipulated into or out of existence. More work is needed to fully establish this connection, however.

that allows us to illuminate the selectional restrictions of subjective attitude verbs and supports a pragmatic theory of subjective language, i.e. one that puts counterstances at the analytical center stage.²²

4.3 Assertion

At long last we now return to the initial observation that plain assertions involving predicates of personal taste come with a direct experience implication.

- (1) I have never tried sea urchin.
 - a. # It's (not) tasty.
 - b. I hear it's (not) tasty.
- (2) I have never gone downhill skiing.
 - a. # It's (not) fun.
 - b. It must (can't) be fun.

Our simple proposal is that the assertions in (1a) and (2a) are expressions of radically counterstance contingent doxastic commitments. Since, as we have seen, a rational commitment may be radically counterstance contingent only if it is suitably grounded, and since *tasty* and *fun* select for experiential grounding in particular, it follows immediately that ordinary uses of predicates of personal taste imply that the speaker has direct experience of the item under consideration.

That such assertions *can* express subjective attitudes follows from what we have said so far, together with the widely held view that a felicitous act of assertion counts as an expression of a doxastic commitment to the truth of the proposition asserted (see, among others, [Bach and Harnish 1979](#)). The fact that *tasty* and *fun* embed under *find* tells us that the meanings of these expressions are underdetermined in a way that resists coordination by stipulation, and so are underdetermined no matter the particulars of context. It is therefore always possible, in any context, for assertions of “*sea urchin is (not) tasty*” and “*downhill skiing is (not) fun*” to be assertions of propositions that are radically counterstance contingent relative to the speaker’s doxastic state, and so to the extent that assertions are expressions of doxastic commitments, it is always possible for assertions of these propositions to be expressions of subjective attitudes.

It does not, however, follow from what we have said so far that such assertions *must* express subjective attitudes, rather than plain doxastic attitudes. The proposition that sea urchin is tasty may be entailed by an individual’s doxastic state and yet not be counterstance contingent relative to it; this is the case, for example, if this proposition is believed on the basis of hearsay alone. Our task, then, is to ensure that assertions like

²²It is for this reason as well that we reject the otherwise attractive supposition that what distinguishes English *find* is not a presupposition of direct experience (as suggested by [Stephenson \(2007\)](#) and [Muñoz \(2019\)](#) and discussed in Section 2.1) but rather the presence of an affective attitude (as suggested by [Franzén \(2020\)](#)). As we have said, we are interested in a theory of subjective attitude verbs that does not only cover *find* but also *consider*, which does not come with an affective requirement *per se*, but does give rise to affective inferences when the embedded predicate happens to be evaluative. Thus we take (radical vs. regular) counterstance contingency as key to analyzing subjective attitude verbs, and we derive familiarity inferences from constraints that lexical items impose on the set of counterstances available in context.

(1a) and (2a) are expressions of subjective attitudes *by default*. The pragmatic principle in (60) does what we want.

(60) **Maximize Commitment**

An assertion of p in context is an expression of a (radically) counterstance contingent commitment to p whenever p cannot be taken to be settled by a matter of fact in that context.

Recall from Section 4.1 that a proposition is settled by a matter of fact just in case it is a clear case in context — if all counterstances (for all doxastic agents) in context agree on it, or all agree on its negation — otherwise it is borderline. The force of (60) is therefore to ensure that assertions of borderline propositions are, by default, expressions of subjective attitudes, and as such, come with familiarity inferences for the reasons we articulated in the preceding sections.

The driving idea behind this principle is that a subjective commitment is stronger than a plain doxastic commitment, since the former is just like the latter, except that it holds only when the commitment is counterstance contingent relative to the attitude holder's doxastic alternatives. In making an assertion, a speaker expresses her commitment to the proposition asserted; (60) dictates that this commitment be as strong as meaning plus context allows. This principle targets the expressive force of an assertion rather than its at-issue content or presuppositions, but it is clearly inspired by the Gricean Maxim of Quantity (Grice 1975), which dictates use of a stronger expression over weaker alternatives, and the Maximize Presupposition principle (Heim 1991; Sauerland 2008), which (informally) requires a speaker to opt for an expression with stronger presuppositions if these presuppositions are known to be satisfied in context. There is an ongoing debate about whether these latter two principles are distinct types of constraints, or variants of a more general principle (see Lauer 2016 for discussion), and the status of Maximize Commitment should be part of this investigation. For the moment, we will treat it as an independent pragmatic constraint.

If Maximize Commitment (or something like it) is correct, then we reduce the direct experience implications of plain assertions of predicates of personal taste to those of subjective attitude ascriptions: such assertions give rise to direct experience inferences *insofar* as they are interpreted as expressing counterstance contingent commitments, and when the proposition asserted is understood as borderline in the context, the assertion must be so interpreted. As we observed in Section 4.2, propositions involving predicates of personal taste like *tasty* and *fun* and other essentially underdetermined predicates are *guaranteed* to be borderline, because there are as many ways of answering the question of whether something is tasty or fun as there are judges. Assertions involving such predicates therefore count as expressions of counterstance contingent commitments by default, according to (60), and direct experience inferences follow for the reasons that we articulated in Section 4.2, as a special case of the requirements imposed by the lexically determined grounding conditions on counterstances.

In other cases, whether the proposition asserted can be taken to be settled by matters of fact depends on the particulars of context. Consider, for example, the cases in (61).

- (61) a. Kim is Russian.
- b. Senator Jones won the debate.
- c. Lee is tall.

An assertion of (61a) may be settled by the facts, if, say, Kim is believed to be from Moscow. On the other hand, if Kim is thought to be from Crimea, it may not be, and in this case, an assertion of (61b) carries a distinct subjective flavor. (61b), in turn, may either articulate a subjective impression of how Senator Jones performed in the debate or it may state a plain fact about the post debate poll results (Pedersen 2012; Vardomsakaya 2018). Doing the latter requires that one knows the polls, but not that one actually watched the debate; voicing the subjective opinion without knowing how the debate went, in contrast, would be distinctly odd. And (61b) can be heard either as an objective characterization of Lee as a clear case, or as a subjective opinion about a borderline individual (and a corresponding commitment to where the borderline is; cf. Barker’s (2013) distinction between “metalinguistic” and “descriptive” uses of vague predicates); only the latter presumes familiarity with Lee’s height, in the sense discussed in Section 4.1. What matters in all of these cases is whether the proposition asserted can be taken to be settled by the facts, and when it cannot — when it is possible to agree on all the facts with the speaker and yet disagree about the issue — the belief expressed is a counterstance contingent one, and comes with familiarity inferences.

Everything we have said here straightforwardly generalizes beyond uses of predicates of personal taste. Aesthetic judgments, moral judgments, and the like, all afford a default interpretation as expressions of counterstance contingent attitudes, and as such inherit the familiarity inferences that we have observed when such judgments are attributed using subjective attitude verbs. But it is important to stress once more that beliefs about taste or aesthetic or moral affairs do not *have* to be counterstance contingent. Everything we have said here, in particular, is compatible with the possibility that such beliefs can be based on testimony, and for all we have said here such beliefs may very well amount to knowledge. It is for this reason that the account developed here has no need for the acquaintance principle that was critically discussed in Section 1: what matters to the account is not whether testimonial beliefs about taste, beauty, and the like may amount to knowledge; our claim is that such beliefs cannot be counterstance contingent unless they are properly grounded in direct experience or some other suitable familiarity with the subject matter.

And as outlined in the previous section, we assume that whether grounding requires direct experience or some other form of familiarity is a matter of lexical semantics. In this respect, our account is similar to the family of analyses which treat acquaintance/familiarity inferences as presuppositions (or some other kind of not-at-issue content), such as Pearson 2013; Anand and Korotkova 2018; Ninan 2020; Willer and Kennedy 2020. But the accounts diverge in their basic predictions about when familiarity inferences should appear. In presuppositional accounts, familiarity inferences are effectively constraints on the use of particular lexical items in speech acts, and as such, the default prediction is that they should appear whenever the relevant expressions are used. But, as noted in Section 1, this is not the case: there are numerous classic presupposition “holes” from which familiarity inferences fail to project, including those shown in (62).

- (62) I have never tried sea urchin, but ...
- a. ... it might be tasty.
 - b. ... it must be tasty.
 - c. ... it’s certain to be tasty.
 - d. ... it’s clearly tasty.

- e. ... it's evidently tasty.

Presuppositional accounts can accommodate such cases by introducing special “obviation” mechanisms to block projection of familiarity inferences out of certain contexts, but the challenge for such approaches, at the cost of explanatory adequacy, is to provide a comprehensive theory of when blocking happens, and why it targets only familiarity inferences and not other kinds of not-at-issue content (Muñoz 2021).

In contrast, the analysis proposed in this paper derives familiarity inferences from lexical (and contextual) constraints on the counterstance space, and so the default prediction is that obligatory familiarity inferences should be present only when counterstance contingency is at stake. Sometimes this is so for compositional reasons, as in subjective attitude ascriptions. Other times this is so for pragmatic reasons, as when Maximize Commitment is at play. In the absence of such reasons — whenever counterstance contingency is not at stake — familiarity inferences are not predicted in the first place. This is what we already saw with plain belief attributions using *think* and *believe*:

- (63) I have never tried sea urchin, but ...
 - a. ... I believe it is tasty.
 - b. ... I think it is tasty.

Here the speaker is clearly committed to sea urchin being tasty, but familiarity is not required. Semantic considerations have no purchase since *think* and *believe* do not presuppose counterstance contingency of the opinion at stake, nor does the pragmatic Maximize Commitment principle, since it targets the commitment to the propositions asserted by (63a–b), i.e. to the speaker’s *believing* or *thinking* that sea urchin is tasty. These propositions do not afford a counterstance contingent flavor in the first place, assuming that what one believes or thinks is a matter of fact; even if they did, it would not follow that the content of the speaker’s belief — the proposition that sea urchin is tasty — must be counterstance contingent in context as well (more on this momentarily). So, no familiarity inference is generated, precisely because in this particular context nothing forces the speaker’s commitment to the proposition that sea urchin is tasty to be counterstance contingent, and neither Affective Grounding nor Strong Opinionatedness must be satisfied.

Similar remarks apply to the examples in (62). None of the (broadly) modal operators involved require, as a semantic matter, their complement to be counterstance contingent, as shown by the fact that their complements can be non-discretionary:

- (64) I don’t know what kind of drink this is but ...
 - a. ... it might be fermented.
 - b. ... it must be fermented.
 - c. ... it’s certain to be fermented.
 - d. ... it’s clearly fermented.
 - e. ... it’s evidently fermented.

And Maximize Commitment, again, targets the asserted contents of (62a–e), requiring these contents to be interpreted as counterstance contingent insofar as context allows. But even if such an interpretation is available for some of the modal expressions involved,

their *complements* need not be counterstance contingent as well. For instance, it may very well be a matter of discretion whether some proposition passes as certain or not, since standards for certainty may differ. But this is so *regardless* of whether or not the truth of the complement *itself* is a matter of discretion. In particular, whether some proposition p is certain may be a matter of discretion even if p is a plain matter of fact.

- (65) a. I consider it certain that this drink is fermented.
 b. I consider it certain that there is liquid water under the ice of Europa.

So again, even insofar as some of the examples in (62) come with a speaker’s commitment to sea urchin being tasty, none of the principles we have invoked would require that commitment to be counterstance contingent in context. As a consequence, no familiarity inference is generated in the first place, and hence there is no need for an explanation of how or why the operators in (62a–e) block such an inference from projecting. It is in this sense that the framework developed here need not appeal to an obviation mechanism, in contrast to those offered in Pearson 2013; Anand and Korotkova 2018; Ninan 2020; Willer and Kennedy 2020.

It is worth stressing that nothing we have said here conflicts with the popular claim that a belief that sea urchin *must* be tasty commits a rational subject to the belief that sea urchin *is* tasty, and vice versa. This is perfectly consistent with our claim that the latter commitment may not always be expressible by asserting “*Sea urchin is tasty*,” since doing so is subject to Maximize Commitment, which cannot be satisfied when the commitment is not affectively grounded. Indeed, if we follow von Fintel and Gillies (2010) (see also Korotkova and Anand 2021 for relevant discussion) and take epistemic *must* to come with an *indirectness* requirement, “*Sea urchin must be tasty*” and “*Sea urchin is tasty*” end up with conflicting but complementary assertability conditions: non-affective grounding is required by the former (since affective grounding, in this case, entails direct experience) and affective grounding by the latter. The resulting picture would be that English provides (at least) two alternative means for articulating a taste judgment: a direct one that is available given suitable affective grounding, and an indirect one that requires non-affective grounding and that commits the speaker to p by asserting a commitment to *must p*.²³

Finally, exocentric uses of of predicates of personal taste such as (66) are another case in which Maximize Commitment lacks purchase:

- (66) My cat is happy because her new food is tasty.

²³One technical feature worth flagging here is that whenever some agent believes that sea urchin *must* be tasty ($\Box p$) without affective grounding, at least some of the agent’s counterstances will be committed to $\Box p$ without being committed to p . What this shows is that a commitment of $\Box p$ brings a commitment to p in its wake due to certain constraints that apply to ordinary doxastic agents but not to counterstances, which are first and foremost beholden to constraints on what it takes to be in a position to take a stance. It poses no trouble making this formally precise. Here is one path that stays squarely in classical territory (see Kratzer 2012 and references therein). Treat epistemic *must* (\Box_e) as a universal quantifier over the set of possible worlds compatible with what is known and most plausible given the normal course of events. Suppose that context supplies a function *Best* mapping worlds to those ways the world could that are optimal in these lights: $\llbracket \Box_e \phi \rrbracket^{c,w} = 1$ iff $\text{Best}(w) \subseteq \llbracket \phi \rrbracket^c$. Since things may not go as they normally go — and so in some cases $w \notin \text{Best}(w)$ — it poses no trouble modeling a counterstance as committed to $\Box p$ without being committed to p . But ordinary doxastic states can be further restricted — say they only allow “proper” points of evaluation w such that $w \in \text{Best}(w)$ — and so if committed to $\Box p$ *must* also be committed to p . Other analyses of would do the job just as well, including the dynamic analysis of *must* as a test over minimal spheres of possible worlds in an information carrier that has been proposed by Willer (2017).

(66) does not articulate the speaker’s tastes and sensibilities but “anchors” the application criteria of *tasty* to the cat’s taste. This intuition can be elaborated in a variety of ways, but one promising idea is that language sometimes allows us to articulate a perspective other than one’s own and that in such cases context shifts so that certain parameters are anchored to the attitudes of the individual whose perspective we are articulating. Epistemic *might*-claims, for instance, are by default evaluated against what is known by the speaker (DeRose 1991); but we can sometimes use *might* to articulate what is epistemically possible given someone else’s state of mind (Egan, Hawthorne, and Weatherson 2005), and in those cases it is the attributee’s knowledge that is contextually relevant, rather than the speaker’s. Likewise, we suggest that exocentric uses of predicates of personal taste shift the context so that it is someone else’s state of mind — in (66), the cat’s — that is articulated. Since the speaker is thus not expressing a commitment to the proposition that the food is tasty, Maximize Commitment has no purchase.²⁴

5 Loose Ends

This section explores several consequences and implications of our counterstance-based approach to subjective language. We begin with a brief discussion of two issues that are of general relevance to any theory of subjective meaning: how it is embedded in a broader compositional theory of meaning, and what it says about disagreement. We then look at a potential challenge for our account involving the interaction of our denotations for *find* and *consider* and the principle of Maximize Presupposition, and argue that, far from being a problem, this interaction provides the basis for an account of the interpretive differences between these verbs that we noted in Section 2.1. And finally, we conclude by examining additional data involving *consider*, which leads us to further articulate what it means to adopt a “pragmatic stance” and thereby render an attitude counterstance contingent.

5.1 Composition and projection

For reasons of space, we have set aside here two important questions about how subjective predicates interact with binary connectives and quantifiers. First, Sæbø (2009) uses patterns of acceptability in complex complements of subjective attitude verbs to argue that only a type-theoretic analysis — which we discussed and rejected in Section 2.1 — can accurately predict when perspectival content projects and when it does not. In Willer and Kennedy 2022, we address Sæbø’s argument and show how a counterstance-centered account may be extended to account for the embedding behavior of complex complements under subjective attitude verbs.

²⁴An assertion of (66) seems to express a commitment to the food being tasty *to the cat*, but clearly in order to know that one does not need to sample the food. Exocentric uses do carry their own distinct familiarity inference in that they suggest that the individual whose tastes and sensibilities matter has experience of the item under consideration. So for instance, while an utterance of (66) does not suggest that the speaker has sampled the food, it does suggest that the cat has. We cannot address the issue in detail here, but the basic suggestion is that due to the shiftiness of exocentric uses, a proposition p expressed using a predicate of personal taste in context c with experiential anchor a must be radically counterstance contingent with respect to $\kappa_c(a, p)$, and so familiarity inferences hold of the anchor, as they do for the subject of a subjective attitude ascription. See Willer and Kennedy (2020) for a more detailed discussion of such uses.

And second, Ninan (2020) points to the challenge of accounting for how familiarity inferences project from under the scope of quantifiers. In Willer and Kennedy 2022, we suggest that a quantified construction of the form ‘ $Q_x(\phi, \psi)$ ’ raises its scope ψ as an issue and that an issue is (radically) counterstance contingent just in case one of its resolutions (i.e. one cell in that issue) is; subjective attitudes (roughly) presuppose the counterstance contingency of the issue raised by the complement. While explicitly designed to respond to Sæbø’s argument, this setup may also serve as a foundation for responding to Ninan’s challenge. In outline: suppose that worlds are partial assignments of objects to the extension and anti-extension of subjective predicates; each counterstance contingent proposition in an issue then implies familiarity with some proportion of objects. Taking inspiration from Ninan’s 2020 supervaluationist proposal, we can then say that an attitude attribution implies familiarity with some particular object just in case each resolution compatible with that attitude implies that the attributee is familiar with that object; it implies familiarity with some quantity Q of objects just in case each resolution compatible with that attitude implies that the attributee is familiar with that quantity of objects. For instance, every complete answer to the question “What is tasty?” compatible with the belief that something on the table is tasty must put some object on the table in the extension of *tasty*, thus accounting for the intuition that “Something on the table is tasty” implies that the speaker has sampled some object or other on the table. Spelling out all the relevant technical details — which would include stating suitable semantic clauses for natural language quantifiers — is a task that we will leave to another day, but we submit there is no reason to be skeptical that the needed work can be carried out.²⁵

5.2 Faultless disagreement

We have also remained silent on the phenomenon of so-called “faultless disagreement.” This refers to the intuition that if Kim says that sea urchin is tasty and Lee responds that it is not tasty, they disagree and, moreover, neither of them need be “at fault” (see, for instance, Kölbel 2004; Lasersohn 2005, 2017; Glanzberg 2007; Stephenson 2007, 2008; Stojanovic 2007; Moltmann 2010; Sundell 2011; Barker 2013; Pearson 2013; Zakkou 2019). We have set this 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, Plunkett and Sundell 2013 and 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 and Khoo and Knobe’s (2018) account of moral disagreements for such a conception). Here we just point out that our concept of counterstance contingency is clearly relevant for the broader understanding 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.

²⁵Readers familiar with Ninan’s 2020 account will recall that it crucially relies on a distinction between failing to have sampled some item and having sampled but not liking it. This critical difference is captured in the story sketched here by treating worlds as partial and thus by distinguishing between failing to put some object x in the extension of β and putting x in the anti-extension of β . Suitable semantic entries for natural language quantifiers then allow us to leverage this distinction so that they effectively mirror Ninan’s attractive proposals.

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 verbs. This prediction is supported by observations in [Vardomsakaya 2018](#) and experimental results reported by [Khoo and Knobe \(2018\)](#) which show that whether a disagreement counts as faultless is crucially dependent on context and not tied to lexical items.

A comprehensive articulation of the implications of our proposals for faultless disagreement does not only require a precise account of what this phenomenon is supposed to be. Prior to that, we need a comprehensive story of the way that counterstances — and discourse alternatives more generally — interact with the norms that govern how we assert, reject, and evaluate utterances in everyday discourse. We must leave such a story for another day.

5.3 Maximize Presupposition

It is a consequence of the definitions of (radical) counterstance contingency in (40) that whenever a proposition is radically counterstance contingent, it is also counterstance contingent, but not vice versa. This means that the presuppositions of *find* are strictly stronger than those of *consider*. Given our assumption that the at-issue meanings of these verbs are the same, this raises the question of whether our initial observation that predicates like *tasty* can embed under both *find* and *consider* is a problem:

- (67) a. I find this sea urchin tasty.
 b. I consider this sea urchin tasty.

The worry is that it is generally the case that whenever two expressions have the same at-issue meaning but one has stronger presuppositions, use of the weaker term is disallowed in contexts in which the presuppositions of the stronger term are satisfied:

- (68) a. Both/#every one of the twins is happy.
 b. The/#a sun rose at 5.45 this morning.
 c. Kim knows/#believes that Paris is in France.

It is contrasts like these which led [Heim \(1991\)](#) to formulate a principle, now known as Maximize Presupposition (and discussed in Section 4.3), which forbids use of an expression when there is an alternative with the same at-issue content but stronger presuppositions whenever those presuppositions are satisfied (see also [Percus 2006](#); [Sauerland 2008](#); [Lauer 2016](#)). If our claims in Section 3.2 about essentially underdetermined predicates like *tasty* are correct, then every context is such that a presupposition of radical counterstance contingency for propositions involving such predicates can be satisfied. The question for us is why Maximize Presupposition does not thereby render use of *consider* with predicates like *tasty* unacceptable.

One answer could be that in fact *find* and *consider* are not identical in their at-issue content, as argued, for example, by [Muñoz \(2019\)](#) (but see note 9), and so are not alternatives in the sense that matters for Maximize Presupposition. Another could be that Maximize Presupposition does not extend to these kinds of cases, for whatever reason. (Perhaps contingency presuppositions are exempt.) However, we would like to

suggest that there is reason to believe that, in fact, Maximize Presupposition is as active in (67a–b) as it is in (68a–c) — which is to say that when the *consider* option is acceptable, the *find* variant is not.

The evidence for this comes from interpretive differences between *find* and *consider* that we already noted in Section 2.1. Recall first that pairs like (15a–b), repeated below, show that *find* and *consider* disambiguate between “qualitative” and “quantitative” interpretations of predicates like *long*: (15a) has only the former interpretation, and (15b) has only the latter.

- (15) a. Kim finds the 8-hour flight from Chicago to Frankfurt longer than the 13-hour flight from Chicago to Tokyo.
 b. Kim considers the 8-hour flight from Chicago to Frankfurt longer than the 13-hour flight from Chicago to Tokyo.

And second, we saw that pairs like (11a–b) differ in meaning, with the *find* sentence expressing a more “sensual” judgment and the *consider* sentence a more “intellectual” one; we hear the same difference in (67a–b).

- (11) a. I find you fascinating/irresistible/sexy.
 b. I consider you fascinating/irresistible/sexy.

Our explanation for these differences relies on the fact that whether a proposition is counterstance contingent, radically counterstance contingent, or neither is both a matter of lexical semantics and *also* a contextual affair. The contrast between the examples in (15a–b) is plausibly explained in the former terms, with qualitative vs. quantitative interpretations reflecting (possibly productive) lexical ambiguity or polysemy (see e.g. Bierwisch 1989; Kennedy and McNally 2010; Kennedy 2016; Willer and Kennedy 2020). What exactly this ambiguity consists in need not concern us here; what matters for our purposes — and what is intuitively correct — is that qualitative readings involve essential underdetermination and quantitative readings involve incidental underdetermination. If this is correct, and if incidentally underdetermined predicates are consistent with counterstance contingency but not radical counterstance contingency, then the contrast in (15a–b) follows: *consider* has only the quantitative reading because the qualitative reading, consistent with radical counterstance contingency, forces use of *find*.

The same type of explanation is in principle possible for the cases in (11a–b), but we find a polysemy/ambiguity hypotheses less plausible for these predicates.²⁶ Instead, we suggest that contrast between the examples in (11a) and (11b) reflects the context sensitivity of the counterstance selection function. Intuitively, use of *consider* suggests that the attitude holder’s judgment does not hinge (entirely) on her affective/experiential

²⁶One reason to think that the qualitative vs. quantitative uses of e.g. *long* reflects a lexical distinction is that they give rise to different familiarity inferences: (15a) requires Kim to have actually experienced the flights, but (15b) only requires that she know how much time they (and their related activities) take. In contrast, predicates like *tasty*, *fascinating*, *irresistible* and *sexy* appear to come with the same kinds of familiarity inferences regardless of whether they are associated with “sensual” judgments in the complement of *find* or “intellectual” judgments in the complement of *consider*, *pace McNally and Stojanovic (2017)*:

(i) # I find/consider sea urchin tasty, though I’ve never tried it.

state. But this is just what we expect if the judgment is counterstance contingent but not radically so, as predicted by Maximize Presupposition. If the judgment were a pure matter of affect, it would be radically counterstance contingent; use of *consider* signals that it is merely counterstance contingent, and so the judgment does not (entirely) hinge on affect. Instead, the speaker signals that there are *reasons* for holding the opinion under consideration, reasons that, we might add, the speaker takes to be such that they would cause anyone who is properly attuned to see things the same way. Although a counterstance space with this structure is never obligatory for an essentially underdetermined predicate, it is always in principle possible, given the context sensitivity of the counterstance selection function.

If this explanation is on the right track, then we predict acceptability of *consider* with *find*-embeddable predicates to track the possibility of accommodation of a context in which the attitude holder’s judgment does not hinge on her affective state. This prediction appears to be correct. Consider, for example, (69a-b).

- (69) a. I find that film hilarious.
- b. I consider that film hilarious.

In a context in which two friends are discussing a film they have just seen, it is considerably more natural for one of them to use (69a) than (69b) to express her judgment, since here their recent experiences are salient. But in a context in which two professional film critics are discussing how to describe the film in a review, (69b) is perfect.

There is, no doubt, more to say about these cases. But the difference in meaning between the examples in (11a) and (11b) is real, and we know of no other account of subjective attitude verbs that simultaneously captures the similarities between *find* and *consider* with respect to selection and familiarity inferences and also provides a principled basis for explaining these differences.

5.4 On taking a “pragmatic stance”

Earlier we said that counterstances arise because speakers are aware of the fact that what we say and believe is a practical affair, in the sense that one’s stance on an issue may depend on (perhaps unarticulated) practical decisions, intentions, and plans. The paradigmatic cases are, of course, those of straightforward semantic underdetermination: whether one wants to call Pluto a planet or treat six feet as tall seems to be a clear matter of linguistic choice. We will not defend at length here the claim that judgments of taste or beauty or morality also have a pragmatic element to them even though, as we have suggested, their application is ultimately dictated by how one is affected by a given object or event, as there seems to be a robust pre-theoretical sense in which immersion in contingent social practices influences one’s moral, aesthetic, and gustatory standards.

A thornier class of cases is presented by examples such as the following:

- (70) a. The ancient Greeks considered stars holes in the sky.
- b. The ancient Greeks believed stars to be holes in the sky.
- (71) a. Mathematicians consider Goldbach’s conjecture unprovable.
- b. Mathematicians believe Goldbach’s conjecture to be unprovable.

The use of *consider* in (70a) and (71a) seems perfectly felicitous, and yet the attitudes ascribed pertain to matters of fact in any intuitive sense of the word. How, then, could the presupposition of counterstance contingency possibly be satisfied here?

The felicity of cases such as (70a) and (71a) is compatible with our account, given a suitable conception of what constitutes a pragmatic stance, and assuming that speakers are flexible when it comes to identifying a set of facts in discourse that can be interpreted in opposing ways.²⁷ Indeed, it is natural to hear (70a) as signaling that there were some observational facts that the ancient Greeks interpreted in a distinct way, based on how they “saw the world,” including their views about astronomy and scientific inquiry more generally. *They* took the way the stars looked at night to be sufficient to conclude that they are holes in the sky; *we* can see the stars the same way and yet draw the opposing conclusion. Relatedly, (71a) suggests that mathematicians treat Goldbach’s conjecture as unprovable, and that this stance based on what they know about the theorem and previous attempts to prove it, but that these facts alone do not actually force the issue in one way or the other. Hence we have a “leap from the facts.”

To make sense of (70a) and (71a), it is crucial that the factual background on which counterstances must agree be suitably tailored. Specifically, to make sense of (70a), certain facts that we (think we) know today about the stars cannot be included among the matters of fact on which all counterstances must agree in context. And insofar as one wants to talk about undiscovered mathematical facts in the first place, there either exists a proof of Goldbach’s conjecture or there does not — so again some factual issues cannot be among those on which all counterstances must agree, if the attitude ascribed in (71a) is supposed to be counterstance contingent. Our model is flexible enough to accommodate these possibilities, but of course it is legitimate to wonder how the required counterstance spaces become available in context.

We do not pretend to have a comprehensive answer to this question, but one natural suggestion would be that even apparent matters of fact can become matters of discretion in context if there is no consensus on how the issue can be settled once and for all, or little confidence that evidence can be found that would satisfy all parties involved. This is why the use of *consider* is natural when describing attitudes that are entrenched in foreign scientific paradigms or other kind of “worldviews” that stand in the way toward finding some common ground, as in the case involving the ancient Greeks, or as in (72a) and (72b):

- (72) a. Joseph Priestley considered oxygen dephlogisticated air.
- b. Many Republicans consider the 2020 US presidential election stolen.

In contrast, the *consider* is less acceptable in the following cases, at least in a typical context in which there is taken to be some common agreement on how the matters can be resolved:

- (73) a. # Ruby considers the red envelope to contain \$1000.

²⁷In contrast, it is much less clear that such cases are compatible with Coppock’s (2018) outlook-based account, at least not without adopting the conceptual foundations of counterstance theory. Possible worlds are complete statements of what is a matter of fact, and since it is arguably a fact whether or not the stars are holes in the sky, and whether or not Goldbach’s conjecture is provable, we do not expect the truth-value of the complements of *consider* in (70a) and (71a) to vary across outlooks corresponding to a single possible world. These propositions are therefore non-discretionary, and should fail to embed under *consider*.

- b. # Some scientists consider there to be life on Mars.

The use of *consider* improves, however, if context supports the belief that Ruby would remain convinced that there are \$1000 in the envelope even if she were to open it and find it empty; or the belief that the the scientists in question have some peculiar conception of what constitutes “life.” So again whether some issue counts as a matter of fact in context — and thus is a potentially counterstance contingent affair — seems to depend on the presence of a shared background methodology for conclusively resolving the issue.²⁸

It does not count against this proposal that the ancient Greeks might not have thought about their astronomical views as matters of discretion: as we have already stressed, it is the speaker, not the attributee, who presupposes the counterstance contingency of the attitude ascribed. Nor are we suggesting that all counterstances are created equal: the speaker may very well think that some stances are intrinsically better than others. The use of a subjective attitude verb in attributing an opinion is thus not necessarily to suggest that the issue is a “wash:” rather, it is to highlight that the relevant facts on the ground can be interpreted in opposing ways due to differing, broadly pragmatic outlooks. What is true is that the variety of pragmatic stances that underwrite the phenomenon of counterstance contingency is quite broad and may include complex background principles that, inter alia, implicitly guide our opinion-forming practices through commonsense assumptions and a conception of what counts as exemplary inquiry. They are not exclusively manners of speaking. Still, a broad sense counterstance contingency as grounded in pragmatic differences may be preserved, given the mundane fact that norms governing belief formation and retention are themselves relative to frames of reference.²⁹

6 Final thoughts

We have argued that an analysis of subjectivity as a distinguished kind of contingency — counterstance contingency — can be leveraged into a unified account of the selectional restrictions of subjective attitude verbs as well of the distinct familiarity inferences that subjective attitude attributions and subjective predications bring into play. It must, of course, be admitted that counterstance contingency is a complex phenomenon and in particular that our account of how counterstances become available in context requires quite a bit of flexibility and is very sensitive to the particulars of the case. Such twists and turns may just be what is commanded by the empirical complexity of subjectivity in natural languages. An alternative suggestion would be that *find* and *consider* track a notion of subjectivity that resists an informative analysis after all — or quite simply that we have to do better. The fact would remain that the subjectivity at play here is formally well captured as a distinct kind of contingency that manifests in the possibility of

²⁸Would a proof of Goldbach’s conjecture not be enough to change a mathematician’s mind, just as finding no money in an envelope would convince an ordinary person that there is no money in it? Presumably yes, but what counts as a correct and complete proof is not a straightforward affair, at least when it comes to complex mathematical theorems. As a matter fact, a considerable number of proofs in mathematics remain controversial today, Mochizuki’s alleged proof of the abc conjecture being a prominent example.

²⁹What we say here is of course inspired by Kuhn’s (1962) discussion of paradigms but the suggestion that our opinion-forming practices are inherently shaped by socially constructed frames of reference traces back at least to the writings of the American pragmatists: see in particular Lewis 1923 but also for instance James 1890, Chapter XXVII, James 1907, Lecture V, and Peirce 1905.

opposing stances. Moreover, we have identified at least some formal constraints on what it takes for an attitude to be thus contingent — Strong Opinionatedness and Affective Grounding — that account for the variety of familiarity inferences that subjective attitude ascriptions (and subjective assertions) trigger. It may be that we do not fully understand (yet) why some opinions pass as subjective in context while others do not, but, we submit, we can illuminate the formal structure that subjective attitude ascriptions seem to track. If someone else is able to provide a firmer conceptual grounding for our proposals, and provide a better story of how counterstances become (un)available in discourse, more power to them. We submit that this is a task worth pursuing: if the proposal developed here is at least somewhat on the right track, the notion of a counterstance is bound to play an important role in our best theory of linguistic meaning.

References

- Anand, Pranav, and Natasha Korotkova. 2018. “Acquaintance Inference and Obviation.” In *Proceedings of Sinn und Bedeutung 22*, ed. Uli Sauerland and Stephanie Solt, vol. 1 of *ZASPiL 60*, 55–72. Berlin: Leibniz-Centre General Linguistics.
- Bach, Kent, and Robert M. Harnish. 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.
- Barker, Chris. 2013. “Negotiating Taste.” *Inquiry* 56(2-3): 240–257. <http://dx.doi.org/10.1080/0020174X.2013.784482>.
- Bierwisch, Manfred. 1989. “The Semantics of Gradation.” In *Dimensional Adjectives*, ed. Manfred Bierwisch and Ewald Lang, 71–261. Berlin: Springer.
- Bouchard, David-Étienne. 2012. *Long-Distance Degree Quantification and the Grammar of Subjectivity*. Ph.D. thesis, McGill University, Montreal.
- Braun, David. 1998. “Understanding Belief Reports.” *Philosophical Review* 107(4): 555–595. <http://dx.doi.org/10.2307/2998375>.
- Budd, Malcolm. 2003. “The Acquaintance Principle.” *British Journal of Aesthetics* 43(3): 386–392. <http://dx.doi.org/10.1093/bjaesthetics/43.4.386>.
- Bylinina, Lisa. 2017. “Judge-dependence in Degree Constructions.” *Journal of Semantics* 34(2): 291–331. <http://dx.doi.org/10.1093/jos/ffw011>.
- Byrne, Alex. 2001. “Intentionalism Defended.” *Philosophical Review* 110(2): 199–240. <http://dx.doi.org/10.1215/00318108-110-2-199>.
- Callahan, Laura Frances. 2018. “Moral Testimony: A Re-Conceived Understanding Explanation.” *Philosophical Quarterly* 68(272): 437–459. <http://dx.doi.org/10.1093/pq/pqx057>.
- Chalmers, David J. 1996. *The Conscious Mind*. New York: Oxford University Press.
- Condoravdi, Cleo. 2002. “Temporal Interpretation of Modals: Modals for the Present and for the Past.” In *The Construction of Meaning*, ed. David I. Beaver, Stefan Kaufmann, Brady Clark, and Luis Casillas, 59–88. Stanford, CA: CSLI Publishers.

- Coppock, Elizabeth. 2018. “Outlook-based Semantics.” *Linguistics and Philosophy* 41(2): 125–164. <http://dx.doi.org/10.1007/s10988-017-9222-y>.
- DeRose, Keith. 1991. “Epistemic Possibilities.” *Philosophical Review* 100(4): 581–605. <http://dx.doi.org/10.2307/2185175>.
- Dretske, Fred. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. “Epistemic Modals in Context.” In *Contextualism in Philosophy*, ed. Gerhard Preyer and Georg Peter, 131–169. New York: Oxford University Press.
- Enoch, David. 2014. “A Defense of Moral Deference.” *Journal of Philosophy* 111(5): 229–258. <http://dx.doi.org/10.5840/jphil2014111520>.
- Evans, Gareth. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.
- von Fintel, Kai, and Anthony S. Gillies. 2010. “Must ...Stay ...Strong.” *Natural Language Semantics* 18(4): 351–383. <http://dx.doi.org/10.1007/s11050-010-9058-2>.
- Fleisher, Nicholas. 2013. “The Dynamics of Subjectivity.” In *Proceedings of Semantics and Linguistic Theory 23*, ed. Todd Snider, 276–294. Ithaca, NY: CLC.
- Franzén, Nils. 2018. “Aesthetic Evaluation and First-Hand Experience.” *Australasian Journal of Philosophy* 96(4): 669–682. <http://dx.doi.org/10.1080/00048402.2018.1425729>.
- . 2020. “Evaluative Discourse and Affective States of Mind.” *Mind* 129(516): 1095–1126. <http://dx.doi.org/10.1093/mind/fzz088>.
- Giannakidou, Anastasia, and Alda Mari. 2015. “Mixed Non-Veridicality and Mood Choice with Emotives.” In *Proceedings of the Fifty-first Annual Meeting of the Chicago Linguistic Society*, 181–195. Chicago, IL: Chicago Linguistic Society.
- . 2017. “A Unified Analysis of the Future as Epistemic Modality.” *Natural Language Semantics* 36(1): 85–129. <http://dx.doi.org/10.1007/s11049-017-9366-z>.
- Glanzberg, Michael. 2007. “Context, Content, and Relativism.” *Philosophical Studies* 136(1): 1–29. <http://dx.doi.org/10.1007/s11098-007-9145-5>.
- Goodman, Noah D., and Michael C. Frank. 2016. “Pragmatic Language Interpretation as Probabilistic Inference.” *Trends in Cognitive Science* 20(11): 818–829. <http://dx.doi.org/10.1016/j.tics.2016.08.005>.
- Grice, H. Paul. 1975. “Logic and Conversation.” In *Syntax and Semantics, Vol. 3, Speech Acts*, ed. Peter Cole and Jerry L. Morgan, 41–58. New York: Academic Press.
- Grinsell, Timothy. 2017. *Semantic Indecision*. Ph.D. thesis, University of Chicago.
- Heim, Irene. 1991. “Articles and Definiteness.” In *Semantics. An International Handbook of Contemporary Research*, ed. Armin von Stechow and Dieter Wunderlich, 487–535. Berlin: Walter de Gruyter & Co.

- Hills, Alison. 2009. "Moral Testimony and Moral Epistemology." *Ethics* 1(94): 445–454. <http://dx.doi.org/10.1086/648610>.
- . Forthcoming. "Aesthetic Testimony, Understanding and Virtue." *Noûs*. <http://dx.doi.org/10.1111/nous.12344>.
- Hirvonen, Sanna. 2014. *Predicates of Personal Taste and Perspective Dependence*. Ph.D. thesis, University College London.
- Hopkins, Robert. 2007. "What is Wrong With Moral Testimony?" *Philosophy and Phenomenological Research* 74(3): 611–634. <http://dx.doi.org/10.1111/j.1933-1592.2007.00042.x>.
- . 2011. "How to Be a Pessimist about Aesthetic Testimony." *Journal of Philosophy* 108(3): 138–157. <http://dx.doi.org/10.5840/jphil201110838>.
- James, William. 1890. *The Principles of Psychology*, vol. 2. New York: Henry Holt and Company.
- . 1907. *Pragmatism: A New Name for Some Old Ways of Thinking*. New York: Longmans, Green, and Co.
- Kamm, Frances Myrna. 2007. *Intricate Ethics: Rights, Responsibilities and Permissible Harm*. Oxford: Oxford University Press.
- Kant, Immanuel. 1781. *Critik der Reinen Vernunft*. Riga: Johann Friedrich Hartknoch.
- Kaplan, David. 1989. "Demonstratives." In *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein, 481–563. New York: Oxford University Press.
- Kennedy, Christopher. 2013. "Two Sources of Subjectivity: Qualitative Assessment and Dimensional Uncertainty." *Inquiry* 56(2-3): 258–277. <http://dx.doi.org/10.1080/0020174X.2013.784483>.
- . 2016. "Two Kinds of Subjective Meaning." In *Subjective Meaning: Alternatives to Relativism*, ed. Cécile Meier and Janneke van Wijnenbergen-Huitink, 105–126. Berlin: De Gruyter. <http://dx.doi.org/10.1515/9783110402001-007>.
- Kennedy, Christopher, and Louise McNally. 2010. "Color, Context and Compositionality." *Synthese* 174(1): 79–98. <http://dx.doi.org/10.1007/s11229-009-9685-7>.
- Kennedy, Christopher, and Malte Willer. 2016. "Subjective Attitudes and Counterstance Contingency." In *Proceedings of Semantics and Linguistic Theory 26*, ed. Mary Moroney, Carol-Rose Little, Jacob Collard, and Dan Burgdorf, 913–933.
- Khoo, Justin, and Joshua Knobe. 2018. "Moral Disagreement and Moral Semantics." *Noûs* 52(1): 109–143. <http://dx.doi.org/10.1111/nous.12151>.
- Kölbel, Max. 2004. "Faultless Disagreement." *Proceedings of the Aristotelian Society* 104(1): 53–73. <http://dx.doi.org/10.1111/j.0066-7373.2004.00081.x>.
- Korotkova, Natasha, and Pranav Anand. 2021. "'Find', 'Must' and Conflicting Evidence." In *Proceedings of Sinn und Bedeutung 25*, ed. Patrick Georg Grosz, Luisa Martí, Hazel Pearson, Yasutada Sudo, and Sarah Zobel, 515–532.

- Kratzer, Angelika. 2012. *Modals and Conditionals*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199234684.001.0001>.
- Kuhn, Thomas Samuel. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Laserson, Peter. 2005. "Context Dependence, Disagreement, and Predicates of Personal Taste." *Linguistics and Philosophy* 28(6): 643–686. <http://dx.doi.org/10.1007/s10988-005-0596-x>.
- . 2017. *Subjectivity and Perspective in Truth-Theoretic Semantics*. Oxford: Oxford University Press.
- Lassiter, Daniel, and Noah Goodman. 2014. "Context, Scale Structure, and Statistics in the Interpretation of Positive-form Adjectives." In *Proceedings of Semantics and Linguistic Theory 24*, ed. Todd Snider, Sarah D'Antonio, and Mia Weigand, 587–610.
- Lassiter, Daniel, and Noah D Goodman. 2015. "Adjectival Vagueness in a Bayesian model of Interpretation." *Synthese* 194(10): 3801–3836. <http://dx.doi.org/10.1007/s11229-015-0786-1>.
- Lauer, Sven. 2016. "On the Status of 'Maximize Presupposition'." In *Proceedings of Semantics and Linguistic Theory 26*, ed. Mary Moroney, Carol-Rose Little, Jacob Collard, and Dan Burgdorf, 980–1001.
- Lewis, Clarence Irving. 1923. "A Pragmatic Conception of the *A Priori*." *Journal of Philosophy* 20(7): 166–177. <http://dx.doi.org/110.2307/2939833>.
- Lewis, Karen S. 2020. "Metasemantics without semantic intentions." *Inquiry* 0(0): 1–29. <http://dx.doi.org/10.1080/0020174X.2020.1847184>.
- Livingston, Paisley. 2003. "On an Apparent Truism in Aesthetics." *British Journal of Aesthetics* 43(4): 260–278–392. <http://dx.doi.org/10.1093/bjaesthetics/43.3.260>.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199682751.001.0001>.
- . 2016. "Vagueness as Indecision." *Aristotelian Society Supplementary Volume* 90(1): 255–283. <http://dx.doi.org/doi.org/10.1093/arisup/akw013>.
- . 2020. "Indeterminacy as Indecision Lecture III: Indeterminacy as Indecision." *Journal of Philosophy* 117(11/12): 643–667. <http://dx.doi.org/10.5840/jphil202011711/1242>.
- McDowell, John. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- McNally, Louise, and Isidora Stojanovic. 2017. "Aesthetic Adjectives." In *The Semantics of Aesthetic Judgment*, ed. James Young, 17–37. Oxford University Press.

- Meskin, Aaron. 2004. "Aesthetic Testimony: What Can We Learn from Others about Beauty and Art?" *Philosophy and Phenomenological Research* 69(1): 65–91. <http://dx.doi.org/10.1111/j.1933-1592.2004.tb00384.x>.
- Moltmann, Friederike. 2010. "Relative Truth and the First Person." *Philosophical Studies* 150(2): 187–220. <http://dx.doi.org/10.1007/s11098-009-9383-9>.
- Muñoz, Patrick. 2019. *On Tongues: The Grammar of Experiential Evaluation*. Ph.D. thesis, University of Chicago.
- . 2021. "The Acquaintance Inference as an Evidential Effect." Ms., University of Chicago.
- Ninan, Dilip. 2014. "Taste Predicates and the Acquaintance Inference." In *Proceedings of Semantics and Linguistic Theory 24*, ed. Lucas Champollion and Anna Szabolcsi, 290–304.
- . 2020. "The Projection Problem for Predicates of Taste." In *Proceedings of Semantics and Linguistic Theory 30*, ed. Joseph Rhyne, Kaelyn Lamp, Nicole Dreier, and Chloe Kwon, 753–778.
- Peacocke, Christopher. 1983. *Sense and Content*. Oxford: Oxford University Press.
- Pearson, Hazel. 2013. "A Judge-Free Semantics for Predicates of Personal Taste." *Journal of Semantics* 30(1): 103–154. <http://dx.doi.org/10.1093/jos/ffs001>.
- Pedersen, Walter. 2012. "Faultless Disagreement and Semantic Relativity." Unpublished ms., McGill University.
- Peirce, Charles Sanders. 1905. "Issues of Pragmaticism." *The Monist* 15(4): 481–499. <http://dx.doi.org/10.5840/monist19051544>.
- Percus, Orin. 2006. "Antipresuppositions." In *Theoretical and Empirical Studies of Reference and Anaphora: Toward the Establishment of Generative Grammar as an Empirical Science, Report of the Grant-Aid for Scientific Research (B), Project No. 15320052*, ed. Ayumi Ueyama, 52–73. Japan Society for the Promotion of Science.
- Plunkett, David, and Timothy Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13(23): 1–37. <http://dx.doi.org/2027/spo.3521354.0013.023>.
- Qing, Ciyang. 2020. *Semantic Underspecification and Its Contextual Resolution In the Domain of Degrees*. Ph.D. thesis, Stanford University.
- Qing, Ciyang, and Michael Franke. 2014. "Gradable Adjectives, Vagueness, and Optimal Language Use: A Speaker-oriented Model." In *Proceedings of Semantics and Linguistic Theory 24*, ed. Lucas Champollion and Anna Szabolcsi, 23–41.
- Reis, Marga. 2013. "Dt. finden und „subjektive Bedeutung“." *Linguistische Berichte* 2013(236): 389–426.
- Sæbø, Kjell Johan. 2009. "Judgment Ascriptions." *Linguistics and Philosophy* 32(4): 327–352. <http://dx.doi.org/10.1007/s10988-009-9063-4>.

- Sauerland, Uli. 2008. “Implicated Presuppositions.” In *The Discourse Potential of Underspecified Structures*, ed. Anita Steube, 581–600. Berlin: De Gruyter. <http://dx.doi.org/doi:10.1515/9783110209303.4.581>.
- Schmidt, Lauren, Noah Goodman, David Barner, and Joshua Tenenbaum. 2009. “How Tall Is *Tall*? Compositionality, Statistics and Gradable Adjectives.” In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 31, 3151–3156. Red Hook, NY: Curran Associates.
- Schroeder, Mark A. 2010. *Noncognitivism in Ethics*. New York: Routledge.
- Stephenson, Tamina. 2007. *Towards a Theory of Subjective Meaning*. Ph.D. thesis, MIT, Cambridge, MA.
- . 2008. “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste.” *Linguistics and Philosophy* 30(4): 487–525. <http://dx.doi.org/10.1007/s10988-008-9023-4>.
- Stojanovic, Isidora. 2007. “Talking about Taste: Disagreement, Implicit Arguments, and Relative Truth.” *Linguistics and Philosophy* 30(6): 691–706. <http://dx.doi.org/10.1007/s10988-008-9030-5>.
- Sundell, Timothy. 2011. “Disagreements about Taste.” *Philosophical Studies* 155(2): 276–288. <http://dx.doi.org/10.1007/s11098-010-9572-6>.
- Umbach, Carla. 2016. “Evaluative Propositions and Subjective Judgments.” In *Subjective Meaning: Alternatives to Relativism*, ed. Cécile Meier and Janneke van Wijnbergen-Huitink, 127–168. Berlin: De Gruyter. <http://dx.doi.org/10.1515/9783110402001-008>.
- Vardomsakaya, Tamara. 2018. *Sources of Subjectivity*. Ph.D. thesis, University of Chicago.
- Willer, Malte. 2017. “Lessons from Sobel Sequences.” *Semantics and Pragmatics* 10(4): 1–57. <http://dx.doi.org/10.3765/sp.10.4>.
- Willer, Malte, and Christopher Kennedy. 2020. “Assertion, expression, experience.” *Inquiry* 0(0): 1–37. <http://dx.doi.org/10.1080/0020174X.2020.1850338>.
- . 2022. “Perspectival Content and Semantic Composition.” In *Perspectives on Taste*, ed. Jeremy Wyatt, Julia Zakkou, and Dan Zeman, 207–230. New York: Routledge.
- Williamson, Timothy. 1996. “Knowing and Asserting.” *Philosophical Review* 105(4): 489–523. <http://dx.doi.org/10.2307/2998423>.
- . 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/019925656X.001.0001>.
- Wollheim, Richard. 1980. *Art and Its Objects*. 2nd ed. Cambridge: Cambridge University Press.

- Xiang, Ming, Christopher Kennedy, Weijie Xu, and Timothy Leffel. forthcoming. “Pragmatic Reasoning and Semantic Convention: A Case Study in Gradable Adjectives.” *Semantics and Pragmatics*.
- Xiang, Ming, Alex Kramer, and Christopher Kennedy. 2021. “Semantic Adaptation in Gradable Adjective Interpretation.” Unpublished ms., University of Chicago.
- Zakkou, Julia. 2019. “Denial and Retraction: A Challenge for Theories of Taste Predicates.” *Synthese* 196(4): 1555–1573. <http://dx.doi.org/10.1007/s11229-017-1520-y>.