

Disagreement, Contextualism, and Coordination

Xinyi Yu

UCL

A thesis presented for the degree of Master of Philosophy

I, Xinyi Yu, confirm that the work presented in this thesis is my own.

Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

The present work is concerned with disagreement over predicates of personal taste (PPTs), such as judgments that a symphony is beautiful or a dish is tasty. These disputes, though grounded in subjective standards, often nonetheless present themselves as genuine disagreements: speakers treat their judgments as incompatible and seek resolution, even while acknowledging their first-personal character. An adequate account must therefore explain both the subjective basis of taste and the interpersonal dynamics that make these exchanges persist as disagreements.

Plunkett and Sundell have proposed a metalinguistic negotiation framework, on which such disputes are understood as moments when speakers press competing standards for how a term like “beautiful” should be applied. I develop this proposal by integrating it into a Stalnakerian conception of context, which distinguishes between the private standards fixed by private Kaplanian contexts and the public Stalnakerian context shared by interlocutors. On this view, disagreement is not simply a matter of expressing opposing standards, but of attempting to influence what is taken up into the common ground—whether one’s own standard or a revised, mutually scrutinized alternative.

The resulting framework shows how contextualist semantics can preserve the first-personal authority of taste judgments while also explaining why disputes over beauty invite argument and revision. In this way, the account provides a plausible picture of how aesthetic discourse can support both individual perspectives and the possibility of genuine disagreement, without committing to an objectivist metaphysics of PPTs.

Impact Statement

The present work examines disagreement over predicates of personal taste, drawing on Plunkett and Sundell's metalinguistic negotiation framework and situating it within a Stalnakerian conception of context. The aim is to show how introducing a public common ground explains why such disputes involve pressure toward coordination rather than mere exchanges of private verdicts.

It is hoped that this account provides a new perspective on debates between objectivist, contextualist, and relativist approaches, while clarifying how aesthetic discourse can preserve individual perspectives yet still sustain practices of collective reasoning.

Acknowledgements

I would like to thank Daniel Rothschild for initial guidance on formulating the project. I would also like to thank Sam Carter for his supervision and support throughout the later stages.

Contents

1. Introduction.....	6
1.1 Objectivism.....	7
1.2 Faultless Disagreement.....	9
1.3 Objectivism Strikes Back—Humean Absolutism.....	12
1.4 The Beauty-Pleasure Principle.....	15
1.4.1 BPP And Humean Absolutism.....	17
1.5 “Beautiful” and the Predicates of Personal Taste.....	19
2. Non-objectivist Theories.....	22
2.1 Contextualism.....	23
2.2 Relativism.....	28
2.2.1 Disagreement Revisited.....	31
3. Assessment of MacFarlane’s Relativism.....	37
3.1 Retraction.....	37
3.2 Why Disagree?.....	43
3.3 Coordination of contexts.....	47
4. Common Ground.....	52
4.1 The Stalnakerian Context.....	52
4.2 Presupposition of Commonality Approach.....	56
4.3 Metalinguistic Negotiation.....	60
4.4 A Stalnakerian Twist.....	63
4.5 Five Desiderata.....	68

Chapter 1

Introduction

You are playing Bruckner's 8th Symphony on your turntable, eager to share with your friends what you consider to be one of the greatest symphonies ever written. As the music soars through the room, you lean back, close your eyes, and murmur, "It is beautiful," savoring the sweeping strings. Across from you, your friend Paul—whose musical judgment you normally trust—presses his lips together, raises an eyebrow, then frowns. When the final movement ends, he finally speaks: "It is not beautiful. It's just loud."

Clearly there is a disagreement—but what, exactly, are you disagreeing about? For the past two decades, philosophers of language have debated the nature of such disputes over predicates of personal taste. The puzzle is this: few are willing to say that the disagreement concerns whether the symphony objectively instantiates the property of beauty. Yet in the absence of such an objective property, it becomes unclear how these disputes constitute genuine disagreements at all. If each of your utterances reflects only your own standards of taste, then your statement only reflects the-beauty-of-Bruckner's-symphony-according-to-your-standard, while Paul's reflects its non-beauty-according-to-his. Without some shared standard, it looks as if you and Paul are not disagreeing but simply talking past one another.

In the shift from genuine disagreement to talking past each other, what slips away is a shared standard against which both judgments can be set. If each of you is only reporting how the music strikes you, there is no clash between your claims—just two parallel reactions. But when you disagree about whether the symphony is beautiful, you take yourselves to be in conflict, not merely stating private impressions. Without a shared basis for the judgment, that sense of genuine contestation slips away, and the exchange begins to look less like disagreement and more like two people talking past one another.

This thesis takes up that challenge. Its central aim is to develop a non-objectivist, specifically contextualist, account that preserves the possibility of *genuine* disagreement in disagreement over PPTs, even when our assertions do not track a universally shared standard.

The first chapter will begin by examining the most straightforward way of securing genuine disagreement: objectivism. On this view, the two parties genuinely disagree because their assertions express contradictory propositions—one claims that the symphony initiates the objective property of beauty, the other denies it. However, I will argue that although objectivism could accommodate genuine disagreements, it fails to capture the subjective nature of the judgment of taste. This critique provides motivation for exploring a

non-objectivist framework that can preserve both genuine disagreement and the personal dimension of aesthetic judgment.

1.1 Objectivism

Let us begin by clarifying what disagreement amounts to. While there is controversy in the literature on disagreement over predicates of personal taste (hereafter PPTs), a widely accepted baseline definition is as follows:

Genuine disagreement (GD): *If in a discourse, Speaker A and Speaker B really disagree with one another, then there is some proposition p such that some utterance by Speaker A expresses p and some utterance by Speaker B expresses q , where q entails not- p .*¹

This principle captures a central intuition: that disagreement involves a conflict between the contents of what is asserted—specifically, that the propositions expressed by the two speakers cannot both be true. Consider, for instance, the exchange in our leading case (hereafter *Symphony*):

You: “This symphony is beautiful.”

Paul: “This symphony is not beautiful.”

Paul’s response appears to directly negate your assertion. On a natural interpretation of the semantic content of each utterance, you attribute an aesthetic property to the symphony, while Paul denies that it possesses that property. Their respective contents are thus logically incompatible, which thereby satisfies the condition stipulated by GD: Paul’s assertion expresses a proposition that entails the negation of yours.

The plausibility of GD is further supported when we consider cases where it correctly rules out disagreement. Take the following dialogue:

Alphie: “I’m in Madison.”

Betty: “I’m not in Madison.”

Suppose Alphie says this while located in Madison, and Betty replies over the phone from Ann Arbor. Although Betty’s utterance is syntactically the negation of Alphie’s, we do not take them to be disagreeing in any substantive sense. This is because the pronoun “I” refers to different individuals in each utterance, and thus the propositions expressed are consistent. GD accounts for this intuitively: since the propositions are not logically inconsistent, there is no genuine disagreement.

Thus, GD not only captures clear cases of disagreement, such as disputes over aesthetic judgments, but also helps distinguish them from merely apparent ones. This makes it a plausible starting point for analyzing disputes in discourse involving PPTs.

¹ This formulation is adapted from Sundell (2010, p. 269).

Given this, objectivism enjoys a *prima facie* advantage: it straightforwardly accommodates GD by positing that aesthetic predicates express objective properties that can be affirmed or denied across contexts. I now turn to a closer examination of its plausibility.

Following MacFarlane (2014), I take objectivism to be the view that:

- (a) “beautiful” is true of some things, false of others, and
- (b) whether “beautiful” is true or false of something on a given occasion does not depend on the idiosyncratic standard of taste of the speaker (or any other salient individual).²

Objectivists thus hold that “beautiful” functions similarly to other predicates that describe the world, such as “green” or “hexagonal,” which characterize objects independently of their relation to the speaker. The disagreement in *Symphony* is therefore analogous to a disagreement over whether something is green or hexagonal—that is, it concerns whether Bruckner’s 8th Symphony initiates the objective property of beauty. More specifically, an objectivist would say that in *Symphony*, what you say is the standard-of-taste-neutral proposition “Bruckner’s 8th Symphony is beautiful,” and what Paul says is the negation of this proposition, and your disagreement is over which proposition is true.³

Before evaluating objectivism as an approach to aesthetic disagreement, I shall clarify its commitments. First, objectivism about beauty is compatible with the idea that beauty is defined in relation to humans. An objectivist may claim that something is beautiful if it has features that are pleasing to the aesthetic sensibilities of a normal human under normal conditions of perception. In this case, while beauty may hold little relevance for non-humans, it would still constitute an objective property. This can be illustrated by an analogy to adjectives like “green”: many philosophers maintain that greenness remains an objective property, even though it is understood in terms of human visual perception.

Furthermore, objectivism about beauty could also accommodate a degree of contextual sensitivity while maintaining that beauty is an objective property. Much like adjectives such as “green,” “tall,” or “fast,” which vary in application depending on context, objectivists can allow the threshold for what counts as “beautiful” to shift across different situations. For example, an objectivist might claim that “fast” is an objective predicate grounded in measurable speed, even though its application varies depending on the comparison class: a runner may be called “fast” at a school track meet but not when compared to Olympic athletes. Likewise, an objectivist can maintain that beauty is objectively grounded, even while acknowledging that the threshold of what qualifies as “beautiful” may shift depending on the context of evaluation. A kindergarten child’s

² Note that on this characterization, non-objectivism is compatible with there being an objectively correct standard of beauty, as long as our use of “beautiful” does not always track it.

³ This is adapted from Schafer (2011). Schafer’s characterisation is originally of ‘absolutism’ instead of ‘objectivism’, yet this is only a notational difference (at least for the purpose of this work).

painting might be deemed beautiful relative to other works by kindergarteners, but not when compared to art works in the Louvre.

As MacFarlane notes, objectivists can accept this kind of context-sensitivity in application so long as the underlying comparative relation—such as “[more beautiful] than”—remains independent of individual taste.⁴ Just as objectivists about fastness accept different contextual thresholds for “fast” while appealing to a stable underlying property (speed), objectivists about beauty can claim there exists an objective aesthetic scale that grounds judgments of beauty, even if the contextual thresholds for applying the term may vary. Given this characterization, objectivism, crucially, implies that there is a solid, objective ranking about the degree of beauty, irrespective of each person’s individual taste.

Importantly, while objectivism allows for context-sensitivity—acknowledging that what counts as beautiful may vary by situation—it does not allow the threshold for beauty to vary between individuals or groups based on subjective standards. An objectivist can say that a child’s painting is beautiful in a kindergarten context but not in the Louvre, yet this is because the context itself shifts the objective standard, not because beauty is relative to personal taste.

In the *Symphony* case, this distinction matters. An objectivist would say you and Paul are not simply applying different personal thresholds for beauty. On the objectivist view, there is a single, contextually determined but absolute standard for what counts as beautiful. Their disagreement is therefore genuine: they differ on whether Bruckner’s 8th meets that objective threshold.

1.2 Faultless Disagreement

Given this, we can now assess the plausibility of objectivism. Does the disagreement in *Symphony* resemble a disagreement over whether something is “green”? Many philosophers in the debate over disagreements of taste reject this idea, and some of them appeal to the phenomenon of *faultless disagreement* as a key reason for their skepticism. Although there is considerable debate over how to define “disagreement,” I will set those controversies aside for now. This is because objectivism is supposed to conflict with nearly all standard characterizations of faultless disagreement. I will return to these definitional issues—exploring different understandings of faultlessness later in this section, and various accounts of disagreement in the next chapter. For now, it is helpful to begin with an intuitive conception of “faultless disagreement”. Following Max Kölbel, a leading proponent of the notion, we can define faultless disagreement as a situation involving two thinkers, A and B,

⁴ 2014, p. 2. MacFarlane’s original claim is on the ‘tasty’. I will explain later why I think it could equally apply to ‘beautiful’.

and a proposition p , such that: (a) A believes (or judges) that p , while B believes (or judges) that *not*- p , and (b) neither A nor B has made a mistake (or is at fault).⁵

This phenomenon seems to be at odds with objectivism. If, as objectivists claim, your statement in *Symphony*—“Bruckner’s 8th Symphony is beautiful”—expresses a standard-of-taste-neutral proposition, then this proposition would simply be true or false. In that case, you and Paul are disagreeing over a factual matter: whether the symphony truly initiates the property of beauty (given the relevant threshold). Regardless of how much you like it or how much Paul dislikes it, one of you must be flat-out mistaken about its real beauty.

However, does faultless disagreement truly exist in aesthetic disputes? Some—such as a passionate admirer or a staunch critic of Bruckner⁶—may not view *Symphony* as an instance of faultless disagreement. Nevertheless, I will not dwell on the debate over its existence here, as I will later argue that faultless disagreement does not necessarily pose a clear disadvantage for objectivism as a solution to the problem of disagreement. That said, I do believe there are cases of aesthetic disagreement that strongly support the intuition that faultless disagreement is possible. Consider *Symphony**:

Amy: Bruckner’s 8th Symphony is more beautiful than his 7th Symphony.

Alfred: No, it’s not the case that Bruckner’s 8th Symphony is more beautiful than his 7th Symphony.

This exchange, I believe, elicits an even stronger intuition of faultless disagreement, and objectivists would treat it as an equally objective dispute, just like *Symphony*.⁷ Objectivists may attempt to deny that it exemplifies faultless disagreement, arguing instead that such disagreement is simply impossible. However, this response is hardly compelling. As Kölbel (2004) rightly observed, we treat disagreements differently depending on the subject matter. In factual disputes—such as determining the number of people in a room—we assume that at least one party must be wrong, prompting further inquiry. By contrast, in disagreements over PPTs, such as whether roses are beautiful, we do not necessarily demand resolution. Aesthetic disagreements allow for an acknowledgment of differing perspectives without assuming error, making it possible to “agree to disagree” in a way that factual disagreements do not permit.

This is not to say that all disagreements over PPTs are faultless. People may vary in their intuitions about faultless disagreement: some may think a large portion of such disagreements are faultless, while others may find only a few to be so. Nevertheless, it is

⁵ Kölbel 2004, pp. 53–4

⁶ Or Isidora Stojanovic. See Stojanovic (2008) for an argument against the notion of faultless disagreement in general.

⁷ One might question whether comparative judgments are genuinely aesthetic. However, I see no relevant difference between *Symphony* and *Symphony** that would affect our discussion of objectivism. As mentioned in the last section, objectivists need to maintain that the ranking of beauty (whether something “is more beautiful than”) remains fixed independently of one’s standard of taste.

difficult to deny that faultless disagreement is at least *possible*. If you share this intuition, then objectivism seems inadequate as an account of aesthetic disagreement, as it fails to accommodate the possibility of faultless disagreement in such cases.

But is objectivism *necessarily* incompatible with faultless disagreement? Is there no way for objectivists to resolve the tension between their view and this phenomenon? A firm “yes” would be premature without a closer examination of what *faultless* truly means in this context.⁸ MacFarlane (2014) offers three possible interpretations of *faultless* in faultless disagreement: being epistemically warranted, being true, and not violating the constitutive norms governing belief or assertion.⁹ Let us now consider whether objectivism can accommodate *any* form of faultless disagreement.

We begin with the notion that faultless disagreement consists of epistemically warranted disagreement. This means that both parties are justified or warranted in holding their respective views—that is, each believes something their evidence gives them reason to believe. Objectivism does not inherently preclude this kind of faultless disagreement: when you assert that Bruckner’s symphonies are beautiful, your evidence (e.g., your disinterested pleasure in listening to them) provides reasons for you to believe they are beautiful. In contrast, the absence of such evidence in Paul’s mind would give him reason to believe they are not beautiful. This gave us evidence to think that faultless disagreement and objectivism are not in principle inconsistent.

However, faultless disagreement understood in this way is not particularly interesting. Most disagreements, whether in matters of personal taste, or in scientific theories, historical events, etc., plausibly involve some degree of epistemic warrant. This makes epistemic warrant a trivial basis for faultlessness, as it does not distinguish aesthetic disagreements from other forms of disagreement. Moreover, this notion of faultless disagreement is unlikely to be what Kölbel has in mind, since epistemic warrant alone does not render the disagreeing parties “free from mistake.” Under an objectivist framework, beauty is considered an objective property, meaning that even if both parties are justified in their beliefs, one must still be wrong. Thus, while objectivism can allow for epistemically warranted disagreement, this does not fully capture the intuition of faultless disagreement that challenges objectivism as an account of aesthetic disagreement.

Another way to characterize faultlessness is in terms of truth. The natural semantics for “true” dictates that it expresses the same property in every context of use—namely, the property of being true. Given this, it is clear that this characterization of faultless

⁸ One could also examine the various meanings of “disagreement.” However, I will address this point in chapter 2 rather than here, as the specific interpretation of disagreement is not essential for objectivism’s resolution of this tension.

⁹ MacFarlane offers four interpretations of faultlessness in total, with one more being *accuracy*. He defines accuracy as a belief or assertion being true relative to its context—such as a world, time, and agent. However, this difference only matters if one is held to a steady diet of regular, non-centered, non-tensed propositions. Here I take objectivism to only be concerned with such non-relativized content, so I shall treat both *truth* and *accuracy* under the general notion of *truth* to avoid unnecessary complications.

disagreement is incompatible with objectivism, as objectivists would not permit two opposing assertions to both instantiate the property of truth.

But does this pose a fundamental problem for objectivism? I argue that it does not, because the notion of faultless disagreement would be incoherent if faultlessness is understood in terms of truth. If one can coherently regard another's belief as true—using the monadic propositional truth predicate—then one could adopt the same belief without abandoning any existing commitments. This implies that both beliefs can be held simultaneously without contradiction. However, it is plainly incoherent to say, “I disagree with you, but what you believe is true.” In light of this, objectivism cannot be said to be in disadvantage if they cannot account for faultless disagreement, for this notion, understood in this way, is incoherent.

Given this, it seems that Kölbel's conception of faultless disagreement can neither refer to epistemically warranted disagreement nor to disagreement over truth—the former being trivial and the latter incoherent. This leaves one remaining alternative: understanding faultlessness as not violating the constitutive norms governing belief or assertion. The key issue here is determining what norm governs belief or assertion in disagreements over PPTs. MacFarlane, in originally proposing this alternative interpretation of faultlessness, suggests that “there is no absolute answer to the question ‘what is the norm governing assertion?’, but only a ‘perspective-relative’ answer.”¹⁰ Central to his answer is the *Reflexive Truth Rule*:

*“An agent is permitted to assert that p at context c_i only if p is true as used at c_i and assessed from c_i .”*¹¹

This rule also appears to align with Kölbel's position, as Kölbel (2004) endorses the principle that “it is a mistake to believe a proposition that is not true in one's own perspective.”¹² Under this view, “fault” refers to violations of constitutive norms.

Yet objectivists cannot accept that constitutive norms are relative. Instead, they could propose an absolute alternative, stripping away context-dependent complexities. One way to do so is suggested by MacFarlane (2012) through *The Truth Norm*, which holds that “One should assert that p only if p is true.”¹³ However, if objectivists adopt this norm as constitutive of assertion, they face a significant challenge in accounting for faultless disagreement. If A asserts p and B asserts *not- p* , only one of them can be making a true assertion, meaning the other must have violated the constitutive norm by asserting something untrue.

¹⁰ 2014, p. 103

¹¹ Ibid.

¹² p. 70

¹³ Williamson (2000: ch. 11) proposes an alternative principle: the *Knowledge Rule*—“At a context c , assert that p only if you know that p at c .” However, I find MacFarlane's (2014: 111–113) argument for the superiority of the Truth Rule over the Knowledge Rule more compelling.

1.3 Objectivism Strikes Back—Humean Absolutism

Fortunately, objectivism is not limited to the *Truth Norm* as the only non-relative alternative. Baker and Robson (2015) propose a way to preserve faultlessness while rejecting the idea that success norms for aesthetic discourse require believing only what is true. Specifically, they defend a variant of objectivism known as *Humean Absolutism*, in which the success norm for aesthetic thought and discourse is not defined in terms of truth but rather in terms of alignment with the judgment of an ideal critic. In this section, I will examine their argument in detail, as its success would demonstrate that objectivism can plausibly accommodate the phenomenon of faultless disagreement.

To begin, Humean Absolutism is grounded in three key principles inspired by Hume's *Of the Standard of Taste* (though Baker and Robson (2015) do not claim that Hume himself explicitly endorsed these views). The first principle (HA1) states that the truth of an aesthetic judgment is determined by the convergence of *ideal critics*—judges possessing strong sense, refined sentiment, extensive practice, comparative experience, and freedom from prejudice.¹⁴ On this view, an aesthetic proposition *p* is true if and only if ideal critics jointly agree that *p*, and such anchoring of aesthetic truth in the verdicts of qualified evaluators rather than subjective or perspective-relative preferences plausibly allows this approach to preserve its objectivist nature.

Importantly, unlike standard objectivist theories that might appeal to a single perfectly objective judge, Humean Absolutism holds that aesthetic truth is established through *convergence* among ideal critics. This opens the possibility that ideal critics may sometimes fail to converge in their judgments, leading to the second principle (HA2): there are aesthetic cases in which ideal critics will disagree—where some judge that *p* and others that not-*p*. Baker and Robson take inspiration from Hume's observation that even the most refined judges may vary in their assessments due to differences in temperament, cultural background, or historical context. As Hume puts it:

[T]here still remain two sources of variation [...] The one is the different humours of particular men; the other, the particular manners and opinions of our age and country. [...] [W]here there is such a diversity in the internal frame or external situation as is entirely blameless on both sides, and leaves no room to give one the preference above the other.¹⁵

Lastly, to account for faultless disagreement, Baker and Robson postulated *paracompleteness* in aesthetic truth by introducing a third principle (HA3) to account for aesthetic falsehood: an aesthetic proposition *p* is false if and only if ideal critics converge in

¹⁴ Hume (1757/1875, sect.23)

¹⁵ Ibid., sect. 28.

judging *not-p*. Just as aesthetic truth depends on agreement among ideal critics, aesthetic falsehood depends on their agreement in rejection. An aesthetic proposition *p* is false if and only if ideal critics converge in judging that *not-p*. The combination of HA1 and HA3 thus allows for a class of aesthetic propositions that are neither true nor false—cases where ideal critics diverge, and neither *p* nor *not-p* is determinately correct.

The possibility of such truth-value gaps leads directly to the principle of *HAF*: If ideal critics fail to converge on *p* and fail to converge on *not-p*, then faultless disagreement over *p* is possible. This principle asserts that when ideal critics themselves do not agree on the truth or falsehood of *p*, there is space for individuals to rationally disagree about *p* without either side necessarily being in error.

This raises a potential challenge: if neither *p* nor *not-p* is true in cases of aesthetic disagreement, does this mean that individuals who believe *p* or *not-p* are making a mistake? One implication of the *Truth Norm* (*One should assert that p only if p is true*) is that:

(T) It is a mistake to assert a proposition that is not true.

This suggests that if *p* lacks a truth value, believing *p* (or *not-p*) would be erroneous. If this were correct, faultless disagreement would be impossible, as both parties in a dispute would be making mistaken assertions.

To resolve this, Baker and Robson propose an alternative success norm:

(HAT) It is a mistake to assert an aesthetic proposition that no ideal critic would endorse, and conversely, if at least one critic endorses this aesthetic proposition, then it is permissible to assert it.¹⁶

Since *HA2* allows for cases where ideal critics diverge, both parties in an aesthetic disagreement can make legitimate assertions—neither violating the constitutive norm of assertion.

To summarize, Humean Absolutism can account for faultless disagreement while maintaining an objective framework by redefining the constitutive norm for aesthetic discourse. Instead of requiring belief in only *true* aesthetic propositions (as traditional objectivism might), it shifts the standard to adherence to the judgment of at least one *ideal critic*. This adjustment allows for faultless disagreement in cases where ideal critics themselves fail to converge on a single verdict.

Two clarifications are in order. First, this account does not suggest that all aesthetic disagreements are faultless. On the contrary, it offers a principled way to delineate the boundaries of faultlessness. Disagreement is faultless only when ideal critics themselves diverge—meaning it does not extend to all aesthetic disputes. Some aesthetic claims may be

¹⁶ The original HAT (hereafter HATO) proposed by Baker and Robson does not contain anything after “and conversely”, and they take the second half of my characterization of HAT as an implication of the first half. This inference, however, is not valid: “no ideal critic endorses a proposition” is a necessary but not sufficient condition for “being mistaken in asserting that proposition.” That is, the contrapositive does not follow from HATO. However, we can safely amend HATO to include both directions, since both are independently supported by HA1–3.

settled by the clear consensus of ideal critics, making disagreement over them untenable. For example, if ideal critics universally agree that Bruckner's symphonies are aesthetically superior to Justin Bieber's latest album, then asserting the opposite would constitute an error. However, in cases where ideal critics remain divided—such as the relative artistic merit of Impressionism versus Surrealism—Humean Absolutism permits reasonable and faultless disagreement.

Second, Baker and Robson do not claim to establish the truth of Humean Absolutism, but merely its plausibility.¹⁷ This parallels my aim in presenting this account—not to prove that objectivism is true, but to show that if Humean Absolutism qualifies as a form of objectivism, then there is a way for objectivism to accommodate faultless disagreement, provided faultlessness is understood in terms of the constitutive norm governing assertion. Consequently, I will not further scrutinize this account or engage with potential counterarguments at this stage. If one sympathizes with the idea of faultless disagreement and finds Humean Absolutism implausible, they may reject objectivism as a solution to the problem of disagreement, *punkt*. If, however, they find it plausible, then faultless disagreement does not undermine objectivism as a theory of aesthetic disagreements. That said, even if objectivism can accommodate faultless disagreement in this way, it does not follow that objectivism (or Humean Absolutism) offers a viable solution to the broader problem of aesthetic disagreement, as deeper challenges remain that threaten its overall viability.

1.4 The Beauty-Pleasure Principle

Objectivism—whether in its standard form or as Humean Absolutism—maintains that there is an objective fact of the matter about whether something is beautiful, and our use of “beautiful” always tracks this objective property. This commitment alone makes objectivism an inadequate solution to the problem of disagreement in PPTs: It either entails systematic error on the part of speakers—making it inconsistent with the *Beauty-Pleasure Principle* (hereafter *BPP*), which governs our assertions about beauty—or, to remain consistent with *BPP*, it must attribute a kind of aesthetic chauvinism to speakers, implying that they are either chauvinistic or clueless when making aesthetic judgments.¹⁸ Both outcomes are undesirable. I will now explain this dilemma in detail.

It is widely accepted that judgments of taste are grounded in feelings of pleasure or displeasure.¹⁹ This is the standard way we determine whether something is ‘beautiful.’ While one may sometimes make aesthetic judgments based on induction or expert testimony, these ultimately rely on subjective aesthetic experiences. For example, suppose a trusted music

¹⁷ p. 444

¹⁸ This argument is adapted from MacFarlane (2014, §1.1).

¹⁹ See SEP, entry Aesthetic Judgment, sec.1.1.

critic—for example, Richard Wagner—highly praises Beethoven’s *Ninth Symphony*. If we have never heard the piece ourselves, or if we have listened to it and felt no pleasure, it would seem odd to confidently judge it as ‘beautiful’ solely on Wagner’s authority. Thus, our judgments of beauty should be restricted to cases where we have direct experience of the object and should be reserved only for those things that genuinely bring us pleasure. This leads to the following principle:

(BPP) If you directly experience something (its sound, look, taste, etc.), you call it ‘beautiful’ only if it pleases you, and ‘not beautiful’ if it does not.²⁰

Objectivism is not in principle inconsistent with BPP: if everyone took pleasure in the same aesthetic objects under normal conditions, then pleasure could plausibly indicate a shared objective property of beauty. However, in reality, aesthetic responses vary significantly between individuals and cultures. Just as one person may find Bruckner’s 8th Symphony beautiful while another finds it not beautiful, our diverse experiences vindicate BPP and challenge the idea that pleasure reliably tracks an objective property of beauty. As Ginsborg (2020) notes, our ordinary experiences of beauty “are in the first instance experiences of pleasure, liking or enjoyment,” but not mere recognition of properties.²¹ This subjective mode of appreciation plays a central role in aesthetic judgment, a role that objectivist theories, whether standard or Humean, tend to marginalize. Standard objectivism treats aesthetic predicates as truth-tracking; Humean Absolutism ties them to the verdicts of an ideal critic. By treating “beautiful” as an objective property, objectivism severs the link between our subjective experience and aesthetic judgment, thereby rendering BPP—our ordinary criterion for calling things beautiful—radically defective.

An objectivist might attempt to reconcile this tension by arguing that each of us implicitly believes our own aesthetic sensibilities are attuned to *the* property of beauty, even if others’ are not. But this would attribute an implausible, unreflective chauvinism to every competent speaker. On what basis could one claim that their own sense of beauty is uniquely reliable? Moreover, our actual practices suggest we do not usually hold such chauvinistic attitudes. While we do not judge something to be beautiful or ugly solely based on other’s testimony, it is natural to place less confidence in or to re-examine, an aesthetic judgment or its object based on counter-testimony from others (e.g. whom we think are authority).²² Furthermore, it is entirely reasonable for someone to enroll in an art appreciation class to refine their taste—something that would be inexplicable if we were all chauvinistic of the infallibility of our own aesthetic judgments.²³

²⁰It’s worth noting that the “pleasure” referenced in BPP need not be limited to sensory pleasure—it can also encompass “disinterested pleasure,” as Kant suggests. Under this broader reading, BPP offers a necessary and sufficient condition for a judgment of taste.

²¹ pp. 8-9

²² This doubt and reexamination is defended as an *intuition* in Robson and Sinclair (2023).

²³ Even in cases like this, we continue to apply ‘beautiful’ in accordance with BPP. Consider an example from MacFarlane (2014): a student enrolls in an art appreciation course, fully aware that their tastes will evolve. On the first

To summarize, objectivism faces a dilemma: either (1) it renders BPP unreliable, thereby attributing systematic error to ordinary speakers, or (2) it assumes that every competent speaker unreflectively believes their own standard of beauty is the objective one. Neither option is particularly appealing. One might wonder, however, whether Humean Absolutism fares better than standard objectivism in this regard, which may potentially preserve the plausibility of objectivism in general.

1.4.1 BPP And Humean Absolutism

Indeed, at first glance, Humean Absolutism appears more compatible with BPP: a Humean absolutist could accept BPP, provided that one's use of aesthetic predicates align with an ideal critic's, thereby allowing individuals to use them when an object pleases them. Unlike standard objectivism, this position does not require untenable chauvinism, as Humean Absolutists are well aware that even ideal critics may diverge in their evaluations of the same object.

Yet, this does not mean that Humean Absolutism salvages objectivism from the tension with BPP. There are two key issues to consider: its claim to objectivity on one hand, and its compatibility with BPP on the other hand.

First, it is unclear if Humean Absolutism is genuinely an objectivist account. In section 4.1, I argued that *if* Humean Absolutism qualifies as objectivism, then objectivism can accommodate faultless disagreement. However, this "if" is not a trivial one that can be accepted without scrutiny. One criterion of objectivism is whether 'beautiful' is true or false of something on a given occasion does not depend on the idiosyncratic standard of taste of the speaker (or any other salient individual). On the surface, Humean Absolutism seems to satisfy this, as the truth or falsity of a beauty attribution does not hinge on any particular speaker's standard.

However, Humean Absolutism explains faultless disagreement by introducing a truth gap: a proposition like "Bruckner's 8th Symphony is more beautiful than his 7th" may be neither true nor false if it aligns with the judgment of some, but not all, ideal critics. However, this move sidesteps a crucial issue: it does not actually establish the truth or falsity of beauty attributions. Yet this binary truth-value is an essential aspect that we seek to preserve in aesthetic propositions. If a disagreement is faultless because both assertions lack truth-value altogether, then the disagreement itself risks becoming vacuous—why argue at all if our respective claims are neither true nor false?

Could Humean Absolutists then shift the focus from truth to proper use, redefining the criterion as "whether the use of 'beautiful' is proper on a given occasion does not depend on the idiosyncratic standard of taste of the speaker (or any other salient individual)"? This

day, when asked which painting they find most beautiful, they are unlikely to withhold judgment—despite anticipating that their opinion will change over time.

shift may preserve the binary nature of aesthetic assertions. However, this raises another problem: while Humean Absolutism detaches aesthetic judgment from personal standards, it does so only by relativizing it to ideal critics—who themselves may not share a fixed, objective standard. Differences in temperament, cultural background, and historical context mean that even ideal critics may diverge in what they regard as beautiful.

Nevertheless, Humean Absolutists may respond that this variation doesn't threaten objectivity. After all, objectivism is only a semantic thesis that requires that when we assert something to be "beautiful" on a given occasion, this "beautiful" picks out an objective property. Humean Absolutism should have no trouble in accounting for that: it ensures that "beautiful" always refers to the property that "being judged beautiful by an ideal critic." But this invites a deeper question: who are these ideal critics, and are they context-invariant? Baker and Robson do not offer a clear account. If we judge a child's painting beautiful in a kindergarten context, must our judgment align with art critics who only appreciate the art works at the Louvre? Surely not, for otherwise our assertion would not be proper. The context suggests a different standard—and likely, a different set of "ideal critics" with the sensitivity to appreciate the value in children's art. In light of this, Humean Absolutism could be interpreted in a non-objectivist sense: across different contexts, who qualifies as an ideal critic may shift, and consequently, the property that "beautiful" picks out also shifts.

One might object: just as objectivism involves a fixed threshold for beauty, could we not say that in any given situation, there is a fixed group of ideal critics determined by the context? We certainly could. But doing so implies that anyone who disagrees with the aesthetic judgment in that context must fall outside the contextually determined group of ideal critics—because how do we make sense of the idea that, in properly judging a child's painting to be beautiful, the ideal critics selected by that context, who are supposed to represent the contextually appropriate evaluative standard, could include individuals who themselves deny the painting's beauty? Therefore, given the degree of context-sensitivity that objectivism allows, Humean Absolutism is more plausibly understood in non-objectivist terms.

The second problem is that, even if Humean Absolutism qualifies as a form of objectivism, it does not clearly resolve the underlying tension between objectivism and BPP. Within the Humean framework, the aesthetic judgments of ordinary speakers (non-ideal critics) remain alienated from their own evaluative experiences.²⁴ When a non-ideal critic judges something to be beautiful, that judgment is not ultimately grounded in their own pleasurable response but is instead validated only by whether it aligns with the standard of an ideal critic. This external standard displaces the speaker's own sense of beauty as a source of

²⁴ Humean Absolutism might be able to accommodate BPP if we focus solely on ideal critics, since one could argue that their judgments are guided by their own preferences. However, this compatibility cannot plausibly extend to the general population, as applying the same norm universally would collapse Humean Absolutism into an expressivist account rather than an objectivist one.

justification, even though it is precisely their first hand experience of aesthetic pleasure that motivates the judgment in the first place. As such, Humean Absolutism fails to explain why this pleasurable experience is still necessary for non-ideal critics to make judgments of beauty—why, for example, someone would not call something beautiful unless they found it pleasing themselves. If the speaker’s experience plays no justificatory role in determining the correctness of their judgment, its necessity remains a mystery. In this way, BPP—which treats aesthetic pleasure as both necessary and sufficient for making a judgment of taste—is rendered unreliable. Therefore, even under Humean Absolutism, the ordinary speaker is estranged from their own criteria for aesthetic judgment, and their use of “beautiful” is systematically undermined by deference to a standard external to their experience.

To conclude, standard objectivism either attributes systematic error to ordinary speakers or assumes an unreflective chauvinism in their judgments. Humean Absolutism does not salvage the view, for it is not genuinely objectivist: rather than securing objectivity, it effectively concedes that no objective property of beauty underlies our aesthetic discourse. Still, the failure of objectivism in this form does not settle the broader question of whether there could be an objective standard of beauty. What matters for present purposes is that ordinary uses of beauty-predicates do not seem to semantically track such a standard. The difficulty for objectivism lies in assuming that they do. Given this, a non-objectivist framework is needed to account for disagreement over PPTs, one that does not foreclose the possibility of objectivity but refrains from building it into the semantics of our everyday judgments.

1.5 “Beautiful” and the Predicates of Personal Taste

Before proceeding, I would like to preemptively defend the view that there is no categorical *semantic* difference between aesthetic adjectives (e.g., *beautiful*) and standard PPTs (e.g., *tasty*, *funny*). In previous sections, I have used aesthetic disagreement and disagreements about PPTs (hereafter PPT) interchangeably. For those familiar with the literature, this might seem concerning, as standard discussions of PPTs typically focus on terms like *tasty* or *funny*, and it is not immediately obvious that this category includes terms of traditional aesthetic interest like *beautiful*. However, I argue that this concern is misguided: there is no fundamental semantic distinction between aesthetic predicates and PPTs.

The intuition that there is such a distinction can be captured, following Sundell (2017), in terms of what he calls *aesthetic pressure*. Aesthetic pressure refers to the normative demands that an assertion involving certain evaluative terms places on others. When we say that something is *tasty* or *funny*, we typically do not impose much pressure on our listener to agree. If someone disagrees, we might simply acknowledge a difference in preference without expecting justification. By contrast, when we describe something as *beautiful*, we do place pressure on our interlocutor to agree or to justify their dissent. Such assertions implicitly “lay

claim to correctness” (Zangwill, 2005), demanding either endorsement or reasoned opposition.

This intuition forms the basis of *the standard view*, which suggests that high-pressure terms like *beautiful* aim at a form of objectivity, making claims that are not tied to individual taste, whereas low-pressure terms like *tasty* are naturally interpreted as being about the speaker’s personal standards. This distinction would explain why aesthetic disputes often seem deeper and less easily defused than disagreements about PPTs. However, I argue that this does not amount to a categorical semantic difference.

If the distinction between aesthetic predicates and PPTs were truly built into their semantics, we would expect high-pressure terms to always generate deep and intractable disputes, while low-pressure terms would always allow for easy resolution. But real-world discourse does not behave this way. As convincingly argued by Sundell (2017), aesthetic pressure is not static but dynamic, shifting based on conversational context rather than being an inherent property of specific terms. A term like *tasty* may carry very little normative weight in casual conversation, but in certain contexts, it can provoke prolonged and serious disagreement. Likewise, while *beautiful* is often assumed to demand consensus, it can also be used in a way that allows for subjective divergence.

Consider this exchange about taste:

Alphie: *Vegemite is tasty.*

Betty: *No, you’re wrong. Go try it again. Have you noticed how salty it is?*

Alphie: *I have. I still think it’s tasty.*

Betty: *You only like it because of your fond memories of Australia.*

Alphie: *That’s not true. Have you noticed how the maltiness of the vegemite complements the flavor of the bread?*

Betty: *Yes, but I don’t want something that makes the bread even maltier. I want something that adds some richness or sweetness, like good old butter or jam. Can’t we agree about anything?*²⁵

This exchange demonstrates that discussions about *tastiness*—supposedly a low-pressure term—can escalate into serious and persistent debates. The participants argue about specific properties of Vegemite and attempt to persuade each other, just as they might in a dispute over an aesthetic judgment.

Conversely, consider a conversation about *beauty*:

Alphie: *Titanic is a terrible movie.*

Betty: *I can see why you feel that way, but to me, it’s beautiful.*

Alphie: *Fair enough.*

Here, rather than engaging in deep disagreement, the speakers acknowledge differing standards and move on. One might think this dialogue is disappointing, since they could

²⁵ This and the following example are from Sundell (2017).

have dug deeper why each holds their respective view, but this exchange is not bizarre or infelicitous. If *beautiful* were inherently a high-pressure term that always compelled agreement, we would expect Alphie and Betty to continue debating. Instead, they settle the issue much as they would in a disagreement over *tastiness*—by recognizing subjective differences. These examples illustrate that both PPTs and aesthetic predicates can function as either high-pressure or low-pressure terms depending on context. Some debates over *tastiness* could become highly charged, while some disagreements over *beauty* could dissolve into personal preference. This variability suggests that aesthetic pressure is not a function of semantic structure.

Chapter 2

Non-objectivist Theories

In this chapter, I examine two leading non-objectivist theories—contextualism and relativism—as candidates for explaining the nature of disagreement over PPTs, while preserving the plausibility of BPP.²⁶ The first half of this chapter proceeds in three stages. First, I discuss standard contextualism, which holds that judgments of taste express different contents in different contexts of utterance. Next, I turn to nonindexical contextualism—a more refined variant that, while often classified as a form of relativism by theorists other than MacFarlane, remains importantly distinct. It allows for variation in the truth value of the same proposition across different contexts of assessment, thus aligning it more closely with relativism in spirit. Finally, I examine MacFarlane’s own relativism, which introduces a crucial refinement: it locates truth not solely in the context of utterance, but also in the context of assessment. Each theory offers distinct resources for capturing the disagreement that characterizes disputes over taste, and the analysis in this chapter engages closely with MacFarlane’s detailed exposition in *Assessment Sensitivity* (2014).

The remainder of this chapter is to assess whether these theories can adequately capture the varieties of disagreement at play in aesthetic discourse. I begin by arguing that standard contextualism struggles to explain what I previously termed “genuine disagreement,” since it allows for conflicting assertions to express compatible contents across contexts, which reduces disagreements to mere talking past each other. I then examine whether nonindexical contextualism and relativism fare better. While both appear to capture a more substantive sense of disagreement, I argue that the notion of “genuine disagreement” itself, as discussed in chapter one, is ambiguous and potentially question-begging. To sharpen the analysis, I adopt a framework developed by MacFarlane (2014), which distinguishes among three forms of disagreement: practical noncotenability, doxastic noncotenability, and preclusion of joint accuracy. Through this lens, I argue that relativism appears to be superior to its rivals: standard contextualism accounts only for practical

²⁶ Another widely discussed approach is expressivism, but I will not examine it in this thesis. Standard expressivism faces well-known challenges, such as the Frege–Geach problem. While some non-standard forms of expressivism attempt to resolve this by incorporating relativist compositional semantics, they begin to resemble forms of relativism rather than providing a distinct alternative. There may be other viable expressivist strategies, but addressing them falls outside the scope of this thesis. More importantly, I believe there is good reason to think that disagreement about whether something is beautiful involves more than mere divergence in attitude. I thus set expressivism aside in this discussion.

noncoterability, nonindexical contextualism captures both practical and doxastic noncoterability, and relativism accommodates all three.

2.1 Contextualism

In general, contextualism about PPTs maintains that such predicates pick out different properties in different contexts. While theorists may differ on what exactly constitutes the relevant context, the standard view takes it to involve the speaker's (or other salient individuals') aesthetic standard. On this view, a sentence like "Bruckner's Eighth Symphony is beautiful," when uttered by A, is true just in case Bruckner's Eighth Symphony meets A's aesthetic standard—e.g., if A likes the way it sounds.

There are two main ways to interpret how this context dependence works. On the first interpretation, "*beautiful*" is treated as a one-place predicate, with the context supplying an implicit indexical reference to the speaker. The logical form might be represented as $P(x)$, where P stands for "beautiful-on-my-aesthetic-standard" or "beautiful-relative-to-the-context." On this view, the context determines how the predicate applies without introducing an additional argument. On the second interpretation, "*beautiful*" is treated as a two-place predicate, represented as $P(x, y)$, where x is the object and y is the relevant person or aesthetic standard. Here, "Bruckner's Eighth Symphony is beautiful" means something like "Bruckner's Eighth Symphony is beautiful-for-A," though the second argument is typically implicit and filled in by context.²⁷

Despite this difference in formal representation, the two views are functionally very similar. Both account for the context sensitivity of predicates like "*beautiful*" by tying their truth conditions to an aesthetic standard provided by the context. This is much like how adjectives such as "*tall*" function: someone can be tall for a gymnast but not for a basketball player, with the comparison class supplied by context.

Thus, when someone says "Bruckner's Eighth is beautiful" without qualification, contextualism holds that they are ascribing a context-sensitive property implicitly anchored to a particular aesthetic perspective. The choice between one-place and two-place predicate analyses is largely a technical matter. For the purposes of this discussion, I remain neutral between them. What is central to contextualism is that PPTs like "*beautiful*" are semantically dependent on a speaker- or context-supplied aesthetic standard.

Compared to objectivism, one major advantage of contextualism is the ability to accommodate BPP: it allows aesthetic judgments like "beautiful" to be appropriately grounded in the speaker's own pleasurable experience. Because the standards for applying aesthetic predicates shift with context—often reflecting the speaker's evaluative

²⁷ MacFarlane (2014)

perspective—contextualism preserves the central role of pleasure in making aesthetic judgment. This flexibility also enables contextualism to naturally explain faultless disagreement. For instance, in the case of the Symphony, your judgment that Bruckner’s 8th is beautiful is true if it accords with your experience of pleasure, without implying that your friend Paul, who judges it not to be beautiful, is thereby making a mistake.

However, standard contextualism is often deemed inadequate as a theory of disagreement of PPTs. A main line of objection is that it cannot accommodate disagreement. Since it does not take two disagreeing parties to express inconsistent propositions, it would reduce disagreement to mere talking past each other. And the other major line, put forward by MacFarlane (2014), is that it cannot accommodate retraction—cases where the speaker rejects the assertion they previously made when their standard of beauty changes.

Let us begin by examining the tension between standard contextualism and the phenomenon of disagreement. According to contextualism, PPTs—such as “beautiful”—are context-sensitive: they pick out different properties depending on the speaker’s context (i.e. their aesthetic standard). Thus, when A says, “This painting is beautiful,” and B responds, “This painting is not beautiful,” A’s utterance expresses ‘this painting is beautiful-relative-to-A’s-standard’, and B’s utterance expresses “the painting is not beautiful-relative-to-B’s-standard”. Since each speaker is attributing a different property to the painting, the two utterances are not strictly incompatible—they can both be true in their respective contexts. This allows contextualism to preserve the idea that both parties may be “right in their own way,” thereby accounting for the apparent faultlessness of such exchanges.

However, this advantage comes at a significant cost. By treating each speaker’s utterance as expressing a distinct proposition, contextualism undermines the sense that a genuine disagreement is taking place. In the symphony case, for instance, it would be misleading to gloss the exchange by saying, “You think Bruckner’s Eighth Symphony is beautiful, but Paul thinks it is not,” since this would imply that both parties are speaking to the same content. More specifically, it would misrepresent your belief as addressing the property assigned by your standard of beauty to the word “beautiful,” rather than reflecting the property assigned within Paul’s context.

To illustrate, compare aesthetic predicates like “beautiful” with paradigmatic context-sensitive terms such as “local.” If Abby says, “Paul’s favorite bar is local” (meaning local to London), and Sam replies, “It’s not local” (meaning local to Paris), there’s no real contradiction. Both speakers can agree that the bar is local to one place but not the other, and a dialogue like the following sounds perfectly natural:

Abby: Paul’s favorite bar is local.

Alan: I agree, but it isn’t local to me.

However, the aesthetic case seems quite different. If you say, “Bruckner’s 8th Symphony is beautiful,” and Paul replies, “I agree, but it isn’t beautiful to me,” the exchange sounds jarring and confused. On the other hand, the following disagreement feels appropriate:

You: Bruckner’s 8th is beautiful.

Paul: I disagree, though it might be beautiful to you.

By contrast, if we return to the “local” case and imagine:

Abby: Paul’s favorite bar is local.

Alan: I disagree, though it might be local to you.

—this would suggest a misunderstanding but not a genuine disagreement. These patterns suggest that standard contextualism may struggle to capture our intuitive sense that people genuinely disagree about judgments of beauty, even when they are aware that those judgments are grounded in differing personal standards. This in turn, suggests that if we understand aesthetic disagreement to represent a case of the intuitive understanding of genuine disagreement, in which two parties are expressing inconsistent propositions, something more than contextual variation of personal standard of beauty is at play.

One way for contextualists to respond to this challenge is to deny that disagreement over personal taste—such as judgments of beauty—is genuine disagreement. After all, people often say “*De gustibus non disputandum.*” However, there are two reasons to think that disagreement, at least, can occur in such cases.

First, there are situations where linguistic denials—such as “No,” “You’re wrong,” or “That’s false”—are entirely felicitous responses to aesthetic assertions. Consider *Symphony* again:

You: Bruckner’s 8th Symphony is beautiful.

Paul: No, it’s not beautiful.

This response sounds natural and directly contradicts the original claim. However, if Paul were merely expressing that the symphony does not strike *him* as beautiful, this kind of direct denial would seem out of place. Consider:

You: Bruckner’s 8th sounds beautiful to me.

Paul: #No/That’s false, it doesn’t sound beautiful to me.

The infelicity of Paul’s second reply suggests that, in *Symphony*, Paul was not merely making a claim about how Bruckner’s 8th does not sound good to *him*, but rejecting something about/in what you said.

In absence of any argument, the default candidate for the object of rejection is the asserted proposition. Yet, of course, contextualists may respond that rejection markers such as “No,” “That’s false,” or “I disagree” can target other elements of discourse. For example, they may point out that such markers can sometimes reject the content of reported speech (e.g., “Sahin said you had a car”—“No, I don’t have a car”), a presupposition (e.g., “Your wife is very beautiful”—“No, we’re not married”), or a shared background assumption, as in

Grice's famous case: "Either Wilson or Heath will be the next Prime Minister"—"I disagree; it will be either Wilson or Thorpe." In these cases, what is being rejected is not the literal proposition asserted, but some implicated or presupposed content.²⁸

However, without a compelling argument, these alternative targets of rejection seem to be difficult to extend to aesthetic disagreement cases. In a dialogue such as:

You: Bruckner's 8th Symphony is beautiful.

Paul: No, it's not beautiful,

the second speaker appears to directly negate the proposition asserted. This is unlike cases involving presupposition or reported speech, and intuitively resembles straightforward contradiction.

That said, I do not believe these possible targets of rejection exhaust the contextualist's resources, nor are they necessarily inapplicable to disagreements about personal taste. In Chapter 3, I will explore other plausible candidates for what might be at issue in aesthetic disagreement—candidates that could offer a more promising escape route for contextualism—and develop a detailed account of how such a view might account for genuine disagreement. For now, however, it is reasonable to conclude that, in the absence of a more developed defense, contextualism struggles to explain the felicity of denial in disagreements over personal taste.

The second reason why there can be said to be disagreement in matters of personal taste is that such disagreements can escalate into heated debates. People don't just express different reactions to a work of art; they often try to convince each other, and offer reasons why, say, Bruckner's 8th Symphony is (or isn't) beautiful: "Beautiful? It's bloated and overindulgent! Do you also think bombast is beauty? Doesn't its lack of clarity undermine its aesthetic value?" Such arguments are not merely about differing subjective responses; they suggest that the participants take themselves to be contesting the correctness of each other's aesthetic judgment. Importantly, we typically do not argue about what is merely beautiful *to* someone—just as we don't argue about someone else's preferences or what they find enjoyable. That we do argue in these cases speaks against the standard contextualist interpretation, which treats aesthetic predicates like "beautiful" as context-sensitive expressions tied to individual standards.

The contextualist might reply that our perception of disagreement is illusory—that we are simply mistaken in taking ourselves to be contradicting one another. However, if contextualism requires attributing this level of systematic error to competent speakers, its theoretical advantage over objectivism becomes questionable. After all, the primary objection to objectivism was precisely that it saddled ordinary speakers with pervasive error in their aesthetic discourse. Yet the contextualist who denies the reality of disagreement must similarly explain why speakers systematically take themselves to be disagreeing. To do so, they

²⁸ Cf. MacFarlane (2014)

must assume that speakers are implicitly committed to an objectivist semantics they do not in fact endorse—an “inchoate objectivism” that renders their discourse semantically confused. But if this is how we interpret aesthetic talk, then the contextualist must still account for why speakers continue to assert claims like “That’s beautiful” based solely on personal response, despite widespread recognition of the diversity of such responses. The explanation would therefore have to appeal to the same kind of unreflective chauvinism that was found objectionable under objectivism, now coupled with a semantic error that objectivism itself does not posit.²⁹

If the preceding argument holds, then genuine disagreement over matters of personal taste is possible, and contextualism falls short in accounting for it. One potential response for the contextualist, however, is to argue that disagreement need not always involve logically inconsistent propositions. If alternative forms of disagreement can be identified, contextualism may be able to capture at least some of them. The next section explores this strategy in more detail. As I will argue, however, this move ultimately fails to provide contextualism with a significantly stronger position in accounting for disagreement when compared to its theoretical rivals.

Another challenge for contextualism is the phenomenon of retraction. Consider the case of a speaker who, years ago, confidently asserted that a particular painting was beautiful. Now, with broader exposure and a more refined sensibility, they no longer find the piece beautiful. When reminded of their earlier statement—“But you said that painting was beautiful”—the speaker may respond, “I take that back.” They would not be inclined to say, “It was beautiful then, but it isn’t anymore,” as that would imply the painting itself somehow changed. Nor would they say, “At the time I just meant it was beautiful to me,” because they didn’t mean that. Rather, they took themselves to be genuinely disagreeing with those who thought the symphony lacked beauty.

This poses a problem for contextualism. On a contextualist account, the earlier utterance (“This painting is beautiful”) expressed a proposition relative to the speaker’s standard then, while the later utterance (“This painting is not beautiful”) expresses a different proposition, relative to a changed standard that the speaker now has. As such, the two statements are compatible, and no retraction should be necessary. This tension is clearer when compared to uncontroversially context-sensitive claims. For instance:

Alphie [in Phoenix]: You can get a swamp cooler at any local hardware store.

Alphie[in Boston the next day]: Nobody sells swamp coolers around here.

Betty: But you said you can get one at any local hardware store!

Alphie: I take that back.³⁰

²⁹ MacFarlane 2014, pp. 11-2

³⁰ Ibid., pp. 13-4

Here, the retraction feels out of place. Since both statements are true in their respective contexts, no retraction is needed.

To summarize, while contextualism captures the intuition that the usage of “beauty” is determined by the speaker’s standard of taste (BPP), it does so at the cost of making genuine disagreement and retraction inexplicable. An alternative account is thus called for.

2.2 Relativism

I shall now turn to relativism, focusing specifically on John MacFarlane’s influential version, which centers on the notion of assessment sensitivity. His account, developed in detail in *Assessment Sensitivity* (2014), will serve as the primary target of critique in the remainder of this chapter. Before presenting MacFarlane’s position, however, it is important to distinguish it from closely related views—particularly the one he refers to as *nonindexical contextualism*. While MacFarlane classifies this view as a form of contextualism, it has often been treated by others as a branch of relativism. I shall elaborate on this deviance of the usage of the term later in this section. For consistency, I will follow MacFarlane’s terminology and refer to it as nonindexical contextualism.

Nonindexical contextualism can be summarized as follows: On this view, the truth of PPTs like “beautiful” is use-sensitive. That is, while the proposition expressed by an assertion remains constant, the truth of the assertion is evaluated relative to parameters determined by the context of use. These parameters include not only standard ones like world, time, and speaker, but also a contextually supplied aesthetic standard. Crucially, the proposition does not contain an implicit indexical such as “for me” or “according to my taste”; rather, the relevant standard is fixed externally by the context.

This contrasts with standard (indexical) contextualism, which holds that the content of aesthetic statements shifts with the speaker’s individual standard of taste at the time of utterance. On that view, the proposition expressed by “Bruckner’s 8th symphony is beautiful” may differ from speaker to speaker, as it implicitly includes something like “according to my standards.” By contrast, nonindexical contextualism maintains that the proposition expressed remains the same across contexts, though the truth value of the assertion can vary depending on the aesthetic standard made salient by the context of use.

We shall now clarify two things: what is meant by *content* in this framework, and how *context* functions within it. Following MacFarlane, we adopt the Kaplanian understanding of content and the model of context articulated in *Demonstratives*. Kaplan offers a nuanced account of the content expressed by sentences containing indexicals. Fortunately, we need not replicate his full analysis here, as judgments of taste typically do not contain indexicals. As Kaplan notes, “Words and phrases with no indexical element express

the same content in every context; they have a fixed character.”³¹ Accordingly, a sentence like “*Bruckner’s symphony is beautiful*” expresses a standard-of-taste-invariant, agent-neutral proposition—*Bruckner’s symphony is beautiful*—rather than a relativized content such as *Bruckner’s symphony is beautiful to Abbey*, as standard contextualism would have it.

As for the understanding of *context*: in Kaplan’s framework, a context of use is a possible occasion of utterance, specified by parameters such as world, time, agent, and location.³² This is often represented as a “centered possible world,” with a particular agent, time, and place at its center.³³ The context determines the *content* expressed by an utterance. By contrast, Kaplan defines a circumstance of evaluation as “both actual and counterfactual situations with respect to which it is appropriate to ask for the extensions of a given well-formed expression. A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well.”³⁴ On his view, a sentence is true at a context of use just in case the content it expresses at that context is true at the corresponding circumstance(s) of evaluation.³⁵ Thus contents have extensions only relative to circumstances of evaluation, but their truth at a context is absolute once the relevant world and time are fixed by the context.

Nonindexical contextualism extends this Kaplanian framework by treating the speaker’s aesthetic standard as an additional coordinate of the circumstance of evaluation. Once that standard is fixed by the context of utterance, the proposition expressed by an utterance has an absolute truth value, which preserves contextualism’s commitment to truth being determined at the context of utterance. While some philosophers (e.g. Kölbel 2008) have argued that this extension collapses into relativism, MacFarlane counters that this is not the case: in Kaplan’s terms, the standard of taste here is supplied by the context of use, so no assessment sensitivity is involved.

This helps clarify a difference in terminology between my usage and MacFarlane’s. Previously, I characterized contextualism as the view that predicates of taste express different contents in different contexts, which means such expressions are context-of-use-indexical. Under that definition, nonindexical contextualism would not count as a contextualist view, since it holds content fixed and varies only the truth value across circumstances of evaluation. However, MacFarlane classifies it as contextualist for two reasons: (1) although taste expressions are not indexicals, their truth-value still depends on features of the context of use; and (2) like standard contextualism, nonindexical contextualism assigns absolute truth values to propositions once the context of utterance is fixed. Nevertheless, this terminological

³¹ Kaplan 1989, p. 521

³² Ibid., p. 494

³³ MacFarlane 2014, p. 53

³⁴ Kaplan 1989, p. 502

³⁵ Ibid., p. 522

discrepancy does not affect the substance of my thesis, as nonindexical contextualism is discussed here solely as a foil to MacFarlane's own relativism.

Above is the basic outline of nonindexical contextualism. Intuitively, it is a more plausible candidate than both objectivism and standard contextualism: it preserves the plausibility of BPP, since it allows the speaker's own standard of taste to determine the extension of the proposition expressed by an aesthetic judgment (e.g., "this painting is beautiful"). It also seems to account for disagreement, since two individuals who assert "Bruckner's 8th is beautiful" and "Bruckner's 8th is not beautiful" can be understood as making conflicting claims about the same proposition, namely one whose truth depends on the contextually supplied standard of taste.

However, despite this apparent advantage over the other two theories, nonindexical contextualism does not constitute the ideal theory, for it struggles to account for retraction. The core problem is that, under nonindexical contextualism, once the truth value of an assertion is fixed by the context of utterance, subsequent changes in taste do not alter its normative status. This is because nonindexical contextualism is governed by the *Reflexive Truth Rule*, which states:

"An agent is permitted to assert that p at context c_i only if p is true as used at c_i and assessed from c_i "³⁶.

In other words, an assertion was permissible if it is true given the speaker's standards and context at the time of speaking (context of use). There is no further requirement to revisit it, even if the agent's evaluative standards change.

This is where the difference between nonindexical contextualism and relativism kicks in. While both theories accept the *Reflexive Truth Rule*, relativism supplements this with a crucial addition: a *Retraction Rule* that governs what should be done when one's standards shift over time. The rule states:

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

This norm captures what nonindexical contextualism fails to accommodate—that agents may be *required* to retract past assertions when their current evaluative standpoint yields a different truth value. The distinction between nonindexical contextualism and relativism marked by this additional rule becomes especially vivid when we consider not only how assertions are evaluated at the time they are made, but also how they are treated retrospectively.

But why does retraction matter in the first place? For MacFarlane, retraction is not merely a matter of changing one's mind—it is a normative speech act that undoes the commitments introduced by a previous utterance.³⁷ To retract an assertion is to publicly

³⁶ Ibid., p. 103

³⁷ Ibid.

withdraw one's endorsement of its content, thereby releasing oneself from the obligation to defend it and releasing others from the entitlement to rely on the speaker's authority. Just as retracting a question cancels the addressee's obligation to respond, and retracting an offer cancels the recipient's entitlement to accept it, retracting an assertion cancels the speaker's prior assertoric commitment.

Relativism provides a natural account of this practice by recognizing that truth is *assessment-sensitive*—that is, it depends not just on the context in which a statement is made, but also on the context from which it is evaluated. When a speaker's present standards diverge from those in place at the time of the original assertion, relativism holds that the assertion may no longer be true relative to the current context of assessment. In such cases, the Retraction Rule obliges the speaker to retract.

Nonindexical contextualism, since it does not encompass such sensitivity, cannot explain this obligation. A speaker with a changed taste should now judge her past assertion to be false—but that judgment alone is insufficient grounds for retraction. This is because she would only be obligated to retract if she also believed that the assertion had been false by her own standards at the time, which she may no longer endorse but must still treat as once sincerely held. MacFarlane likens this to the temporalist, who can coherently say that her past assertion “It is noon” is now false without being obligated to retract it, since it was not false when it was uttered.

This comparison between nonindexical contextualism and relativism brings into sharp relief the need for a framework that allows truth to be relativized not only to the context of use but also to the *context of assessment*. Without this dual relativity, we cannot make sense of the normative pressures that accompany changing standards.

In light of the preceding discussion, we are now in a position to articulate what MacFarlane's relativism entails. Like indexical contextualism, it allows the truth (extension) of an expression of a judgment of taste to be sensitive to the speaker's evaluative standard. However, unlike nonindexical contextualism, it does not treat this sensitivity as fixed solely at the context of use. Instead, it extends the consideration to the context of assessment. Let us now turn to its treatment of disagreement over PPTs, and examine its alleged superiority over competing theories.

2.2.1 Disagreement Revisited

Recall that in Chapter 1, I defined *genuine disagreement* as follows: there is some proposition p such that Speaker A asserts p and Speaker B asserts q , where q entails not- p . This definition has served relatively well throughout the preceding discussion. However, at this stage, it risks appearing overly simplistic and insufficiently examined. Two concerns arise. First, on what basis can we take this to be the definitive criterion for disagreement? Second, if this definition is accepted without further analysis, it becomes difficult to assess the comparative merits of

different semantic theories solely as accounts of disagreement. For instance, relativism and nonindexical contextualism might appear equally successful under this criterion, even though the former encodes a more sophisticated and explanatory treatment. To evaluate relativism's account of disagreement properly—and to understand what is at stake—it is necessary to develop a more fine-grained understanding of what disagreement, particularly those over PPTs, consists in.

Let's start from the first question. Can we take disagreement to be precisely what has been formulated in GD? According to Sundell (2011), the answer is no. He argues that GD is ambiguous and ultimately question-begging.

Its ambiguity stems from the conflation of two distinct phenomena associated with disagreement: (i) the *felt conflict between speakers*—cases where speakers take themselves to be at odds—and (ii) the *felicity of linguistic denial*—cases in which one speaker can appropriately respond to another with an explicit denial (e.g., “No, it's not”).

While these two phenomena often coincide, they can come apart. For instance, in the exchange:

Sam: I like Dave.

Sara: I don't like Dave.

Although Sara clearly adopts an opposing stance, it would be infelicitous for her to reply:

Sara: “#No, I don't like Dave,”

Yet she might naturally and appropriately say:

Sara: “Well, I don't like Dave. In fact, I think you're dead wrong about him.”

This illustrates that the presence of a felt disagreement—an impression that the speakers are in conflict—does not guarantee the felicity of denial. Importantly, such cases are not rare anomalies; they are common and linguistically natural. In this way, the two phenomena come cleanly apart, and any theory of disagreement must therefore clarify which of them it aims to explain.

One might attempt to preserve GD by appealing to an intuitive notion of “real disagreement,” hoping to bypass this ambiguity. But as Sundell points out, this move fails. If the term “really” is doing any meaningful work in the phrase “really disagree,” it must imply that the disagreement concerns the truth of a proposition—specifically, that one speaker asserts p and another asserts a proposition that entails $\neg p$. However, the consistency or inconsistency of the propositions expressed is exactly what is under contention in disagreements involving PPTs. Thus, to appeal to GD is to presuppose the very thing it is meant to analyze.

Accordingly, we should not assume that all genuine disagreements must involve contradictory propositions. Rejecting this assumption, however, does not yet tell us what disagreement should consist in. A useful starting point is the two plausible candidates offered

by Sundell: the felicity of denial and the felt conflict. The question then arises: are both of these genuine forms of disagreement, or does only one capture the phenomenon adequately?

Sundell argues that genuine disagreement should be understood in terms of the felicity of denial. He proposed the following definition:

Disagreement =df The relation between speakers that licenses linguistic denial.

Yet this definition may appear unsatisfying. Sundell does not explain why disagreement should be anchored in the felicity of denial rather than in the felt conflict. What, for instance, rules out cases of felt conflict without denial as genuine disagreement? And if such cases do not involve disagreement, what explains the conflict that nonetheless arises?

These questions, although left implicit, are not unanswerable. Sundell's rationale seems to be that cases of mere divergent attitudes—for example, one person liking Dave while another dislikes him—differ from uncontroversial disagreements such as those over factual matters, like how old Dave is. To understand this difference, let's first clarify the nature of felt conflict.

What underlies such conflict is what MacFarlane (2014) calls *practical noncotenability*: two attitudes are in tension when an agent cannot coherently adopt both without a change of mind. If Sam likes Dave and Sara dislikes him, or if one person admires the symphony while another is repulsed by it, their *attitudes* are incompatible in a way that would render holding both practically incoherent.

Yet just like Sundell, we might not want to classify cases which exemplify only practical noncotenability as disagreements. Suppose A wants only the icing of a cupcake, and B wants only the sponge. A and B have practically noncotentable attitudes, but they seem to be in a happy state of harmony. Disagreement, by contrast, feels like a dispute: to disagree with someone is not just to have a different attitude, but to have a kind of tension whose resolution depends on one or both parties changing their minds.³⁸ And that is why Sundell is right to anchor disagreement in the felicity of denial, rather than in the felt conflict alone.

However, even if we grant Sundell's criterion, it does not, by itself, explain what makes denial felicitous in such cases. Our broader goal is to understand the nature of disagreement over PPTs by identifying the different forms disagreement can take. Merely pointing to denial as a marker of disagreement does not explain what underlying features make denial appropriate. For example, in the case of Sara and Sam, we see that it is the noncotenability of their evaluative attitudes that drives the sense of disagreement. Thus, rather than treating felicity of denial as foundational, it is more productive to identify the structural sources—like noncotenability—that give rise to it.

In Chapter 6 of *Assessment Sensitivity*, MacFarlane distinguishes, in addition to practical noncotenability, two further forms of disagreement that arise in disputes over PPTs:

³⁸ 2014, p. 123

doxastic noncotenability and preclusion of joint accuracy. The former can be seen as extensions of the basic notion of practical noncotenability. In the case of doxastic noncotenability, two individuals are said to disagree when one cannot coherently adopt the other's belief without abandoning or revising their own. For instance, if George believes that all bankers are rich, while Sally believes that Vern is a poor banker, George cannot accept Sally's belief without giving up his own.

Another form of disagreement—preclusion of joint accuracy—requires more explanation. While doxastic noncotenability concerns the coherence of holding two beliefs simultaneously, this second kind targets particularly the disagreement involving nonstandard content, like centred proposition (which has truth values relative to a world, a time, and an agent as “center”). While it might not be distinct from doxastic noncotenability for one who endorses regular, non-centered, non-tensed propositions, this distinction becomes especially important when we consider *assessment-sensitive* or *relativized* contents—those whose truth values depend on contextual features like the speaker, time, location, etc. This makes the distinction particularly relevant for evaluating the strength of relativism in contrast to nonindexical contextualism.

To understand this, we must first clarify what *accuracy* means, and how it differs from *truth*. The content of a belief or assertion can be true or false, but when we ask whether the belief or assertion itself is accurate, we are asking whether its content is true relative to its context. These two notions are different: even if we now judge the content of David's past belief—say, “I am eating a sandwich”—to be false from our current perspective, we can still regard it as having been accurate, because it was true relative to its context: <actual world, yesterday, David>. MacFarlane further formalizes this as follows:

Accuracy. An attitude or speech act occurring at context c_1 is accurate, as assessed from a context c_2 , just in case its content is true as used at c_1 and assessed from c_2 .³⁹

With this understanding of accuracy in place, we can now examine the difference between doxastic noncotenability and preclusion of joint accuracy more closely. Consider the case where Andy believes the centered proposition “I am eating a sandwich,” and David believes “I am not eating a sandwich.” These beliefs are doxastically noncotenable—Andy cannot adopt David's belief without abandoning his own—but both beliefs may nonetheless be accurate if Andy is indeed eating and David is not, each at their respective times and locations. Conversely, even beliefs that are cotenable—i.e., that can be held together without incoherence—may still preclude one another's accuracy. Suppose Andy believes at 2 p.m. that “I am eating a sandwich,” and at 3 p.m., David believes “No one was eating a sandwich an hour ago.” These beliefs are not logically incompatible and could be jointly held. However, if Andy was in fact eating at 2 p.m., then David's belief must be inaccurate.

³⁹ Ibid., p. 127

Thus, while noncotenability captures a sense of doxastic tension, *preclusion of joint accuracy* seems to capture an arguably more robust notion of disagreement—one that aligns with our intuition that if two people truly disagree, they cannot both be right (relative to a single context). In this stronger sense, disagreement involves a clash in which the accuracy of one person’s belief or assertion rules out the possibility of the other’s being accurate.

Before we proceed to evaluate the strength of relativism in addressing the problem of disagreement, we must first examine whether disagreements involving PPTs genuinely exhibit the forms of disagreement discussed above.

Consider the case of *Symphony*. It certainly involves *practical noncotenability*. Even if Paul’s statement ‘Bruckner’s 8th is not beautiful’ does not contradict the semantic content of your assertion—perhaps because each of you asserts a different proposition anchored to your own aesthetic standards—it remains the case that you and Paul express incompatible evaluative attitudes. You cannot coherently adopt Paul’s attitude without revising your own, and vice versa.

Nevertheless, it would be implausible to assume that *practical noncotenability* exhausts the forms of disagreement operative in cases such as *Symphony*. As discussed in the critique of contextualism, there are many situations where linguistic denial—through expressions like “No,” “You’re wrong,” or “That’s false”—emerges as a perfectly felicitous response to aesthetic assertions. Such denial is difficult to account for if the disagreement involves only a clash of attitudes. Just as it is infelicitous to respond to “I like Dave” with “No, that’s false”, it is puzzling why similar denials sound entirely natural in aesthetic disagreements, unless something more than attitudinal conflict is at stake. In these instances, the felicity of denial suggests the presence of *doxastic noncotenability*. The use of propositional anaphora—such as “that’s false” or “I disagree with that”—requires a referent that cannot plausibly be an attitude alone, so the most plausible candidate seems to be belief or propositional content. In this regard, standard contextualism cannot secure this form of disagreement.

What about preclusion of joint accuracy—can disagreement over taste, such as the case of *Symphony*, exemplify this form of disagreement? It seems plausible that it can. In many aesthetic disputes, participants do not merely see themselves as expressing divergent attitudes or incompatible beliefs; they take themselves to be *correcting* or even *refuting* one another. If Paul is ultimately persuaded by your assessment of *Bruckner’s 8th Symphony*, he is likely to feel compelled to retract his earlier claim, treating it not just as different but as *inaccurate*.

If this is correct, then disagreement over taste encompasses more than just practical and doxastic noncotenability; it also includes cases where only one party’s assertion can be correct. In such cases, nonindexical contextualism reveals a limitation. While it can accommodate both practical and doxastic noncotenability—by allowing that incompatible

beliefs may each be accurate relative to different standards or contexts—it cannot explain the need to retract a prior assertion once one is persuaded by another’s judgment. That is, it fails to secure the preclusion of joint accuracy. Relativism, by contrast, does. On the relativist view, the accuracy of an assertion is evaluated relative to a single context of assessment (typically that of the assessor), so from any such standpoint, only one of the conflicting claims can be correct.

Relativism thus appears better equipped to explain the full range of disagreement phenomena found in disputes over taste—more so than either standard contextualism or nonindexical contextualism. But does this mean we have arrived at the best solution to the problem of disagreement over PPTs? Not necessarily. In the next chapter, I will argue that MacFarlane has not decisively shown that relativism is superior to its alternatives, and that its apparent advantages may be less secure than they initially appear.

Chapter 3

Assessment of MacFarlane's Relativism

The previous chapter argued that MacFarlane's relativism appears to fare better than its non-objectivist alternatives, primarily because it can (1) account for the phenomenon of retraction and (2) accommodate a wider range of disagreement types observed in disputes over taste. In this chapter, however, I argue that relativism does not offer an ideal account of disagreement, nor does it clearly outperform its competitors.

First, I contend that MacFarlane fails to establish that retraction is obligatory, thereby undermining the claim that relativism holds a clear advantage over other theories on this point. Second, I argue that both nonindexical contextualism and relativism face more fundamental explanatory shortcomings: neither can adequately explain what the motivation of disagreement is when speakers assess the same utterance relative to distinct standards under their semantic framework, nor can they make sense of the coordination of contexts that makes such disagreements meaningful. Without a plausible account of how conversation aims at mutual alignment, these views fall short of capturing the interpersonal dynamics of discourse and the motivational structure behind genuine disagreement.

3.1 Retraction

MacFarlane places considerable weight on the *retraction rule* in his defense of relativism, particularly to distinguish it from nonindexical contextualism. Recall his formulation:

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

This rule shall now be closer scrutinized.

In MacFarlane's view, this rule implies a normative requirement: one *ought* to retract an earlier assertion when one no longer takes it to be true from one's current evaluative perspective. But this raises an important question that MacFarlane does not adequately address: what justifies this normative obligation?

As mentioned in the last chapter, MacFarlane himself thinks that the obligation of retraction is motivated, since it removes the speaker's responsibility to defend the earlier assertion, and it signals to others that they should no longer rely on that assertion as representing the speaker's view. This ties closely to MacFarlane's view on assertion: As MacFarlane (2005, 2009) explained, to assert a proposition is to undertake a commitment to defend the proposition if challenged. More formally:

Assertoric Commitment: In asserting that p at a context c_u , one commits oneself to providing adequate grounds for the truth of p (relative to c_u), in response to any appropriate challenge, or (when appropriate) to deferring this responsibility to another asserter on whose testimony one is relying. One can be released from this commitment only by withdrawing the assertion.⁴⁰

On this view, assertion is a normative act that binds the speaker to defend the asserted content under appropriate conditions. Retraction, then, serves as a mechanism for *releasing* the speaker from that commitment and for *signaling* that the commitment no longer holds.

To illustrate this, we might compare assertions to public signage: imagine a restaurant placing a sign outside that reads: “Free can of soda with sandwich purchase.” Upon seeing this sign, customers reasonably expect the promotion to be honored. The restaurant, in turn, is committed to fulfilling this offer. When the promotion ends, the sign ought to be removed. Leaving it in place, despite the change in circumstances, would mislead passersby into thinking the offer is still valid and the restaurant remains bound by it. In the same way, the assertion of a proposition—such as “Bruckner’s 8th is beautiful”—invites interlocutors to take the speaker as committed to defending that claim. If the speaker no longer endorses it, failing to retract it may mislead others into assuming the commitment still holds.

However, this analogy raises an important conceptual distinction: between *releasing oneself* from a prior assertoric commitment and *informing others* that the commitment has been withdrawn. The former appears to be an internal or private act—something one might do merely by changing one’s mind or standards. The latter, by contrast, is a public act that plays a communicative role in discourse. This raises a key question: Is retraction primarily about the private act of releasing commitment, or the public act of signaling that release to others?

While MacFarlane’s formal definition of assertion—and thus retraction—appears to concern only the (private) act of releasing commitment, I believe there are strong reasons to think that retraction is not just about internal decommitment but about communicative engagement with others.

We typically do not need to notify *ourselves* that we no longer endorse a belief or an earlier assertion—we simply revise our view and move on. Once our evaluative standard has changed, we take the original commitment to be dissolved on its own: we no longer believe the original claim, and we no longer feel bound to defend it. It would be odd to think we have to speak out loud to ourselves that “I take it back” in order to release us from our

⁴⁰ 2005, p. 38

original commitment. Thus in this private sense, the assertoric commitment has already lapsed.⁴¹

But if retraction were only this kind of private revision, it would lack practical conversational utility. Such a withdrawal, left unexpressed, fails to inform interlocutors that the commitment no longer holds. As long as others remain unaware of the change, they may continue to treat the speaker as bound by the earlier assertion and entitled to be challenged for it. In this way, retraction only becomes conversationally relevant when it is publicly performed. To see this clearly, imagine someone once asserted, in conversation, that “Bruckner’s Eighth is not beautiful,” and later underwent a shift in their aesthetic standards. It would be odd to suggest that they could meaningfully retract the earlier claim by quietly saying to themselves, “I take that back.” Such a private act is both unnecessary and ineffective, as it does nothing to alter how others perceive the speaker’s commitments. If the withdrawal is not made known, others will continue to treat the original assertion as representing the speaker’s view.

In this light, retraction matters only insofar as it serves to communicate a change in one’s evaluative standard to others—specifically, by signaling that one no longer endorses a previously asserted claim. However, if we understand the primary function of retraction in this communicative way, then certain problems arise for the idea that retraction is normatively required. I identify two such problems, one terminological, and the other theoretical. If these problems hold, they call into question the necessity of retraction and instead support the view that retraction is merely *permissible*. And if retraction is only permissible, then one of relativism’s chief theoretical advantages over other non-objectivist approaches disappears, as other views could accommodate retraction just as well.

The terminological issue is that, MacFarlane treats retraction as a distinct speech act, akin to saying, “I take that back,” which functions as a kind of undoing of the original assertion. In terms of the analogy discussed earlier, retraction resembles the act of removing a sign that was previously displayed—an explicit withdrawal of a public commitment. However, this narrow conception of retraction is quite demanding: it involves not only signaling a change in one’s evaluative standard but also a specific method to achieve this effect—performing the act of rescinding the original assertion.

⁴¹ Here we should distinguish between believing to oneself, asserting to oneself, and asserting to others. Clearly, if one merely comes to hold a different belief, no act of retraction is needed (see MacFarlane (2014), pp. 104-5). The same holds, I suggest, for so-called “private assertions.” It is unclear why one would need to explicitly retract an assertion made to oneself in order to release oneself from a prior commitment—such commitment, if it ever existed, seems indistinguishable from belief. Since assertion is defined in terms of undertaking a commitment to defend a proposition when appropriately challenged, and we do not typically challenge ourselves in the relevant sense, the very idea of a private assertion is questionable. Even if such challenges happen, it would take the form of an internal dialogue, which more closely resembles evaluating conflicting beliefs and lacks the public structure. It would simply be quite odd to “take back” a private assertion to ourselves. Retraction, then, properly applies only to assertions made to others.

Yet it is unclear why the second element—the symbolic act of “taking it back”—is essential for fulfilling the conversational function MacFarlane attributes to retraction. One can readily agree that the first element—signaling that one no longer endorses an earlier assertion—is important, as it directly serves to release the speaker from the assertoric commitment tied to a claim they no longer uphold. But why must this release take the form of a specific performative act of retraction? If the core function is simply to indicate a shift in evaluative perspective, then other speech acts can serve the same purpose. A speaker might say, for instance, “I used to find Bruckner’s Eighth not beautiful, but I don’t think so anymore—my tastes have changed.” They might even explicitly decline to retract the earlier assertion, saying, “Still, I wasn’t wrong at the time—I simply lacked the refined sensibility to appreciate it.”⁴² These utterances disavow the original commitment and clarify the speaker’s current view without requiring the formal act of “taking it back” in MacFarlane’s narrow sense.

In this way, what I call the broad notion of retraction—merely signaling a change in standard—appears to fulfill the same conversational function as MacFarlane’s narrower notion, which additionally involves a symbolic undoing of the original assertion. If that’s the case, then explicit disavowal of an assertion is not uniquely necessary to dissolve assertoric commitment. And once this broader understanding of retraction is granted, it becomes unclear why other non-objectivist theories couldn’t also accommodate retraction in this sense.

However, this concern is not the central challenge to the necessity of retraction. MacFarlane might well accept that the broader notion is the one that matters. The real difficulty for contextualist and non-indexical contextualist accounts lies elsewhere—specifically, in explaining why retraction is appropriate in cases involving PPTs but not in cases involving indexicals, where retraction would seem misplaced. I will return to this issue shortly. First, however, we turn to a deeper theoretical concern regarding the necessity of retraction itself.

As Ninan (2016) argues, the rationale for the obligation of retraction—namely, that it releases the speaker from a prior assertoric commitment and signals to others that the commitment no longer holds—does not transfer straightforwardly to the relativist framework. This is because it’s unclear that anyone was misled in the first place. Consider an aesthetic example: in context c_u , Anna asserts, “Bruckner’s Eighth Symphony is beautiful,” and this assertion is permitted because the symphony genuinely meets her aesthetic standard at the time. If c_u is a normal context, her audience—say, Ben—understands the conditions under which her assertion is properly made, namely, the reflexive truth rule, and he will assume that Anna is trying to adhere to this rule. In this case, Ben can reasonably conclude that Anna believes Bruckner’s 8th meets her own standard in c , and—assuming Anna is

⁴² Marques 2015, p. 3346

trustworthy—Ben can further infer that Bruckner’s 8th does in fact meet Anna’s standard according to her then-current evaluative outlook in c_u .

Now suppose that by context c' , Anna no longer finds Bruckner’s Eighth beautiful. According to the Relativist Retraction Rule, she is now normatively required to retract her earlier assertion. But why should she be? It’s not that she misled Ben—his belief that *the symphony met Anna’s standard of beauty in c_u* remains perfectly accurate. Nor was the assertion impermissible at the time it was made; it was true and appropriate relative to the context of use. The fact that Anna’s standards have changed does not undermine the correctness of Ben’s belief or the propriety of her earlier assertion.

The upshot of this point is that, under a relativist semantic framework, asserting “Bruckner’s 8th is beautiful” is not straightforwardly analogous to putting up a sign that says “Free can of soda with sandwich purchase,” where the public reasonably assumes the offer remains valid until the sign is removed. The key difference lies in the nature of the content: the message on the sign is not assessment-sensitive, whereas PPTs are. If we were to preserve the analogy while respecting this distinction, the sign would need to read something like “Free can of soda with sandwich purchase (valid on 22nd June 2025),” with the parenthetical is one possible way to specify the context of use and assessment.⁴³ Understood this way, the restaurant is only committed to the truth of the promotional claim *relative to that specific context*, just like an utterer is only committed to defend p’s truth *relative to a context* when uttering p. It becomes unclear, then, why the sign must be taken down after the offer expires, since no one would be misled about its validity beyond the explicitly marked context. Similarly, in the relativist framework, it is unclear why a speaker must retract a properly made assertion just because their tastes have changed.

In light of these considerations, it appears that the *obligatoriness of retraction* is far from securely established. While relativists have often claimed that their framework uniquely predicts the need to retract when standards shift—arguing that the truth of an assertion is relative to a context of assessment—this claim loses traction once we acknowledge that retraction is not always required. The more pressing question, then, is not whether retraction is obligatory, but whether contextualism and nonindexical contextualism can adequately explain its permissibility—particularly given that, as mentioned briefly above, they face a challenge in explaining the asymmetry between retracting indexicals and retracting PPTs.

⁴³ One might object that the analogy is imperfect because our standards of taste tend to remain relatively stable over short periods, and thus they are not immanently shifting in the way temporal contexts are. This could make it seem inappropriate to liken PPTs to context-sensitive signage that varies with time. However, the analogy is not meant to equate the predictability of change between these two cases, but rather to illustrate the structural role that contextual parameters play in determining meaning. In the case of the soda sign, time functions as an implicit contextual parameter—e.g., the offer is valid on this date—which is only specified to make the message context-sensitive. For assertions involving PPTs, the relevant parameter is the speaker’s standard of taste. While it may be difficult to find an exact analogue to an aesthetic standard in the case of signage, the point of the analogy is simply to show how assessment-sensitive content must be indexed to a context in order to be properly understood or evaluated.

To unpack the challenge more closely: It is perfectly natural for a speaker to say, “I was wrong to say Bruckner’s Eighth is not beautiful,” if they later come to appreciate the symphony. This kind of retrospective disavowal feels appropriate in light of changed aesthetic standards. By contrast, it would be deeply odd for someone to say, “I was wrong to say I am eating a sandwich,” when they were indeed eating a sandwich at the time. Yet on the contextualist views, both assertions are truth-conditional in a similar way: either their truths are evaluated relative to the context of utterance (as would claim by the nonidentical contextualists), or their content is fixed by the context of utterance (as would claim by the standard contextualists). In both cases, the original assertions remain true, and retraction would seem unwarranted. So why is one kind of retraction intuitively permissible, while the other is not?

To address the puzzle, we must look beyond semantic content and turn to the pragmatic dimensions of discourse. This is where Knobe and Yalcin’s (2014) experimental findings become especially instructive. In their study, a speaker makes an epistemic modal claim—“Joe might be in Boston”—based on the evidence available at the time. Later, it is revealed that Joe is actually in Berkeley. Participants judged it appropriate for the speaker to retract the earlier claim, even though they did not judge the original statement to have been false. The crucial takeaway here—one not limited to modals—is that retraction can be licensed even when no *semantic* error has occurred.

This result has broader theoretical implications. If retraction is not always motivated by falsity or truth-value failure, then we cannot treat it as a reliable semantic diagnostic. As MacFarlane himself acknowledges, “there is nothing inherently wrong with retracting an assertion one still thinks is true.”⁴⁴ This suggests that retraction is not necessarily a matter of truth. I suggest that retraction can be seen as reflecting a speaker’s shifting commitment to the evaluative standard underlying the original assertion, but not to a change in the semantic features of the proposition asserted. That is, what changes is not the truth-conditions or the content itself, but the speaker’s current stance toward the standard presupposed at the time of utterance. This pragmatic shift in commitment explains why a speaker might disavow a past assertion even while acknowledging it was semantically appropriate in its original context.

This helps explain the familiar asymmetry between retraction of PPTs and indexicals. Take an assertion like “Bruckner’s Eighth is not beautiful,” made at age 15. If the speaker’s aesthetic sensibility changes over time, they may later say, “I was wrong to say that,” not because the earlier assertion was false in context, but because they no longer stand by the evaluative standard that made it appropriate. In this sense, the retraction reflects a change in normative alignment rather than a correction of semantic inaccuracy. By contrast, retraction of indexicals like “I am eating a sandwich” does not seem permissible in the same way, as

⁴⁴ 2014, p. 109

indexicals refer rigidly to features of the utterance context—such as the speaker, time, or place—that do not admit of later reassessment. There is no evaluative standard functioning as a parameter provided by the context which might evolve, and hence no shift in normative stance to motivate retraction.

In light of this, we may plausibly conclude that contextualism can accommodate the data about retraction without needing to posit a relativist semantics. What endows the permissibility of retraction is not necessarily that the semantic content has changed, but can also be that the speaker's commitment to the evaluative framework behind it has. This means that there is no clear advantage of relativism over other non-objectivist theories, as the data point that relativism is committed to explain could also be explained by the alternatives.

That said, a further question arises once we look more closely within the domain of PPTs themselves. For while aesthetic judgments such as "*Bruckner's Eighth is not beautiful*" often invite retraction when standards change, ordinary taste judgments such as "*Fishstick is not tasty*" typically do not. Both are equally retractable in principle, and both are warranted if a speaker's evaluative commitments shift. Yet in practice, we seem to feel a stronger pull to retract the former than the latter. This suggests that, beyond the speaker's changing stance toward an evaluative standard, there is also a dimension of pragmatic pressure on speakers to revise their earlier assertions when their standards evolve. Such pressure appears weak in some cases, like gustatory taste, but stronger in others, like aesthetic discourse.

I will not pursue this line further here. The central point has been to explain why retraction feels natural for PPTs but not for indexicals. But the answer for this data point, I suggest, lies not in semantic differences but in the varying roles that different evaluative terms play in social coordination. The next chapter will return to this point.

3.2 Why Disagree?

In the previous section, I argued that retraction—often cited as a key advantage of relativism—fails to establish relativism's superiority over competing theories. In the remainder of this chapter, I will examine deeper challenges to relativism, focusing first on its treatment of disagreement. Specifically, I argue that under relativist semantics, since disagreement over PPTs does not preclude joint reflexive accuracy, the motivation for disagreement becomes unintelligible.

To start, I introduced the notion of preclusion of joint accuracy in the last chapter. Yet, as MacFarlane notes, relativists are not unique in capturing this notion: objectivists, who adopt absolutist semantics and do not distinguish truth from accuracy, can also say that disagreement involves preclusion of joint accuracy. To further distinguish relativism, MacFarlane proposes a generalized notion:

Preclusion of joint reflexive accuracy: the idea that the accuracy of my belief or assertion (as assessed from my context) precludes the accuracy of yours (as assessed from your context).

This allows relativists to say that two speakers—say, one claiming that the symphony was beautiful, and the other claiming it was not—can each make an accurate assertion relative to their own standards, even while their judgments preclude one another’s accuracy when assessed from the same context. By contrast, for the objectivist, these two notions coincide: if your assertion is accurate (true) from your perspective, it is accurate *simpliciter*, and so must be accurate from mine as well.

The relativist, then, does not claim to vindicate disagreement in all the same senses that the objectivist does. They can admit that disagreement about matters like the beauty of something is less robust than paradigm cases of objective disagreement, since it does not preclude joint reflexive accuracy.

The notion of joint reflexive accuracy supports the kind of faultless disagreement that relativism aims to explain—“faultless” in that no speaker violates the constitutive norms of assertion or belief (i.e., the Reflexive Truth Rule), and “disagreement” in the sense of preclusion of joint accuracy. In disagreements over PPTs, such as “beautiful,” both speakers conform to the reflexive truth rule. Yet their assertions cannot both be accurate when assessed from a single context—yielding a genuine disagreement without fault. This is the type of disagreement MacFarlane argues only relativism can account for. Standard objectivism and contextualism can at best explain *warranted* disagreement, while Humean absolutism and nonindexical contextualism cannot explain *preclusion of joint accuracy*, and thus fall short.

Why revisit faultless disagreement here? Because preclusion of joint reflexive accuracy marks a crucial difference between taste-based disputes (e.g. Symphony) and paradigm objectivist disagreements (e.g. Tom) within the relativist framework:

Symphony:

Anna: Bruckner’s Eighth Symphony is beautiful.

Ben: No, it’s not.

Tom:

Anna: Tom is 7 years old.

Betty: No, he’s not.

For relativism to explain why *Symphony* exhibits an interesting sense of faultlessness (i.e. faultlessness not only understood in terms of warranted) and *Tom* does not, competent speakers must recognize the assessment relativity of terms like “beautiful.” Without this recognition, we would interpret *Symphony* as a simple contradiction, just like *Tom*, and lose the intuition of faultlessness. This means that if relativism is to claim an explanatory edge over competing theories—by accounting for faultless disagreement in a way others cannot—it must assume that competent speakers grasp that we are not simply attributing the absolute notion of truth to assertions involving PPTs, but the relative notion of accuracy. And this, in turn, requires our ability to cognitively track contextual assessment in discourse.

So how does relativism explain our ability to do this? On MacFarlane's view, this recognition is embedded in the very meaning of the predicate. The assessment relativity of expressions of personal taste is not an extralinguistic inference, but a built-in semantic feature. This is evident in MacFarlane's formal lexical entry for "beautiful," which specifies that its extension depends on an aesthetic standard (b) included in the circumstance of evaluation:

$[\text{"is beautiful"}]^c_{\langle w, t, b, a \rangle} = \{x \mid x \text{ is beautiful-according-to-} b \text{ at } w \text{ and } t\}$ ⁴⁵

Here, the truth of "x is beautiful" depends on the world (w), time (t), and most importantly, the beauty standard (b) of the assessor. In contrast, MacFarlane's lexical entry for "poisonous" contains no such assessor-relativity:

$[\text{"is poisonous"}]^c_{\langle w, t, b, a \rangle} = \{x \mid x \text{ is poisonous at } w \text{ and } t\}$

As observed by Clapp (2015), this difference shows that, according to MacFarlane, competent speakers know (probably tacitly) that "beautiful" involves a relative standard, while "poisonous" does not. Thus, our awareness of assessment-relativity in taste discourse arises from our tacit knowledge of the semantic structure of these predicates. This semantic grounding, in turn, underwrites the relativist explanation for why *Symphony* feels faultless and *Tom* does not.

However, this exposes a significant challenge for relativism: if competent speakers fully grasp the relativist semantics of PPTs, they should recognize that utterances involving terms like "beautiful" are true or false only relative to an aesthetic standard. Suppose Anna says, "Bruckner's Eighth Symphony is beautiful," and Paul replies, "No, it's not." If Paul accepts relativism and assumes Anna is being sincere, then he should also acknowledge that her assertion is true relative to her own context of assessment—that is, among other things, her aesthetic standard—even if it is false relative to his own. In other words, he should recognize that joint reflexive accuracy is possible in this exchange.

But if Paul knows this, it becomes unclear what exactly he is trying to do by contradicting Anna. If he understands her statement as accurate-by-her-lights, then his denial—rather than, say, expressing a difference in taste—lacks a clear communicative aim. The force of his denial, which ordinarily implies that Anna has gotten something wrong, appears misplaced if both parties accept that the disagreement merely reflects differing standards and contexts of assessment.

This is the core of the pragmatic challenge: if speakers are aware that predicates like "beautiful" involve relativized truth, they should also understand that their denial, though true by their own lights, is false by the other's. It's therefore conceivable that they could agree that the artwork is beautiful-relative-to-one's standard but not relative-to-the-other's. Recognizing this, they would be rationally compelled to abstain from expressing contradiction in cases like *Symphony*, unless they have some further goal in mind. The fact

⁴⁵ MacFarlane 2014, p. 150

that such disputes arise and could possibly continue for competent speakers poses a serious problem for the relativist attempt to explain the *motivational basis* of disagreement. As MacFarlane himself puts the challenge, relativism makes exchanges such as *Symphony* “look like a pretty silly game. Why do we play it?”⁴⁶

Importantly, the same issue arises for nonindexical contextualism. While it differs in locating truth relative to the context of utterance rather than assessment, it similarly treats predicates of taste as reflecting context-sensitive standards. If speakers are aware of this structure, they too should interpret each other’s claims as true relative to different contexts and refrain from contradiction unless some further practical aim is at stake. Thus, the pragmatics of disagreement remains equally puzzling under nonindexical contextualism.

So what replies could be given by relativists regarding the seemingly lost point of disagreement? MacFarlane’s own response is that such exchanges are “designed to foster controversy,” where “the point of using controversy-inducing assessment-sensitive vocabulary is to foster coordination of contexts. We have an interest in sharing standards of taste, senses of humor, and epistemic states with those around us.” On his view, we have a natural interest in sharing standards of taste with those around us, and engaging in disagreement—even when we know it stems from differing standards—helps promote this kind of interpersonal alignment. Furthermore, MacFarlane suggests that disagreement simply “feels uncomfortable” and that this discomfort itself may be a feature of human nature that motivates us to seek coordination.

I agree with MacFarlane that one important function of disagreements about taste is to bring about some form of interpersonal alignment. However, I do not think pointing this aspect of disagreement out is sufficient to defend the relativist account—what MacFarlane does not explain is *why* uttering sentences that are alleged to be assessment-sensitive as a matter of their semantics fosters such coordination of context. If the semantics of predicates like “beautiful” or “tasty” are genuinely assessment-sensitive, and if competent speakers are tacitly aware of this, then it becomes unclear *why such utterances should provoke controversy at all*.

Clapp (2015) offers a useful analogy involving spatial-orientation predicates such as “on the left” to highlight a challenge for relativism. Suppose John and Mary are facing opposite directions, and John says, “The café is on the left.” If both are competent speakers and understand that such utterances are relative to orientation, Mary would not reasonably reply, “No, it’s not,” simply because the café is not on her left. Instead, she would recognize that John’s assertion is true relative to his orientation. No genuine controversy arises here because both parties are aware of the context-sensitive semantics at play.

The analogy suggests that if aesthetic predicates like “beautiful” function in a similarly assessment-sensitive way, then competent speakers should likewise recognize that

⁴⁶ 2007, p. 29

utterances such as “Bruckner’s Eighth Symphony is beautiful” are true or false only relative to different aesthetic standards. If both speakers understand this, then outright contradiction—saying “No, it’s not”—seems pragmatically unmotivated. The more appropriate conversational response would be to *clarify* the basis of one’s evaluative outlook or attempt to *coordinate* standards, not to deny the other’s claim as though it were straightforwardly false. In this light, assessment-relative semantics would appear to reduce, rather than foster, the motivation to disagree. This presents a tension with MacFarlane’s claim that such semantics are “designed to foster controversy.”

A sharp reader might challenge the force of Clapp’s analogy on two grounds. First, spatial predicates like “on the left” are not technically assessment-sensitive, at least not in the same way that PPTs are. Their relativity depends on a conventional parameter such as physical orientation (location), not a non-conventional parameter like an aesthetic or gustatory standard. Second, if relativism permits analogies with context-sensitive expressions such as indexicals, it begins to look indistinguishable from standard contextualism, which relativism was intended to go beyond—the central promise of relativism is to account for meaningful disagreement in taste discourse—something contextualism allegedly fails to explain.

These concerns are well-taken but ultimately misdirected. The analogy to spatial predicates like “on the left” is not meant to reduce aesthetic discourse to mere indexicality, nor to suggest that predicates of taste operate with the same semantic parameters. Rather, it serves to highlight a pragmatic tension within the relativist account. If competent speakers understand that aesthetic assertions are assessment-sensitive—true or false relative to an individual’s aesthetic standard—they should not feel that contradicting such assertions serves any rational communicative purpose. In other words, just as competent speakers do not feel compelled to contradict spatial claims they know to be orientation-relative, they should not feel the need to challenge aesthetic judgments they recognize as true-relative-to-another’s-standard. And yet, in real-life discourse, they do.

This reveals a deeper shortcoming in the relativist framework. While relativism provides a descriptive apparatus for categorizing taste disagreements—invoking mechanisms such as doxastic non-cotenability or preclusion of joint accuracy—this semantic machinery does not explain what the disagreeing party aims to accomplish. That is, the semantics leaves unanswered why speakers engage in disagreement at all if they accept that both views may be correct relative to different standards. Thus, Clapp’s analogy helpfully illustrates that, if relativist semantics alone determined how we respond to aesthetic disagreement, we would expect less controversy, not more. That we continue to engage in such disputes suggests that the semantics offered by relativism is insufficient to capture the full dynamics at play.

3.3 Coordination of contexts

In this section, building on the previous discussion, I argue that relativism—along with nonindexical contextualism—fails to capture a central function of disagreement: the coordination of evaluative perspectives across speakers. Disagreements, particularly those involving taste, often unfold as reason-giving exchanges through which interlocutors seek to understand, negotiate, and potentially align their standards. Even when full convergence is not achieved, speakers may feel the pull of the other’s perspective, and their evaluative outlooks may shift incrementally. Of course, not all such disagreements aim explicitly at alignment, but it is easy to imagine cases where such a motivation is present—especially when participants are open to engaging one another’s views. This kind of mutual responsiveness is characteristic of genuine communicative exchange. If we accept, then, that one central function of disagreement is precisely this coordination of evaluative outlooks, MacFarlane’s relativist framework proves inadequate.

Notably, this way of framing the matter also revisits what sort of cases should count as genuine disagreement. Much of the literature focuses on minimal examples such as one person asserting “x is beautiful” and another responding “x is not.” If the exchange ends there, however, it functions no differently than “I like Dave” and “I don’t like Dave”: both simply register opposed attitudes without any attempt to explore, resolve or negotiate them. As argued in the previous chapter, mere practical noncotenability as such does not suffice to account for genuine disagreement, and disagreement over PPTs should not be reduced to such exchanges of attitudes. What matters, and what is reflected in real conversational practice where genuine disagreements happen, are cases in which interlocutors treat their opposed judgments as something to be worked through—by exchanging reasons, probing the basis of one another’s outlooks, and seeking (even if unsuccessfully) some measure of alignment. These are the cases where disagreement becomes a vehicle for the coordination of evaluative perspectives, and thus the cases that any account of disagreement must explain.

While MacFarlane speaks of coordinating contexts, the notion of context operative in his semantics is a Kaplanian one: a formal n-tuple including agent, time, place, and standard of taste—parameters which apply *only to the speaker* at the time of utterance. On this understanding, there is no *shared context* for interlocutors to coordinate. Yet genuine disagreement presupposes the possibility of negotiating a standard that both speakers can treat as authoritative, even if only for the purposes of the conversation. This points toward a different conception of context—not as an indexical snapshot of the speaker’s situation, but as a shared space of discourse. I will return to this alternative conception in the next chapter.

Stalnaker (2014) sharpens this point by pointing out that treating disagreements over PPTs as genuine disagreements carries the implication that a shared standard is being invoked (at least the speakers are adopting such pretense)—one that both parties take the speaker’s judgment to be assessed against. Even where preferences diverge, participants often act as if

their judgments ought to conform to some common norm, and that at most one party may succeed in doing so. As he puts it, disagreement arises only if “the context of utterance, or the intentions of the speaker, determine a standard, at some level of generality, that both parties can be said to be aiming to meet.”⁴⁷ On this view, disagreement is not merely a clash of perspectives but a vehicle for gauging divergence and working toward a coordinated evaluative standpoint.

This is precisely the function that the relativist framework, as typically construed, fails to explain. As shown in the previous section, relativism treats an assertion like “*Bruckner’s Eighth is beautiful*” as expressing a proposition whose truth depends on the speaker’s standard at the time of assessment. But because each speaker operates from within their own context, there is no semantic mechanism for treating the exchange as a shared project of alignment. The result is a model where each speaker is effectively sealed off from the other, and where the persistence of disagreement—particularly in cases where no objective facts are at stake—becomes difficult to explain.

The root of this difficulty lies in MacFarlane’s conception of assertion and context. According to his view, to assert is to undertake a commitment to a proposition’s truth relative to one’s own context of assessment. While this aligns with the relativist semantics, it abstracts away from the role assertions play in real-life discourse—namely, their capacity to guide and reshape shared evaluative perspectives. In disagreements that unfold through reason exchange, assertions serve not merely to register private commitments, but to invite uptake, contestation, and possible revision within a public space of reasons. This requires a richer notion of context, where interlocutors can respond to one another’s reasons and revise their outlooks over time. MacFarlane’s individual-centered notion of context lacks the conceptual machinery to accommodate this interpersonal engagement. As a result, it cannot account for the very coordination of perspectives that gives disagreement its conversational point.

Importantly, the difficulty here is not confined to relativism. Any semantic framework that locates the content and assessment of assertions solely within private contexts will inevitably struggle to capture the coordination function of disagreement. The issue is structural: without invoking a shared or negotiable evaluative space, there is no way to account for the pragmatic role disagreement plays in guiding interlocutors toward mutual understanding.

One might attempt to address this shortcoming by modifying relativism to allow for public contexts of assessment. After all, there is no real conceptual barrier to construing the context of assessment as shared rather than purely private. This adjustment might appear to give relativism the tools to account for the kind of mutual evaluative reasoning that real-life disagreement involves. However, once relativism appeals to public contexts to explain such

⁴⁷ 2014, p. 200

coordination, it forfeits its distinctive theoretical appeal. Competing non-objectivist theories can just as easily accommodate such public resources without invoking the machinery of relativist semantics. If a shared context is already doing the work, then other non-objectivist theories could just explain disagreement as a clash over whether something meets a common standard of beauty in the given situation, which leads to relativism's added semantic complexity becoming gratuitous.

This conclusion points toward the need for a new framework—one that captures the interpersonal dynamics of aesthetic discourse without incurring unnecessary complexity. Before turning to such a proposal, I will summarize the key desiderata that have emerged from the preceding discussion:⁴⁸

1. **(BPP).** Our account should explain why speakers with first-hand experience of how something looks, sounds, or tastes are warranted in calling it beautiful just in case its look, sound, or taste is pleasing to them.
2. **Disagreement.** Our account should explain how there can be genuine disagreements (i.e. not merely cases of expressing divergent attitudes) about whether something is beautiful, even when both parties have first-hand knowledge of its features (how it looks, sounds, or tastes) and are aware that those features are pleasing to one party but not the other.
3. **Faultlessness.** Our account should explain how two parties engaged in such genuine disagreements can nonetheless both be faultless—where faultlessness is understood not merely in terms of epistemic warrant but as implying that neither party is making a mistake.
4. **Retraction Conditions.** Our account should explain why speakers are permissible to retract an earlier assertion of beauty if its qualities no longer please their current tastes—even if those same qualities did please their tastes at the time of the original assertion. It should also leave space for the observed asymmetry: some aesthetic judgments invite stronger pressure toward retraction than others.
5. **Coordination of Contexts.** Our account should explain how conversational exchanges about beauty, even in the absence of objective standards, can serve the interpersonal function of coordinating evaluative perspectives—helping interlocutors refine or align their standards over time.

As we have seen, each of the aforementioned views—objectivism, contextualism, nonindexical contextualism, and relativism—succeeds in meeting some of the key desiderata outlined above, but none satisfies all. Objectivist accounts struggle with *Faultlessness* and *BPP*, as they leave little room for accommodating the personal, first-personal character of taste judgments. Standard contextualism faces difficulty with *Disagreement*, since once contextual parameters are fixed to each speaker, apparent conflicts reduce to talking past each

⁴⁸ Some of the following desiderata has been summarized in MacFarlane (2014).

other; it also plausibly struggles with *Coordination of Contexts*, as it lacks a mechanism for explaining how discourse leads to convergence between distinct individual standards. Nonindexical contextualism and relativism handle Disagreement and Faultlessness more naturally, but at the cost of obscuring what disagreement amounts to when utterances are assessed relative to distinct standards. They also lack the resources to account for the *Coordination of Contexts* that gives such disagreements their structure and purpose. Without appealing to a public context, these views cannot explain how or why interlocutors engage with one another's reasons or expect to be persuaded.

Finally, although all of these views can appeal to a pragmatic mechanism to explain why retraction is generally available, they cannot explain why the *pressure* to retract varies across evaluative domains. This asymmetry is a less central issue than disagreement or faultlessness, but it highlights the need for a fuller account of the pragmatic forces shaping aesthetic discourse.

Chapter 4

Common Ground

In this final chapter, I defend a metalinguistic account of disagreement over PPTs that incorporates a public, Stalnakerian notion of context and which is, at its core, contextualist. The central aim is to demonstrate that this approach can account for the full range of data that motivated relativist treatments of predicates of personal taste, while avoiding the complications and radicality often associated with relativist semantics.

The chapter proceeds in several stages. I begin by introducing an alternative to the private, Kaplanian conception of context—namely, Robert Stalnaker’s notion of context as common ground. This conception emphasizes the shared background of presuppositions that underwrites successful communication and provides a natural starting point for understanding how evaluative discourse operates in practice.

Building on this foundation, I examine the presupposition of commonality account of disagreement, most notably defended by López de Sa. On this view, predicates of personal taste presuppose a shared standard, to which speakers appeal when making aesthetic or taste-based assertions. However, as Baker (2012) has convincingly argued, this approach struggles to explain meaningful disagreement in cases where no such standard plausibly exists or being presupposed. If disagreement can occur even when speakers know their standards differ, it seems misleading to treat the conflict as a breakdown of such a presupposition.

To address this problem, I propose that disagreements over PPTs are best understood as metalinguistic negotiations: exchanges in which speakers contest which contextual standard should be treated as authoritative in the ongoing discourse. This metalinguistic perspective gains traction when combined with a Stalnakerian conception of context as common ground, which provides a public framework in which competing standards are motivated to be advanced, defended, or revised.

The resulting picture, I suggest, is both flexible and robust. By combining metalinguistic negotiation with a Stalnakerian common-ground framework, the account captures the nuanced interplay of disagreement, coordination, and faultlessness. In this way, the framework accommodates the full range of phenomena motivating semantic relativism, but without committing to a radically revisionary semantics.

4.1 The Stalnakerian Context

As mentioned in the last chapter, the primary reason for relativism's failure to accommodate the coordination of context is that Macfarlane endorses a "commitment-untaken" notion of assertion and an individual-centered conception of context. Thus, for any theory of disagreement to satisfy this desideratum, the most natural solution is to adopt a different understanding of assertion—one that centers on shared commitments. A plausible starting point, as suggested by Andy Egan (2007)—himself a self-labelled relativist but effectively a nonindexical contextualist according to MacFarlane—is Stalnaker's (1978) account of assertion, which Egan calls the "belief-transfer" model.

On this model, assertion is not merely the expression of a belief, but a proposal to update the conversational common ground. The essential effect of an assertion, as Stalnaker puts it, "is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed" (1978, p. 86). That is, an assertion aims to make its content part of what is jointly taken for granted. In later work, Stalnaker (1998) elaborates that the point of discourse—at least of one central kind—is precisely the exchange and coordination of such information. Participants begin with a background of shared or presupposed-to-be-shared beliefs and presuppositions, and assertions function to modify that background. Crucially, this picture requires a specific understanding of context. Since the meaning of a (context-sensitive) assertion depends not only on its syntactic and semantic features but also on the circumstances in which it is made, context must serve a dual role: both as the interpretive background relative to which assertions are understood and as the informational target that assertions aim to change.

This conception of assertion thus calls for a public, dynamic notion of context—one that tracks the shared commitments and expectations of interlocutors and provides the basis for ongoing negotiation and update. This is what Stalnaker calls the *common ground*. At its core, the common ground is the set of propositions that participants in a conversation mutually accept—or at least mutually presuppose to be accepted.⁴⁹

This stands in contrast to the Kaplanian conception of context, which is static and individualistic. In Kaplan's framework, a context is modeled as a centered possible world: a world paired with a speaker and a time, where these parameters determine what is said by an utterance. While this kind of context contains all facts relevant for interpreting an utterance, it does not distinguish which facts are mutually recognized or relevant for the purposes of communication. As Stalnaker notes, it includes "all that is the case," but fails to isolate what the participants are actually coordinating on.⁵⁰ Accordingly, it lacks the resources to model the evolving structure of shared presuppositions that underpins successful discourse.

⁴⁹ 2002, p. 704

⁵⁰ 2014, p. 24

This makes the Kaplanian context ill-suited for modeling coordination in discourse—especially in cases of disagreement, where successful communication hinges on interlocutors aligning on contextual standards. Notably, as Stalnaker emphasizes, these two notions of context are not competing theories about the same thing but complementary. Their precise relationship is a subtle issue (see Stalnaker 2014, Chapter 1), though a detailed comparison lies beyond the scope of this thesis. My concern here is to show why a public, Stalnakerian model of context is essential for explaining coordination in disagreement over personal taste. I shall now turn to this model in more detail.

To fully grasp this model, I shall focus on clarifying three interrelated notions central to Stalnaker’s account: context, common ground, and speaker presupposition. First, Stalnaker treats context as a body of available information: the common ground. More specifically, Stalnaker (2002) characterizes it as “the mutually recognized shared information in a situation in which an act of trying to communicate takes place.” This notion of context can be formally represented by the context set: the set of possible worlds compatible with what is mutually presupposed in the conversation.

The context set is suited to play both of the roles previously discussed: it provides a shared interpretive background for resolving context-sensitive content, and it represents the evolving space of mutual commitments that assertions aim to shape. In this way, Stalnaker’s formal apparatus captures how discourse progresses—it allows us to represent both what is assumed for communicative purposes and how assertions function to refine or challenge that assumed space, thus content could affect subsequent context.⁵¹

Within this framework, to assert a proposition is to propose eliminating from the context set all those worlds in which the proposition is false. In doing so, the speaker attempts to update the common ground by reducing the set of live possibilities—those mutually recognized as still in play—to only those in which the asserted proposition holds. In Stalnaker’s words, assertions function to shape the common ground, enabling participants to “distinguish among alternative possible ways that things may be.”⁵²

Notably, the notion of common ground is a propositional attitude concept, rather than a factive one. That is, a proposition may belong to the common ground not because it is true, but because it is mutually accepted—either genuinely or by pretense—for the purposes of the conversation. This concept has an inherently iterative structure: a proposition is in the common ground if each participant accepts it, accepts that the others accept it, accepts that the others accept that they accept it, and so on.

While the common ground can be defined in terms of the propositional attitudes of participants, it is not itself an individual attitude. Nonetheless, the individual propositional attitude of speaker presupposition can be defined in terms of the common ground: an agent

⁵¹ Stalnaker 2014, pp. 24-5

⁵² 1978, p. 85

A presupposes that ϕ if and only if A accepts, for the purposes of the conversation, that it is common ground that ϕ .⁵³ Crucially, these presuppositions need not be genuinely believed by anyone, nor do they need to be true; what matters is that they are mutually treated as the shared informational basis of the exchange.⁵⁴ A speaker may presuppose any proposition she finds conversationally useful, so long as she assumes her audience will be willing to go along with it.

Before turning to a theory of disagreement grounded in this model of context, I will first examine an important mechanism for adjusting the common ground—particularly salient in cases of disagreement, where interlocutors recognize a divergence in presuppositions: the mechanism of accommodation.

Accommodation is the process by which participants in a conversation adjust their presuppositions to maintain alignment about the common ground. Since it is reasonably uncontroversial that successful communication depends on shared assumptions about what is taken for granted, any divergence in these assumptions—when made manifest—can prompt cooperative participants to bring their presuppositions back into alignment. This mechanism is grounded in the norm that participants in a conversation should converge on a shared understanding of the common ground.

Accommodation is typically triggered by a manifest event—an occurrence that is publicly evident to all parties involved. Such events may include observable changes in the environment (e.g. a goat walks into the room) as well as speech acts themselves, provided that communication is proceeding successfully. When a speaker utters a sentence, and the semantics of the language are mutually understood, it becomes common ground that this utterance—with its intended meaning—has occurred. However, speech acts can also reveal defects in the conversational context—for example, when one speaker presupposes something the other does not accept. There are various responses one can take after recognizing a defect in the conversational context. Stalnaker (2014) outlines three such responses: (1) accommodate by adjusting one's own belief to align with the presupposition; (2) attempt to correct the other party's presupposition instead of modifying one's own; or (3) continue with the conversation by accepting the presupposed content for pragmatic purposes, without necessarily endorsing its truth. Each of these is a cooperative strategy and plausible in real conversational exchanges.

Given our focus on disagreement, I shall elaborate briefly on the second kind of response—rejection. In this case, rather than accommodating, the hearer makes it manifest that they do not accept the content that the speaker presupposed, thereby prompting the speaker to revise or drop it. Rejecting a presupposition differs from rejecting an assertion: it involves recognizing and challenging what was taken for granted, rather than what was

⁵³ 2014, p. 25

⁵⁴ Stalnaker 1978, p. 84

explicitly stated. Still, there are linguistic resources available for such correction. For example, one may respond with an utterance that manifestly presupposes the opposite of the original presupposition: if Alice refers to Bert's infant son by saying, "How old is she?", Bert might reply, "He is four months," thereby correcting Alice's presupposition regarding the child's gender.

Such acts of rejection offer another perspective on how we might understand disagreement—particularly disagreement in matters of taste. Given the central role that speaker presuppositions play in structuring conversation, it becomes plausible to frame such disagreements as arising from divergent presuppositions about what is, or should be, part of the common ground. Fortunately, de Sá (2008) advances this very line of thought.

4.2 Presupposition of Commonality Approach

De Sá (2008) suggests that in cases like the familiar dispute over the beauty of a symphony, one or both speakers are presupposing that they do not differ in aesthetic standards. More specifically, according to de Sá's approach, utterances such as "Bruckner's Eighth is beautiful" triggers a presupposition of commonality—specifically, that the interlocutors are alike in the respects that matter to judgments of beauty. In any non-defective conversation where one speaker asserts, "Bruckner's Eighth is beautiful," and the other replies, "No, it isn't," the presupposition of commonality is taken to be part of the common ground. As a result, the two utterances are genuinely contradictory, since provided that they are alike, either they are both aesthetically pleased by Bruckner's 8th or they are both not.

De Sá draws on Stalnaker's theory of speaker presupposition and presupposition triggers. On this account, a presupposition reflects the speaker's belief about what is part of the common ground—that is, what is mutually accepted or taken for granted in the conversational context. But what exactly is a presupposition trigger? This notion is central to de Sá's model. According to Stalnaker (2014):

"To say that a certain expression 'triggers' a presupposition ϕ is to say something like this: the data suggest that the expression is normally inappropriate or infelicitous when used in contexts in which the speaker is not presupposing (taking it to be common ground) that ϕ ."⁵⁵

In short, an expression triggers a presupposition when its felicitous use depends on that presupposition being part of the common ground. Applied to aesthetic discourse, this means that statements like "Bruckner's Eighth is beautiful" presuppose shared standards of taste.

Crucially, the presupposition at the heart of de Sá's model is pragmatic in nature—it is not borne by the sentence itself, but rather by the speakers engaged in the conversation.⁵⁶ This is not immediately obvious, given the phrasing that "an utterance triggers a

⁵⁵ p. 7

⁵⁶ Baker 2012, p. 116

presupposition of commonality,” which may suggest a semantic presupposition. But this is misleading. Unlike conventional triggers—such as *again*, *too*, or *stop*—aesthetic predicates do not systematically encode presuppositions in this way. This raises a question: how should we understand the idea that utterances “trigger” presuppositions, if these are pragmatic rather than semantic in nature?

Regarding this question, de Sá himself does not offer an explicit account, but a clarification can be found in Stalnaker’s broader framework. According to Stalnaker, claims about what sentences presuppose could be understood as claims about what cooperative speakers can normally be expected to be presupposing when they use those sentences. A presupposition trigger, on this view, is an expression or construction that signals—whether explicitly or implicitly—that a certain presupposition is in play.

Crucially, Stalnaker emphasizes that claims about presupposition triggers should be seen as describing surface-level linguistic phenomena: they are observations about when speech is felicitous or appropriate, which may have varied explanations. In some cases, the presence of a presupposition trigger might reflect a genuine semantic feature of a lexical item or some sort of construction that certain presuppositions are required. But in other cases, the source of the presupposition trigger may be purely pragmatic, such that the explanation of such triggers might appeal to general conversational principles combined with the truth-conditions of the expressions involved, which by themselves do not mention the notion of presupposition.⁵⁷

In light of this, de Sa’s account of presupposition could be coherently understood as relying mostly on pragmatic presuppositions. At its core, his account is contextualist. He formulates it as follows:

- (a) For each context *c*, “*Bruckner’s Eighth is beautiful*” expresses the proposition that is true (at index *i*) iff, with respect to *i*, the speaker of *c* is pleased by Bruckner’s Eighth.
- (b) The predicate “*is beautiful*” triggers a presupposition that the participants in the conversation are relevantly similar with respect to aesthetic taste.⁵⁸

If successful, this model would offer a straightforward pragmatic mechanism for explaining the desideratum of *coordination of context*—when context is understood as common ground—without the complex relativist semantics employed by standard relativism. However, two challenges arise concerning its theoretical plausibility, both centered on clause (b).

The first is on the detectability of such presupposition of commonality: in many real-world conversations, we simply find no evidence that the participants are presupposing shared taste. If such presuppositions are not operative, the account risks losing its basis

⁵⁷ Stalnaker (2014), pp. 94–95. Note that Stalnaker is originally discussing presupposition requirements in general instead of presupposition triggers. His conclusions in this passage build on arguments developed in earlier chapters. See Stalnaker (2014), Chapters 2 and 3.

⁵⁸ 2008, p. 304

altogether. Baker (2012) presses this point, arguing that predicates of personal taste do not trigger a presupposition of commonality, since the standard diagnostics for presupposition fail to support such a claim. Baker appeals to three widely accepted tests: von Fintel and Yablo’s “Hey, wait a minute” test, their “...and what’s more...” test, and the awkward cancellation test. Each of these is designed to detect whether an utterance carries a presupposition in the ordinary pragmatic sense. For our purposes, it will suffice to elaborate on just one—the “Hey, wait a minute” test—which is perhaps the most intuitive.

The idea, originally formulated by Yablo and von Fintel, is this: if a speaker presupposes a proposition p in asserting a sentence S , then it should be felicitous for a listener who was previously unaware of p to interrupt with, “Hey, wait a minute, I didn’t know that p .”⁵⁹ Consider the example *Goldbach*:

A: The mathematician who proved Goldbach’s Conjecture is a woman.

B: Hey, wait a minute—I didn’t know someone had proved Goldbach’s Conjecture.

Here, A’s utterance presupposes that Goldbach’s Conjecture has been proven, and B’s interjection is entirely appropriate and intelligible in light of that presupposition. Now consider *Dubrovnik*, a case involving an aesthetic predicate:

A: Dubrovnik is beautiful.

B: Hey, wait a minute—I didn’t know we shared the same standards of beauty.

Unlike *Goldbach*, B’s response in *Dubrovnik* strikes us as odd. While there may be specially constructed contexts where such a reply is reasonable—say, if A and B have just discovered a surprising overlap in taste—in many ordinary contexts, A would be well within their rights to reply, “I never claimed we were alike.” This creates a problem for the presuppositional view: If felicitous uses of aesthetic predicates always presupposed commonality of taste, then B’s response in *Dubrovnik* should *always* be a natural and intelligible protest. The fact that it often isn’t suggests that such presuppositions are not in fact carried by the utterance. The upshot is that, if the “Hey, wait a minute” test fails to detect a presupposition of commonality in cases involving predicates of taste, we have strong empirical reason to question whether such a presupposition exists at all.

This concern about the very existence of a presupposition of commonality can also be motivated independently of such diagnostics. Egan (2014), for instance, presses a more theoretical version of the worry on the ad hocness of such presupposition. He pointed out that, in standard cases of presupposition (cases without predicates of personal taste), we can usually point to some feature of the meaning of an expression that explains *why* the presupposition must be in place—why its absence would render the utterance odd, incoherent, or infelicitous. But with predicates of personal taste, no such story has been convincingly articulated. The worry, then, is that the presupposition of commonality is not generated by any identifiable feature of the expression’s meaning, but rather “tacked on” in

⁵⁹ Von Fintel (2004) and Yablo (2006).

order to explain away certain conversational phenomena. This stands in stark contrast to more paradigmatic presupposition triggers, like “the king of France,” “Lisa’s third daughter,” or “the woman who proved Goldbach’s Conjecture”, etc., where the presuppositional content flows directly from the meaning of the expression itself.

Regarding this concern, one way to defend the presupposition of commonality is to point out the pragmatic nature of such presupposition, which may render the ordinary presupposition test inadequate. Baker’s critique rests heavily on the effectiveness of standard presupposition tests, such as the “Hey, wait a minute” test. While such tests work well for *semantic* presuppositions that follow directly from a term’s conventional meaning, their applicability to *pragmatic* presuppositions, especially those which are held by the speakers not because they’re linguistically encoded, is far less clear. In many cases, we can clearly identify intuitive presuppositions that the test fails to register as such.

Consider the following:

A: The mathematician who proved Goldbach’s Conjecture is a woman.

B: Hey, wait a minute—I didn’t know we both speak English.

B’: Hey, wait a minute—I didn’t know we share the same understanding of “mathematician” / “prove” / “Goldbach’s Conjecture” / “woman.”

Here, B’s responses target presuppositions that are undeniably being presupposed by the speaker—namely, that the speakers share a language and a mutual understanding of key terms. Yet the responses sound odd or out of place, possibly because these presuppositions do not stem from semantic triggers in the way that, for example, the presupposition that *somebody proved Goldbach’s Conjecture* arises in the sentence “*The mathematician who proved Goldbach’s Conjecture is a woman.*” Instead, presuppositions like “*we are both speaking the same language*” or “*we are using the terms with the same understanding*” are too fundamental: they are held because they are conditions of the very possibility of communication, not because they are semantically encoded. This suggests that pragmatic presuppositions of this sort resist detection by the standard diagnostic tests.

However, even if Baker’s objection is set aside on these grounds, the presupposition of commonality account faces a further difficulty: genuine disagreements persist even in cases where it is mutually recognized that no shared standard of taste (thereby no shared understanding) is in place. Intuitively, it is possible for individuals to engage in what appears to be genuine disagreement even when they know in advance that their standards diverge. Imagine someone who holds an eccentric or idiosyncratic view that they recognize as being widely rejected. Such a person could still assert that view in conversation and be met with denial, thereby producing what looks and feels like ordinary disagreement, despite no shared standard being presupposed. More familiarly, consider two friends who have repeatedly clashed over a particular topic—for example, whether Bruckner’s Eighth Symphony is beautiful. Each time the subject arises, they disagree, fully aware that their aesthetic

sensibilities differ. These recurring exchanges manifest disagreement, yet it would be implausible to claim that either party presupposes commonality of taste on each occasion.⁶⁰

It is instructive to compare this with a paradigmatic presupposition case. Suppose A and B have long disagreed over whether the universe was created by a divine being or arose spontaneously. Even if A knows B does not think there is a creator of a universe at all, it would not be odd for A to say, “*The creator of the universe is outside time, which explains how they could bring it into existence.*” Here, the expression “*the creator of the universe*” presupposes that there is such a creator, and this presupposition arises from the expression’s meaning itself.⁶¹ In the *Symphony* case, by contrast, any presupposition of commonality depends not on the meaning of the beauty predicate itself, but more on the interlocutors actually believing there is a shared aesthetic standard. When that belief is absent—as in repeated aesthetic disagreements—the presupposition fails to arise at all. This difference underscores the central problem: unlike standard presupposition triggers, PPTs do not generate presuppositions of commonality as part of their meaning.

In light of these two concerns, López de Sa’s account does not succeed in explaining the phenomena of disagreement. Nevertheless, this pioneering attempt to ground disagreement in a presupposition of commonality or common ground points toward a more promising direction. In what follows, I develop an alternative account that characterizes disagreements over PPTs as instances of metalinguistic negotiation. On this view, the relevant presupposition of commonality is not a presupposition triggered by PPTs that interlocutors already share similar aesthetic standards. Rather, it is a pragmatic presupposition held by the speakers that a common use of the term could be achieved as the result of the conversation.

4.3 Metalinguistic Negotiation

Plunkett and Sundell (2013) propose that many disagreements—including those involving PPTs—are best understood as instances of metalinguistic negotiation. On this view, speakers are not simply trading conflicting factual claims, but are instead negotiating how a term ought to be used in a certain context. To put it in terms of common ground: the disagreeing parties are disagreeing about what is the usage of a term that should be registered to the common ground.

To explain how this works, Plunkett and Sundell draw on Kaplan’s (1989) two-tiered semantics for indexicals. Kaplan distinguishes between a term’s character—its

⁶⁰ Egan 2014, p. 82

⁶¹ It might be contentious to claim that the expression “the creator of the universe” carries the presupposition that such a creator exists, just as some dispute whether “the king of France” presupposes the existence of a king of France. Whether this presupposition flows directly from the expression’s meaning is irrelevant here. What matters is that, even setting aside questions about whether the presupposition is meaning-derived, it would still be apt to describe A’s utterance as appropriate, and to note that in saying “the creator of the universe”, A at least presupposes that there is a unique satisfier of the description.

context-invariant meaning encoded by linguistic convention—and its content, which is what the term picks out in a particular context of utterance. Applied to gradable adjectives and PPTs, this framework allows us to say that such terms have stable characters but context-sensitive contents. For example, the character of *tall* is something like “having a maximal degree of height greater than a threshold supplied by context.” What varies is the contextual setting of the threshold.

Consider one of Plunkett and Sundell’s examples. Oscar and Callie are cooking chili. Oscar says, “That chili is spicy,” while Callie replies, “No, it’s not spicy at all.” Both know exactly what the chili tastes like—so the disagreement is not about the chili’s chemical properties or how it registers on their tongues. Rather, they disagree over whether the term “*spicy*” should apply in this situation. This is therefore not a factual disagreement about the world but a metalinguistic one: a negotiation over which contextual standard of “spiciness” should govern their conversation.

Now, if this case exemplifies genuine disagreement—in the sense that the conversation continues beyond the bare exchange of opposed judgments—then it would be misleading to describe Oscar and Callie as merely reporting their private standards. If that were all that was happening, the dispute should simply stop once each had asserted their own standard, since the two propositions would be trivially consistent. But the fact that such disagreements do not stop there indicates something more: in asserting their claims, Oscar and Callie are also pragmatically advocating for one parameter setting rather than another. In other words, they are proposing, and contesting, which threshold of spiciness ought to structure their discourse. Seen this way, their disagreement is substantive: the claim that one contextual standard is preferable to another is precisely the kind of thing over which two speakers can disagree. Extending this to the *Symphony* case, when you and Paul dispute whether Bruckner’s Eighth Symphony is beautiful, the disagreement may likewise be read as a metalinguistic negotiation—over which evaluative standard of beauty the term should reflect in that conversational context.

A natural question arises: Could a disagreement over PPTs sometimes be a dispute about their characters, rather than merely about their contents? Plunkett and Sundell explicitly acknowledge that some disputes—though not typically those involving PPTs—may take this form. One of their most well-known examples, drawn from Peter Ludlow, concerns the predicate “*is an athlete*”. Ludlow recounts a heated sports-radio debate over whether the racehorse Secretariat should appear on a list of the greatest athletes of the twentieth century:

A: Secretariat is an athlete.

B: No, Secretariat is not an athlete.

Here, unlike cases involving context-sensitive adjectives like *spicy*, there is little reason to think “*athlete*” is semantically context-sensitive. Besides, both speakers already know all the

relevant facts about Secretariat's abilities, achievements, and history, just like how Oscar and Callie know the facts of the matter of chili, so it is also implausible to think this dispute concerns factual matter about the subject. What differs is how each speaker applies the term "*athlete*": one includes non-human animals, the other does not. This pattern suggests that they mean different things by the term. The dispute, then, is metalinguistic: a negotiation over which of two competing concepts—call them C_1 and C_2 —should be the character of the expression "*athlete*" in this context.

Although Plunkett and Sundell do not take this to be the typical shape of disagreements involving predicates of personal taste, it is not hard to imagine structurally similar cases in the aesthetic domain. Suppose two philosophers—one (call him A) believes in Kant's distinction between judgments of agreement and judgments of taste, while the other (B) rejects this distinction and believes aesthetic pleasure can be derived from all kinds of objects. One evening, A and B are celebrating at an elegant restaurant, each enjoying a carefully selected, refined glass of wine. Both are experienced wine connoisseurs, equally attuned to the subtleties of flavor, aroma, and texture, so there is little reason to think they are not perceiving the same sensory qualities of the wine. Now consider the following exchange:

B: This is a beautiful glass of wine.

A: No, it isn't.

As in the Secretariat case, the source of disagreement here is not a clash about the wine's properties, nor about contents determined by context. Rather, A and B appear to disagree over the character of the predicate *beautiful*: for B, its character licenses application to objects like wine, whereas for A, it does not. Each party is thus advancing a competing conception of the term's meaning, and the disagreement unfolds as a negotiation over which character should be operative in their conversation.

However, we don't need to think that disagreements over PPTs are about their character. First, it is uncontentious to say cases like this are anomalous. In many cases, we can more simply characterize the disagreement in terms of content, especially with terms like "tasty" or "funny," which don't leave even as much room for negotiation over character as "beautiful" does. These predicates seem more straightforward, and their disagreements are typically about what is the contextual standard for being considered tasty or funny, rather than about the underlying concept of taste or humor. Even in cases of beauty, when we assume that the other party shares a similar aesthetic sensibility, the most natural way to understand the disagreement is as a dispute over which contextual standard of beauty should apply to the conversation.

Moreover, given the explanatory force of metalinguistic negotiation and its compatibility with a contextualist semantics, it is more natural to treat disagreements over PPTs as pragmatic rather than character-level phenomena. On a contextualist view, PPTs

function much like indexicals: just as “I” shifts reference depending on who utters it, so too *beautiful* shifts depending on whose standard is in play. This makes it unnecessary (and indeed somewhat misleading) to describe PPT disagreements as disputes over character, i.e. over the invariant linguistic meaning, rather than as differences in content, i.e. in the contextual standard applied.

This perspective also helps explain why disputes over character should be regarded as exceptional. Even indexicals, though typically unproblematic, can occasionally become the subject of philosophical disagreement at the level of character. Imagine two philosophers at lunch: one says, “I am hungry,” intending it in the ordinary, first-personal sense. The other, drawing on a particular theory of personal identity, objects: “You’re misusing ‘I’—the term can’t pick out the biological organism but only the continuing psychological subject, and that subject is not the one who feels hunger.” Here, the dispute is about how the character of “I” itself should be fixed. Such cases illustrate that even the simplest indexicals can, in rare contexts, be contested at the level of character. But precisely because they are anomalous, they highlight why it would be misleading to model ordinary PPT disagreements on them.

4.4 A Stalnakerian Twist

Above is the basic outline of the metalinguistic negotiation account. Before turning to evaluate its plausibility in light of the five desiderata we proposed, let us first delve deeper into the account and examine how it fits within a Stalnakerian framework. Building on contextualist semantics, it is clear that the contextual standard carried by each speaker’s utterance is semantically anchored in a Kaplanian context, which is fixed individually. Yet if we are to satisfy the desideratum of coordinating context, we must also appeal to a public context. How this coordination actually works requires clarification. Even within a Stalnakerian conception of common ground, when one speaker asserts “*x* is beautiful” and another responds “*x* is not,” what is added to the common ground is only that *x* is beautiful according to A’s standard and not beautiful according to B’s. Thus, while there is mutual acceptance of what each speaker believes, there is no inherent pressure to coordinate their standards so that they converge on the same one. A parallel can be drawn with indexicals: if two speakers assert “I am here” and “I am not here,” what enters the common ground is simply each speaker’s respective claim about their own location, without any need to align the meaning of “I” or “here.” Why, then, do PPTs appear to generate a distinctive coordination pressure? And, given our ideal of treating conversation as coordination of context, what exactly should be registered in the common ground?

Before offering an answer to these questions, it is important to clarify its scope. Not every use of a PPT gives rise to coordination pressure. It would be a mistake to claim that whenever we utter a PPT-statement, we thereby pressure our interlocutor to align with our standard. To avoid that overreach, my focus will be narrower: I will only address cases of

disagreement. More precisely, my question is not why PPTs in general generate coordination pressure, but why they do so specifically in contexts of disagreement.

Further, by “disagreement,” I do not mean those minimal exchanges that terminate with opposed verdicts—cases where one says “x is beautiful” and the other “x is not.” As argued in the previous chapters, such exchanges merely register practical noncotenability and do not yet exhibit the kind of tension we take genuine disagreement to involve. The cases that matter here are those in which the opposed judgments generate a pressure toward resolution—where the conversation can only move forward if one or both parties reconsider their stance, and where reasons are exchanged with that possibility in view. It is this distinctive tension, not present in every context, that constitutes genuine disagreement in the relevant sense.

To sharpen this, let us return to MacFarlane’s classification of structural sources of disagreement: practical noncotenability, doxastic noncotenability, and preclusion of joint accuracy. Since I am working within a contextualist framework, the relativist notion of accuracy does not play a role here. This leaves doxastic noncotenability as the most promising candidate for capturing the structural core of genuine disagreement, for it marks cases where one cannot coherently adopt the other’s belief without abandoning one’s own. But here a difficulty arises: if disagreement simply consisted in doxastic noncotenability, then even the most minimal exchanges—where one says “x is beautiful” and the other says “x is not”—would already count as genuine disagreements. After all, these judgments cannot both be adopted at once. Yet, as we have seen, such bare standoffs rarely exhibit the kind of tension that characterizes substantive disagreement.

This suggests that mere incompatibility of commitments (doxastic noncotenability) does not amount to the lived phenomenon of genuine disagreement. What turns a clash of commitments into genuine disagreement is the presence of coordination pressure—the felt demand on interlocutors to negotiate, defend, or refine their standards in order to keep the exchange going. Only when doxastic noncotenability is coupled with this pressure do we arrive at disagreement in the full sense I want to capture.

The Symphony case makes this clear within a contextualist framework. By asserting opposing verdicts about Bruckner’s Eighth, Paul effectively proposes that “beautiful” should not apply, while you propose that it should. These claims are not semantically contradictory (i.e. they do not express inconsistent contents), but they embody incompatible normative commitments—what Plunkett and Sundell call *rationaly incompatible contents*. Such incompatibility suffices for doxastic noncotenability: neither party can adopt the other’s stance without giving up their own. Yet the exchange does not end there: It also creates a pressure to respond, to argue, to refine one’s stance, and thus to sustain the exchange. It is this transition from mere incompatibility to active negotiation that marks the difference between a bare standoff and a genuine disagreement.

With this clarified, let us now return to the unresolved question of where this pressure comes from. Plunkett and Sundell locate its source in pragmatics: there are many situations where aligning our standards for taste predicates has direct consequences. If you and Paul disagree about whether Bruckner's symphony is beautiful, reaching agreement may affect whether you both decide to buy tickets to next week's concert. The same holds for terms like "funny" or "tasty," where coordination can guide decisions about what movie to watch, where to dine, or how to plan shared activities.

However, this explanation may be somewhat unsatisfying, for these practical-consequence cases may be too contingent that similar coordination needs can even arise for ordinary indexicals. For example, two people disputing "The pub is here" vs. "The pub is not here" might also need to converge on their use of "here" if they are trying to meet and are using the pub as a landmark. If such parallel cases exist, then the difference between indexicals and taste predicates cannot be explained solely by occasional pragmatic needs. And also it seems like in some cases it'd be more useful and informative for us to maintain different standards: for example if you and Paul are deciding whether both of you will enjoy the concert, it is then more useful for you each to report on whether it is beautiful according to your own standards than to coordinate on one of your standards.

So why do disagreements involving PPTs generate coordination pressure? The answer, I suggest, is that in cases of genuine substantive disagreement—where judgments are not only expressed but also contested and negotiated—it is plausible to see the interlocutors as aspiring to improve their standards of taste. The pressure arises not simply from the need to coordinate on immediate practical decisions, but from the recognition that our evaluative standards are revisable and can be made more reflective through mutual scrutiny. If only private standards matter, there would be little reason to argue or to seek convergence: one could simply report one's verdict and stop there. But our practices of disagreement reveal a further aim, namely, to arrive at a standard that carries broader authority, one that others can share and build upon. Convergence matters, then, because it is part of an aspirational practice of shaping a communal standard which is more robust precisely because it has survived challenges and has been refined in light of reasons.

If we accept this, then it becomes clear that while contextualist semantics restricts us to reporting beauty relative to our own standard, disagreement is nonetheless a practice of mutual improvement. Each speaker puts forward her standard as worth adopting, while remaining open to adjustment in light of the other's reasons. In this way, disagreement serves as a method for testing and improving evaluative standards, aiming to make them more reflective and resilient.

Thus, in the case of *Symphony*, when debating whether Bruckner's Eighth Symphony is beautiful, you and Paul are not merely expressing private verdicts. Since each reasonably takes their own way of applying "beautiful" to be worth consideration, but also

recognizes the potential need to justify and possibly adjust it, the point of adducing reasons, then, is not only defensive but aspirational: to move toward a more robust standard by subjecting it to challenges. And crucially, since the two standards are presumably exclusive as they entail different judgments towards a single object, this aspirational process is inherently comparative. Since the two standards are tested against one another in a shared exchange, the outcome is not simply private improvement on each side, but the possibility of convergence on whichever standard—or modified version of it—proves more resilient under scrutiny. This is why the practice generates coordination pressure: through reasoning, speakers are not just refining their own outlooks in isolation but working toward a shared evaluative outlook.

However, how exactly is such coordination achieved in practice? Imagine a case where you succeed in persuading Paul, and the two of you reach convergence. You direct him to a passage midway through the movement and describe what you hear: a recurring motif that is transformed across sections, a sustained delay and then release of harmonic tension, and an increasing textural density that unifies separate lines into a larger formal gesture. These features, you suggest, provide the rationale for your standard of beauty in this context. In doing so, you are presenting a rationale that makes your application of “beautiful” more reflective and potentially compelling to others. Paul, in accepting your view (if he accepts), is not simply capitulating to your taste; he is revising his standard in light of considerations that he now takes to be worth incorporating. The outcome of such exchanges is therefore the gradual shaping of a more robust, mutually scrutinized evaluative framework.

Two routes to convergence are possible. First, Paul may in fact share your standard—valuing structural coherence and textural density—but simply failed to notice those features in the symphony. Once you make them salient, he recognizes them and acknowledges the work as beautiful by his own lights. Alternatively, Paul may begin with a different standard, perhaps prioritizing the general mood the music conveys—whether calm, uplifting, or tragic. Yet, after your explication, he comes to appreciate that structural features also matter for aesthetic evaluation, and adjusts his standard accordingly, much as students do in a music appreciation class. Cases like this are not anomalous. It exemplifies the familiar shape of a successful critical exchange: a dialogue in which each party not only articulates but also reconsiders the grounds of their standard, and thereby understands more clearly why they count something as beautiful.

Now let us recast this treatment of disagreement in Stalnakerian terms. At the outset, speakers update the common ground with parallel reports: “A takes x to be beautiful” and “B takes x not to be.” If the exchange stopped there, the parties would merely be registering opposed attitudes—a perfectly ordinary occurrence, but not yet a genuine disagreement. The disagreement becomes genuine only when the conversation continues: when each party treats the conflict as something that calls for reasons, negotiation, or revision. In that case, the clash is taken to matter for the participants, which opens the space for coordination.

In such cases, the exchange moves beyond parallel reports since each speaker aims to have her standard guide the conversation, while remaining open to being persuaded by the other's reasoning. The goal, in either direction, is to update the common ground with a more plausible communal standard. Successful resolution typically adds two kinds of propositions to the common ground. First, descriptive ones about the work (for example, *p*: "the passage exhibits motif transformation and cumulative textural build"); this may be genuinely new for the other speaker or merely made salient for the purposes of discussion.⁶² Second, a proposal for the conversational standard (*S*): "for the purposes of this exchange, motivic unity and textural density are the criteria that determine beauty." When Paul accepts *p* and *S*, the context set is reduced accordingly: worlds incompatible with *p* or with *S* as the governing evaluative rule are eliminated, and the common ground comes to include the proposition that, for this conversation, *x* counts as beautiful under *S*.

And here the contrast with indexicals becomes evident. With indexicals like "here" or "now," once the reference is fixed, there is no pressure for the conversation to go further. By contrast, with PPTs, once opposed judgments are in play, speakers can treat them not as inert reports but as commitments to be examined and potentially revised. The pressure is not semantic but pragmatic: to engage in genuine disagreement means negotiating which standard should govern the exchange, since otherwise the discussion would amount to no more than the registration of private attitudes.

Two further clarifications are needed. First, it is clear that not every case of disagreement ends in convergence, yet this does not undermine the account I have advanced. Sometimes interlocutors recognize that no agreement can be reached—whether because their standards are too far apart or because further discussion would no longer be productive. In such cases, they may simply agree to disagree. This is not a breakdown of coordination, but a conversational resolution in its own right: the parties acknowledge their divergence, suspend further dispute, and leave the exchange with clearer awareness of their own and each other's standards. Crucially, a lack of resolution does not make the earlier disagreement any less genuine. Even when convergence fails—as it often does—the exchange still serves an important function. By offering reasons and challenging one another's outlooks, speakers highlight different features, reconsider their importance when formulating a judgement, and refine their standards with greater clarity. The outcome is not always agreement, but usually a more reflective and better-coordinated evaluative practice.

Second, not all predicates of personal taste generate the same intensity of coordination pressure. This observation calls for explanation: why do disagreements about beauty or artistic value provoke deeper, more persistent attempts at resolution than disagreements about food or humor? As discussed in Chapter 1, Sundell describes this in

⁶² The descriptive proposition should not be taken to represent an objective 'fact of the matter' about the work. The speaker may only be describing the ways in which the object registers to them, and can therefore be mistaken.

terms of *aesthetic pressure*: some judgments create stronger expectations of coordination than others. Ordinary taste predicates such as “tasty” invite relatively low pressure; paradigmatic aesthetic judgments, especially about art and beauty, exert much higher pressure; while indexicals sit at the zero point on the spectrum, exerting no such pressure at all.

But what explains this variation? One compelling answer comes from Bourdieu’s *Distinction* (1984), where he famously writes, “*taste classifies, and it classifies the classifier.*” On his view, judgments of taste—what we would call the use of PPTs—serve as markers of cultural capital and instruments of social classification. Social agents, themselves classified by their classificatory practices, distinguish themselves by the distinctions they draw between the beautiful and the ugly, the refined and the vulgar. In making an aesthetic judgment, a speaker not only appraises the object but also—often unwittingly—signals her position in cultural space.⁶³ This signal carries classificatory weight, which in turn creates pressure for others to contest or defend their own standards in response.

Bourdieu’s discussion of music makes the point vivid. He remarks that “nothing more clearly affirms one’s ‘class,’ nothing more infallibly classifies, than tastes in music.”⁶⁴ This is partly because the dispositions required to appreciate and participate in certain musical practices—such as attending concerts or playing “noble” instruments—are rare and socially exclusive, making them potent markers of distinction. But it is also because “musical culture” is not simply a stock of knowledge or experiences; it is culturally framed as the most “spiritual” of the arts. To be insensitive to music is perceived, especially in bourgeois contexts, as a form of coarseness. Other aesthetic judgments—whether about literature, painting, or poetry—may exert somewhat less pressure, but they still function as markers of refinement in a way that judgments about food or humor typically do not.

The upshot is this: the intensity of coordination pressure varies with the classificatory significance of the judgment. At the low end of the spectrum, a dispute over whether licorice is tasty rarely bears on one’s social position, so coordination pressure remains minimal. At the high end, disagreements about beauty or artistic worth carry clear social and cultural stakes, which gives interlocutors stronger reasons to negotiate and defend their standards. This explains why, within the broader category of predicates of personal taste, aesthetic judgments generate especially sustained efforts at coordination.

4.5 Five Desiderata

Having outlined the framework, the next task is to examine whether it can meet the five desiderata identified in the previous chapter. In particular, we want to determine whether it can explain all the data points that motivates relativism, but by relying on pragmatic mechanisms in combination with a standard, simple contextualist semantics. In what

⁶³ p. 6

⁶⁴ p. 18

follows, I consider in turn: faultlessness, coordination of context, BPP, disagreement and retraction.

Start with BPP and the closely related cluster of faultlessness and coordination of context. The contextualist semantics anchors the content of a “beautiful” judgment to the speaker’s evaluative standard—that is, the speaker’s pleasure or aesthetic response is constitutive of what they correctly mean when they say something is beautiful. This secures **BPP**: having first-hand access to the object and experiencing it as pleasing gives the speaker the standing to assert “x is beautiful.” At the same time, this semantic anchoring preserves **faultlessness**: because each utterance is assessed relative to the speaker’s standard, two opposing judgments can both be correct in their respective contexts.

Even though faultlessness is preserved, clashes of verdict do not simply cancel out or leave the conversation inert. What gives such exchanges their distinctive character is that aesthetic assertions in contexts of disagreement function pragmatically as proposals for how the predicate should be applied in the present exchange. In disagreement cases, when one speaker asserts “Bruckner’s Eighth is beautiful,” she does not just report her pleasure (and the standard of taste it reflects); she also puts forward that verdict as a salient stance that others may accept, reject, or contest. It is precisely this foregrounding of an evaluative response that supplies a focal point for interlocutors to coordinate their standards around. In this way, the account explains how **coordination of context** is possible without positing objective standards—while still allowing each party to remain semantically faultless.

Turning to **disagreement** more narrowly: the account explains how disagreement can be genuine even when both parties have first-hand knowledge of how the object looks, sounds, or tastes like. Two speakers may engage with the same work yet diverge in their responses. In some cases, the difference stems from how finely attuned their sensibilities are, even if the perceptual data are the same—for example, one listener may discern the unity of motifs in a symphony, while another, less experienced, treats each motif as an isolated moment and so fails to appreciate the larger structure. In other cases, the divergence arises from how much weight they assign to different features, such as which aspects matter most in judging beauty in the first place. In either scenario, their standards come apart. When each puts forward her judgment as the one that should govern the conversation, the clash becomes doxastically non-cotenable: one cannot simultaneously endorse standard A and standard B as authoritative within a single exchange. Thus the account preserves the intuition that such disputes are genuine disagreements.

Finally, retraction falls into place under the same pragmatic mechanism. As argued in the last chapter, retraction can reflect a change in the speaker’s commitment to the evaluative standard that originally licensed the assertion. If the speaker’s tastes change, she may withdraw the earlier proposal—saying “I was wrong to have said x is not beautiful”—not because the original utterance was semantically defective, but because the speaker no longer

endorses the standard that made it appropriate. This contrasts with indexicals, whose referents are rigidly fixed by utterance context and so do not admit of retrospective reassessment. For predicates of personal taste, by contrast, the underlying standards are themselves subject to change, and retraction simply tracks this evolution in normative alignment.

Crucially, on the account I defend, this also explains why retraction pressures vary across different predicates of taste. In domains such as beauty, where evaluative discourse is bound up with coordination pressures—where our claims implicitly invite uptake and alignment from others, and where disagreement carries social stakes of legitimacy and distinction—we are prone to retract when our outlook shifts. By contrast, in cases like finding licorice tasty, where coordination pressure is minimal and there is little expectation that others should share one’s standard, we are not under the same pragmatic demand to retract. Retraction, then, is not a uniform phenomenon across all PPTs: it tracks the extent to which discourse operates under pressures of coordination and social accountability.

Taken together, then, the contextualist framework meets the desiderata. It preserves BPP, explains faultlessness, and accounts for the coordination of context by showing how evaluative assertions function as proposals for conversational standards. The result is a unified account that captures the data without recourse to relativist semantics.

Concluding Remarks

In this thesis I have examined the leading accounts of disagreement over predicates of personal taste—objectivism, contextualism, and relativism. Each offers important insights but proves inadequate in crucial respects: objectivism secures disagreement only at the cost of denying the first-personal warrant of taste (BPP); standard contextualism preserves this warrant but risks reducing disputes to mere talking past one another; nonindexical contextualism and relativism promise both first-personal warrant and arguably faultless disagreement, yet struggle to explain the coordinating function of disagreement. These limitations motivated an alternative that retains a straightforward contextualist semantics while explaining disagreement through pragmatic mechanisms. Drawing on Plunkett and Sundell’s notion of metalinguistic negotiation and embedding it within a Stalnakerian conception of context, I argued that exchanges about taste are not exhausted by the reporting of individual standards, rather, such exchanges are best seen as moments in which speakers advance, contest, and refine evaluative standards with an eye toward uptake into the common ground. On this view, disagreement is a practice through which interlocutors subject their standards to scrutiny and, under pragmatic pressure, aspire toward more reflective and mutually acceptable ones. Importantly, nothing prevents rival semantic frameworks from adopting this same pragmatic explanation; but once its explanatory resources are acknowledged, the appeal of a more revisionary semantics such as nonindexical contextualism or relativism diminishes.

While developed in the context of PPTs, this proposal raises a broader question: can a contextualist–pragmatic account illuminate disagreement across other domains? Many disputes—in ethics, politics, or even descriptive areas of language—share a similar structure: they seem to resist resolution in purely semantic terms yet carry a felt demand for negotiation and potential convergence. If the framework advanced here can shed light on such cases, it would enrich our understanding of the role of evaluative standards in discourse and contribute to a more general picture of disagreement as a practice of coordination through language. Pursuing this question lies beyond the scope of the present work, but I hope to have set out a framework that makes such an extension possible.

Bibliography

- Baker, C. (2012). Indexical contextualism and the challenges from disagreement. *Philosophical Studies*, 157(1), 107-123.
- Baker, C., & Robson, J. (2015). An absolutist theory of faultless disagreement in aesthetics. *Pacific Philosophical Quarterly*, 98(3), 429-448.
- Bourdieu, P. (1984). *Distinction: A social critique of the judgement of taste*. Harvard University Press.
- Clapp, L. (2015). A non-alethic approach to faultless disagreement. *Dialectica*, 69(4), 517-550.
- Egan, A. (2007). Epistemic modals, relativism and assertion. *Philosophical Studies*, 133(1), 1-22.
- Egan, A. (2014). There's something funny about comedy: A case study in faultless disagreement. *Erkenntnis*, 79(S1), 73-100.
- Ginsborg, H. (2020). Aesthetic Normativity and Knowing How to Go On. *Con-Textos Kantianos* 1 (12):52-70.
- Hume, D. (1875). Of the standard of taste.
- Kaplan, D. (1989). Demonstratives An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes From Kaplan*. Oxford University Press.
- Knobe, J., & Yalcin, S. (2014). Epistemic modals and context: Experimental data. *Semantics and Pragmatics*, 7.
- Kolbel, M. (2004). Faultless Disagreement. *Proceedings of the Aristotelian Society*, 104(1), 53-73.
- Kolbel, M. (2008). Introduction: Motivations for Relativism. In M. Kolbel & M. Garcia-Carpintero (Eds.), *Relative Truth*. OUP.
- López de Sa, D. (2008). Presuppositions of commonality: An indexical relativist account of disagreement. *Relative Truth*, 297-310.
- MacFarlane, J. (2005). The Assessment Sensitivity of Knowledge Attributions. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology volume 1*. OUP Oxford.
- MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies* 132 (1):17-31.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166(2), 231-250.
- MacFarlane, J. (2012). Richard on truth and commitment. *Philosophical Studies*, 160: 445–453.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. OUP Oxford.
- Marquez, T. (2018). Retractions. *Synthese*, 195: 3335-3359.
- Ninan, D. (2016). Review of Assessment Sensitivity: Relative Truth and Its Applications. *Philosophical Review*, 125(3), 429-447.

- Plunkett, D., & Sundell, T. (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, 13(23), 1-37.
- Schafer, K. (2011). Faultless disagreement and aesthetic realism. *Philosophy and Phenomenological Research*, 82(2), 265-286.
- Stalnaker, R. (1998). On the Representation of Context. *Journal of Logic, Language, and Information*, 7(1), 3-19.
- Stalnaker, R. (1999). Assertion. In *Context and content: Essays on intentionality in speech and thought* (pp. 78-95). Oxford University Press on Demand.
- Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25(5), 701-721.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stojanovic, I. (2008). The scope and the subtleties of the contextualism–literalism–Relativism debate. *Language and Linguistics Compass*, 2(6), 1171-1188.
- Sundell, T. (2011). Disagreement about taste. *philosophical studies*, 267-288.
- Sundell, T. (2017). Aesthetic Negotiation. In J. O. Young (Ed.), *The Semantics of aesthetic judgements* (pp. 82-105). Oxford University Press.
- Von Fintel, K. (2004). Would you believe it, the king of France is back. In M. Reimer & A. Bezuidenhout (Eds.), *Descriptions and beyond*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Yablo, S. (2006). Non-catastrophic presupposition failure. In J. J. Thomson & A. Byrne (Eds.), *Content and modality* (pp. 164–180). Oxford: Oxford University Press.
- Zangwill, N. (2005). *The Oxford handbook of aesthetics*. J. Levinson (Ed.). Oxford University Press.