

Objective Truth in Matters of Taste

Mihnea Capraru*

Abstract In matters of personal taste, faultless disagreement occurs between people who disagree over what is tasty, fun, etc., in those cases when each of these people seems equally far from the objective truth. Faultless disagreement is often taken as evidence that truth is relative. This article aims to help us avoid the truth-relativist conclusion. The article, however, does not argue directly against relativism; instead, the article defends non-relative truth constructively, aiming to explain faultless disagreement with the resources of semantic contextualism. To this end the article describes and advocates a contextualist solution inspired by supervaluationist truth-value gap approaches. The solution presented here, however, does not require truth value gaps; it preserves both logical bivalence and non-relative truth, even while it acknowledges and explains the possibility of faultless disagreement. The solution is motivated by the correlation between assertions' being true and their being useful. This correlation, furthermore, is used not only to tell which assertions are true, but also to determine which linguistic intuitions are reliable.

The final publication is available at Springer via <http://dx.doi.org/10.1007/s11098-015-0577-z>.

*<http://mihnea.capraru.org>

1 Introduction

Bittersweet chocolate is tasty to some but not to others, and none seem able to justify their preferences with objective reasons. Since the Greek Sophists and perhaps before, many have invoked the relativity of personal taste in order to argue that truth, too, is relative. If it depends on whom you ask whether bittersweet chocolate is tasty—so goes their argument—then it also depends on whom you ask whether *it is true that* bittersweet chocolate is tasty.

During the last decade, philosophers and linguists have put this argument into a rigorous form that has become quite influential, namely, the Faultless Disagreement Argument for Relativism (FDAR). This form of the argument originates with Max Kölbel in 2002, and occurs among others in Peter Lasersohn (2005), Tamina Stephenson (2007), John MacFarlane (2007; 2014), and Andy Egan (2010).

Here is what FDAR consists in. Take the sentence type, ‘Bittersweet chocolate is tasty.’ Some of us are disposed to utter this sentence type approvingly, while others are disposed to utter and approve its negation, i. e., ‘Bittersweet chocolate is not tasty’:

(1) Amelie: Bittersweet chocolate is tasty.

Bernoulli: Bittersweet chocolate is not tasty.

If we wish to save truth from relativity even while acknowledging that personal taste is relative, then it might seem, at first, that we can adopt a very simple form of contextualism. According to this simple contextualist solution, we can explicate the truth conditions of Amelie’s and Bernoulli’s utterances thus:

(2) (SIMPLE CONTEXTUALISM)

Amelie: Bittersweet chocolate is tasty to me.

Bernoulli: Bittersweet chocolate is not tasty to me.

Notice, however, that in dialog (1) it seems intuitively that Amelie and Bernoulli disagree with each other; in dialog (2), on the other hand, they clearly do not disagree, but simply talk about their different taste standards. If there is genuine disagreement in (1), then the simple contextualist solution (2) is mistaken.

The proponents of FDAR maintain, indeed, that the disagreement in (1) is genuine. Moreover, according to FDAR, no matter how we choose a taste standard σ , Amelie's utterance in (1) cannot have the following truth conditions:

Amelie's utterance is true *iff* bittersweet chocolate is tasty according to σ ,
i. e., *iff* Tasty(bittersweet chocolate, σ).

If the above were true for any choice of σ , then Amelie's utterance would either be objectively true or it would be objectively false, depending whether bittersweet chocolate and σ satisfied the two-place relation *Tasty*. This means that either Amelie would be objectively right and Bernoulli objectively wrong, or the other way. Yet in (1) it seems that each of the two is 'faultless,' at least in the minimal sense that neither is any further from the objective truth than the other.

FDAR concludes that the utterances in (1) have the following truth conditions:

Amelie's utterance is true *iff* bittersweet chocolate is tasty (full stop)
Bernoulli's utterance is true *iff* bittersweet chocolate is not tasty (full stop)

There is, presumably, no fact of the matter as to whether bittersweet chocolate is tasty full stop. Hence the above are not simply truth conditions on possible worlds, but rather, they are conditions on pairs of the form \langle possible world, taste standard \rangle .¹ For instance, the pair \langle actual world, Amelie's tastes \rangle verifies Amelie's utterance and falsifies Bernoulli's, while the pair \langle actual world, Bernoulli's tastes \rangle verifies Bernoulli's and falsifies Amelie's.

¹Depending on what is at issue, the truth conditions can also be relativized to further parameters: \langle possible world, taste standard, ... \rangle . In this paper, however, at issue is only personal taste.

So much about *truth conditions*. As for the *propositions* that determine such truth conditions, they are, of course, relative propositions, i. e., propositions that are true or false only relative to one taste standard or another. We may construe such propositions, for instance, as sets of ordered pairs ⟨possible world, taste standard⟩.

As is customary, we will call the view I just described *relativism*. Notice, however, that there is another prominent view often called ‘moderate relativism’ (e. g., Recanati, 2007). As applied to personal taste, moderate relativism does not relativize *utterance* truth to taste standards; instead, moderate relativism relativizes only the truth of *propositions*, while allowing utterances’ truth conditions to be fixed once and for all by the contexts in which the utterances are made. I. e., according to moderate relativism the following are true:

1. In (1), Amelie expresses the relative proposition that bittersweet chocolate is tasty (full stop)
2. In (1), the context of Amelie’s utterance determines a particular taste standard σ .²
3. In (1), Amelie’s *utterance* is true *iff* Amelie expresses a relative proposition that is true relative to the contextually determined standard, i. e., *iff* it is true, relative to σ , that bittersweet chocolate is tasty.

As we see, moderate relativism is *not* the kind of view that we call relativism in this article. Relativism *simpliciter* denies that utterances are either objectively true or objectively false; in MacFarlane’s terms, relativism *simpliciter* relativizes utterance truth to the *context of assessment*, i. e., to whom we ask and to the time and world-state at which we ask them. Moderate relativism, on the other hand, maintains that each utterance is objectively true or objectively false. (This means, among other things, that moderate relativism is vulnerable to refutation by FDAR.) Henceforth I will always refer by ‘relativism’ to relativism *simpliciter*.

²Depending on the version of moderate relativism, σ could coincide with Amelie’s personal preferences, or it could be something more complex.

This article does not embrace FDAR and relativism; neither does the article, however, argue directly against relativism. Instead, the article defends contextualism constructively, aiming to illuminate the general shape as well as the details of a contextualist explanation of faultless disagreement. Hopefully, this will yield not just another shot in the duel of theories, but a better understanding of the practice of communicating about subjective matters. The latter is a surprisingly consequential problem, and even if relativism is eventually abandoned, it will be its merit to have brought the problem to our attention.

* * *

As we work our way toward our contextualist treatment of faultless disagreement, we will give a prominent role to the connection between truth and usefulness. Truth values have a particular kind of strategically important explanatory power. For illustration consider the following two stories.

First story. Fred has an old scar that has become infected. The family doctor consults a specialist, who asserts, ‘The infection is staphylococcal.’ The family doctor therefore prescribes Oxacillin, which is usually effective against Staphylococcus. The doctors are, indeed, right, and one week later Fred has fully recovered.

Second story. Fred has an old scar that has become infected. The family doctor consults a specialist, who asserts, ‘The infection is staphylococcal.’ The family doctor therefore prescribes Oxacillin, which is usually effective against Staphylococcus. In this story, however, the specialist is wrong. Nevertheless, one week later Fred has fully recovered. ‘Why?’, we ask. As it turns out, the infection was streptococcal, and often Streptococcus too will respond to Oxacillin.

Notice the ‘why’-moment during the second story. When we hear the first story, we don’t encounter a similar moment. When the diagnosis is true, we are not surprised that Fred recovered. When the diagnosis is false, however, we need an additional explanation; there has to be more to the story. The false diagnosis won’t do on its own—it will not contribute to the story

the kind of explanation that a true diagnosis would have contributed.

When people hear true assertions and act on their basis, then—all else being equal—we are not surprised to learn that they have successfully achieved their goals. The reason why we are not surprised is because there is a general explanatory connection between true assertions and useful actions, a connection analogous to that between accurate maps and successful sailing or driving. It is because of this general truth–usefulness connection that we don’t need to ask ‘why’ during the first story.

Truth, then, is normally useful. This is so in two senses, one epistemic and the other statistical:

- (3) [epistemic] If an action A is informed by a true assertion T , and if A succeeds in its goals, then T ’s usefulness is *normal* or *non-accidental*, in the sense that T ’s usefulness needs no further explanation other than T being true.
- (4) [statistical] True assertions are normally useful, in that, all else being equal, actions are more likely to succeed in their goals if informed by assertions that are true, than if the same assertions had been false.

Furthermore, true assertions are normally useful not only to their hearers, but also to their speakers; more on this will follow in Section 4.

To be sure, I am not saying that we should expect *all* true assertions to be useful. Usefulness is not a prerequisite for truth. Take, for instance, the assertion of a true theorem belonging to a very arcane branch of mathematics (so arcane that it cannot even help the career of its discoverer). No matter how true, this assertion will probably not be very useful. Notice, however, in (3) and in (4), that we are talking not merely about assertions that are true, but about true assertions that inform actions. Arcane mathematics does not usually inform actions, and that is why we don’t usually expect it to be useful.

Unlike arcane mathematics, personal taste assertions do, very often, inform actions. People take advice in matters of taste such as food, clothing, or haircuts. We are therefore going to use the truth–usefulness connection as a guide to personal taste assertions’ truth values.

Summing up, I will undertake: a) to render utterances of personal taste objectively, non-relatively true (or depending on the case, false), b) while acknowledging that personal tastes are subjective, and c) while preserving the strategic explanatory connection between truth and usefulness.

This suggests that we need a more sophisticated contextualist theory. One such theory is the one advocated by Keith DeRose, namely, the truth-value gap view (2004; 2009). On this view, when the interlocutors have divergent tastes, as they do in (1), then their utterances are neither true nor false, but rather truth-valueless. On the other hand, when the interlocutors have convergent tastes, as in (5), then their utterances are true (or sometimes false).³

(5) Cora: Bittersweet chocolate is tasty.

Dorian (five minutes later): Yes, I just tried it and you’re right, it *is* tasty!

This is clearly an improvement over simple contextualism, as it allows for faultless disagreement: In cases of faultless disagreement, on this view, both interlocutors make truth-valueless utterances, hence both are equally remote from the objective truth. Nevertheless, I find the truth-value gap view problematic. Take (1) again. Here is a contextualist rendering of Amelie and Bernoulli’s disagreement, where σ is the standard determined by (1)’s context:

Amelie: Bittersweet chocolate is tasty by standard σ .

Bernoulli: No, bittersweet chocolate is not tasty by standard σ .

According to the truth-value gap view, the above are neither true nor false. The following, therefore, is also truth-valueless:

³Note that DeRose is writing primarily about knowledge attributions, not about personal taste; he does, however, state his view as a theory about more than just knowledge attributions, and in any case one would be free to adopt it as such.

Tasty(bittersweet chocolate, σ)

If the above is truth-valueless, then there is no fact of the matter as to whether σ says that bittersweet chocolate is tasty. How, though, could this be? At most, what may be is that σ is *undecided* between tasty and non-tasty—to wit, σ doesn't say that bittersweet chocolate is tasty, but it does not say it is non-tasty either. When that happens, then it is *false*, not truth-valueless, to say that σ renders bittersweet chocolate tasty. Hence *at least* Amelie's utterance must have a truth value. It seems therefore that it is best to preserve bivalence, i. e., to keep all taste utterances either true or false.

In Section 3, therefore, I will develop a bivalent contextualist solution, one on which the gap is not in the truth values, but rather in the verdicts of the context-determined standard σ . That is, in (1) σ says neither that bittersweet chocolate is tasty, nor that it is not; accordingly, both Amelie and Bernoulli are making equally false assertions. In (5), on the other hand, Cora and Dorian are making true assertions. I will provide a mathematical model for the view in Sections 5 and 6.

As we see, one reason for a DeRose-style solution is that it allows us to save faultless disagreement. There is, however, a further reason—namely, our solution allows us to do justice to the truth–usefulness connection. I will explain why in Section 4.

Finally, in Section 7 I will survey the field of linguistic intuitions that have been marshaled to support relativism and contextualism respectively, and I will investigate which intuitions we should trust, and which it is safe, under circumstances, to doubt.

We will thus be able to explain faultless disagreement while reconciling objective truth with subjective taste, preserving bivalence, and safeguarding the strategic connection between truth and usefulness.

2 Relativism and Disagreement

In dialog (2) we have seen a very simple contextualist paraphrase of dialogs such as (1). That simple contextualism was rejected, both by the relativists and by ourselves, because it ‘loses disagreement,’ i. e., it has the interlocutors talking past each other instead of genuinely contradicting one another. As we will see, however, it is not only simple contextualism that loses disagreement. So does relativism, *unless* it posits significant semantic naivety (or ‘semantic blindness’); and if it does the latter, then relativism loses most of the force of the FDAR.

This issue is noticed by Matthew McGrath in 2005 (pp. 493–94) and is explored at length by Isidora Stojanovic in 2007. According to relativism, when Amelie and Bernoulli disagree faultlessly, this is because there is a proposition p such that Amelie asserts that p and Bernoulli denies that p , and such that p is true for Amelie and false for Bernoulli:

Amelie: Bittersweet chocolate tastes great!

Bernoulli: No, it doesn’t.

What, though, can this proposition p be? First, it cannot be a proposition that includes either Amelie’s or Bernoulli’s subjective tastes. If p were such a proposition, then the two would no longer disagree faultlessly, but one of them would be wrong and the other right. We might then paraphrase their exchange thus:

Amelie: Bittersweet chocolate tastes great to me! (presumably right)

Bernoulli: No, it doesn’t taste great to you. (presumably wrong)

Alternatively:

Amelie: Bittersweet chocolate tastes great to you! (presumably wrong)

Bernoulli: No, it doesn’t taste great to me. (presumably right)

For the same reason, *p* cannot include *any* taste standard at all. Instead, *p* must be the *relative* proposition that bittersweet chocolate tastes great full stop. To be sure, this is indeed what the relativists have in mind. The problem, though, is that it is not clear how Amelie and Bernoulli can disagree about this relative proposition *unless* they are unaware of its relativity. If Bernoulli thought there was nothing more to *p* being true than *p* being true *for* one subjective judge or another, then he could tell a) that Amelie is right by her own standards, and b) that Amelie is only ‘wrong’ in the sense that she is wrong by Bernoulli’s standards. Then, however, his disagreement would make just as little sense as in this other case:

Amelie: I like bittersweet chocolate!

Bernoulli: No, I don’t.

Hence Amelie and Bernoulli must be semantically naive; they must be unaware that their disagreement concerns merely a relative proposition. Stojanovic concludes that there is no such thing as faultless disagreement at all. This is because, she maintains, “[s]peakers of English are semantically competent with predicates of taste: they master their meaning and truth conditions” (Stojanovic, 2007:696).

I concur with Stojanovic that Amelie and Bernoulli cannot disagree genuinely *unless* they ignore certain aspects of their utterances’ truth conditions. This means that FDAR requires a certain amount of semantic naivety. Then, however, FDAR loses a good deal of its force, because once we have semantic naivety on the table, we—contextualists—can invoke it too. That is just what I will do; in matters of personal taste I find semantic naivety rather plausible, as I will explain in Section 4. Accepting semantic naivety will allow us to pursue a version of contextualism that saves faultless disagreement, in the sense that it saves the faultless interlocutors from talking past each other.

That will be the topic of the next three sections.

3 Finding the Gap

We want a contextualist theory that allows Amelie and Bernoulli not to talk past each other. One such theory is the truth-value gap view (DeRose, 2004, 2009). On this view, Amelie and Bernoulli's assertions are neither true nor false: for the assertions to have truth values, Amelie and Bernoulli would need to share the same taste standards. On page 7 I have explained why I do not find this view entirely plausible as is. We can develop, however, a view in the vicinity of DeRose's, one that preserves bivalence and shifts the gap from the truth values into the taste standards' verdicts.

Here is what I have in mind. Let σ be the taste standard operating in the context of dialog (1). On the truth-value gap view, σ is such that:

- (6) [By standard σ , bittersweet chocolate is tasty] iff bittersweet chocolate is tasty both by Amelie's personal standard and by Bernoulli's;
[By standard σ , bittersweet chocolate is non-tasty] iff bittersweet chocolate is non-tasty both by Amelie's personal standard and by Bernoulli's.

Let us adopt this supervaluationist aspect of the truth-value gap view. Let us also adopt a further aspect of the view, namely, the following trivial consequence of the quotation/disquotation principles:

- (7) (It is true that bittersweet chocolate is tasty by σ) iff (bittersweet chocolate is tasty by σ).

What we will deny, however, is this final aspect of the truth-value gap view:

- (8) (It is false that bittersweet chocolate is tasty by σ) iff (bittersweet chocolate is *non*-tasty by σ).

Instead, we will adopt this equivalence:

- (9) (It is false that bittersweet chocolate is tasty by σ) iff (it is *not* true that bittersweet chocolate is tasty by σ).

Notice that (9) preserves logical bivalence, whereas (8) would sacrifice it.⁴ (9) leaves no gap in the truth values; the gap is now elsewhere, namely, in the verdicts of the standard σ . By this standard, bittersweet chocolate is not tasty, but it is not non-tasty either. Rather, on the matter of bittersweet chocolate, σ is silent:

It is false that by standard σ bittersweet chocolate is tasty.

It is false that by standard σ bittersweet chocolate is non-tasty.

It is true that it is not the case that by standard σ bittersweet chocolate is tasty.

It is true that it is not the case that by standard σ bittersweet chocolate is non-tasty.

Thus both Amelie's and Bernoulli's assertions are false. We can paraphrase their truth conditions as follows:

Amelie: By standard σ , bittersweet chocolate is tasty. (false)

Bernoulli: By standard σ , bittersweet chocolate is non-tasty. (false)

We can now clarify the sense in which Amelie and Bernoulli disagree. Their assertions are not mutually contradictory, but contrary. That is, they can both be false at the same time, but they cannot both be true.

* * *

Finally, let me briefly emphasize that I do *not* propose these truth conditions for *all* taste assertions; I propose them for cases of faultless disagreement. In certain other cases, taste assertions may have different kinds of truth conditions; of particular interest are the assertions

⁴A logically similar but otherwise unrelated approach is suggested by A. N. Prior to deal with future contingents in 1957:94–103.

that Lasersohn calls ‘exocentric,’ whose truth seems to depend on whichever judge the speaker has in mind (2005:668–74):

Mary: How did Bill like the rides?

John: Well, the merry-go-round was fun, but the water slide was a little too scary.

(Lasersohn, 2005:674)

In this case, John says the truth iff the merry-go-round was fun for Bill; it doesn’t matter whether it was fun for the interlocutors, Mary and John.

4 Taste, Truth, and Usefulness

4.1 Reasons, Pulling and Pushing

To support a claim we can give two kinds of reasons, which we will distinguish as ‘pushing’ and ‘pulling.’ A pulling reason is a matter of what the claim can do for us. For instance, if the claim that p helps us explain or predict phenomena, then these are pulling reasons to believe that p . On the other hand, a pushing reason is one that allows us to see why the claim is true. A proof that p , for instance, is a reason that pushes us to accept that p ; so does a scientific prediction that p (as opposed to a confirmed prediction *based* on p , which is not pushing but pulling).

Here is an illustration. If we adopt the law of gravitation because it predicts planetary orbits or projectile trajectories, then these are pulling reasons. On the other hand, if we adopted it because it followed from generalized relativity, which we had already accepted for independent reasons, then we would adopt it for pushing reasons. (Furthermore, if we adopt generalized relativity because it can explain gravitation, then this is a pulling reason.)

Let us now look at the reasons that support the key aspect we have adopted from DeRose’s truth-value gap view, namely (6):

[By standard σ , bittersweet chocolate is tasty] iff bittersweet chocolate is tasty both by Amelie's personal standard and by Bernoulli's;

[By standard σ , bittersweet chocolate is non-tasty] iff bittersweet chocolate is non-tasty both by Amelie's personal standard and by Bernoulli's.

DeRose's reason to embrace (6) is very good. It is, however, a pulling reason. (6) *allows us to say* that Amelie and Bernoulli's assertions are true (or false) in situations when their standards converge. For instance, if Amelie has the same skating rink preferences as Bernoulli, then (6) allows us to say (*mutatis mutandis*) that Amelie's assertion is false:

Amelie: The skating rink will be fun!

Bernoulli: No, it won't be. You only think it will be fun because you think it's regulation size. But you are misinformed.

The above, as we see, concerns what (6) allows us to do, or to say; in other words, the above is a pulling reason. Now there is nothing wrong with pulling reasons, but we can strengthen the case for a thesis if we can add reasons that push, to wit, if we can say not only what the thesis *can do* for us, but also why we ought to believe it *anyway*.

I maintain that we can find such a basis in what I have called the strategic connection between truth and usefulness. Amelie's and Bernoulli's assertions, if true, should promote their goals. Hence we should be able to determine their assertions' truth conditions if we determine the aspects of reality that the assertions (normally) reflect when the assertions are useful.

4.2 Taste and Usefulness

Since faultless disagreement appears to require semantic naivety, we need to look at the conditions under which taste assertions can be semantically naive, but still (normally) useful. As I will

argue, semantically naive taste assertions are normally useful if and only if the interlocutors share similar taste standards.

First, let us imagine a case very similar to John Perry's Z-land (1986). Z-land is a small island whose inhabitants are unaware that there are more lands than Z-land in the world. This is why Z-landers don't treat weather reports as location-dependent. To them, it is not sunny-in-Z-land; it is simply sunny full stop. Picture, now, a community whose members all share exactly the same tastes, hence never come to realize that tastiness can vary with personal standards. Rather, they believe that tastiness is an objective property of certain external objects. One day Zeus says to Zoe:

Zeus: This salad is tasty.

Then Zoe, who trusts Zeus, will eat the salad, and since she shares the same tastes as Zeus and everyone else, she will enjoy it. For Zeus and Zoe and their compatriots, tastes behave just as location does in Z-land, i. e., they make no difference. Even though Zeus and Zoe are unaware that tastiness is at bottom a relation between foods and subjective standards, they can still communicate about tastiness usefully.

Let us move now to a case a bit closer to real life. Suppose that Yeti and Yeta are similar to Zoe and Zeus in that they believe taste properties to be objective. At the same time, however, Yeti and Yeta are like us in that their personal tastes do *not* coincide entirely. Nevertheless, very often their tastes do overlap significantly, hence very often they can communicate usefully despite their naivety. Suppose, for instance, that Yeti and Yeta share mostly the same tastes in flavored yogurt. (Perhaps they only differ on the strawberry flavor.) Then Yeta can usefully say:

Yeta: Don't eat the green yogurt. It's disgusting! I think it's pistachio.

If Yeti follows her advice he will be well served, because he has similar tastes and he will therefore avoid a bad experience.

Thus far we have looked at people who are entirely naive about the semantics of personal taste predicates. Zeus and Zoe, Yeta and Yeti believe that such words express one-place, objective properties of external objects. In this respect they differ from most of us. When in a reflective mood, most of us will either maintain that such words are relative, or perhaps profess uncertainty. Few would maintain without hesitation, in a reflective mood, that certain foods are objectively delicious.

In a less reflective mood, however, even sophisticated speakers can fall into naive beliefs about whether pistachio yogurt is tasty, whether oak trees look better than palm trees, or whether blue is more beautiful than green. This can happen, for instance, when people are upset, intoxicated, or half asleep. Most importantly, it can happen because, as it seems, we *perceive* taste features as if they were one-place, objective properties. For instance, I perceive certain colors (under certain circumstances) as pleasant; that is, not as pleasant-to-me, but simply as just pleasant. Since we perceive taste features as if they were one-place, it is plausible that sometimes we use personal taste predicates with an intention to communicate about objective properties.

Take Xena and Xenon. Neither believes, on reflection, that green is objectively beautiful or that vanilla soy milk is objectively delicious. In the rush of everyday life, however, they often step down from their reflective pedestals and turn into something closer to Yeti and Yeta. They think and talk as if taste properties were objective, even though when reflective they would know better.

Xena and Xenon, I maintain, are just like us.

Suppose, now, that Xena says to Xenon, just as Yeta might say to Yeti:

Xena: Don't eat the green yogurt. It's disgusting! I think it's pistachio.

Even if Xena and Xenon are in a non-reflective, naive disposition, a disposition in which they think yogurt can be objectively disgusting, they can still communicate usefully: all they need is

for their personal tastes to overlap when it comes to yogurt. If Xena's and Xenon's tastes align, then it doesn't matter that they forgot that yogurt cannot be objectively disgusting. They can get away with semantic naivety.

Thus far I have argued that if the interlocutors have similar tastes, then their naive assertions are normally useful. On the other hand, however, if the interlocutors have *dissimilar* tastes, then their naive assertions are normally useless, or even harmful. Recall that Yeti and Yeta differ on strawberry yogurt; let's say, Yeta likes it but Yeti does not. If Yeta says to Yeti, 'The strawberry yogurt is delicious,' then either Yeti believes her—and then her assertion is harmful—or he disbelieves her—and then the assertion, albeit not harmful, is still useless.⁵

To be sure, Yeti can still benefit from Yeta's assertion, if he learns from it that Yeta *believes* that strawberry yogurt is delicious. In her turn, Yeta could then benefit from Yeti's awareness of her tastes. But this is not the kind of usefulness that is relevant to semantics and to truth conditions. Instead, this kind of usefulness is a matter of Gricean, pragmatic reasoning; *all* assertions can be useful this way, no matter whether they concern taste, planets, or triangles.

4.3 Taste and Truth

Semantically naive speakers can make taste assertions just as useful as those made by speakers more sophisticated. Moreover, even when otherwise sophisticated speakers become naive, they can still make useful assertions.

Notice, now, that it is not by accident that these naive assertions are useful. They are useful when the interlocutors have similar tastes where it matters, and this tends to happen not sporadically or by chance, but systematically. For instance, members of the same community often share tastes because of their shared background. Customers of the same restaurant often share tastes simply because the restaurant caters to customers with those tastes. Friends and family

⁵Notice that this concerns *naive* assertions. Sophisticated assertions can easily be useful even when the speakers have dissimilar tastes, as for instance, in restaurants with semantically sophisticated waiters and customers.

have often developed similar tastes through past interaction. Finally, when people exchange advice about food, pop music, or clothing, they often do so because they have already evinced similar tastes in the past. (We tend not to seek advice from sources we think misleading.) In all these ways, people's tastes come to overlap, and therefore when naive taste assertions are useful, they normally don't just *happen* to be so, but rather they can be expected to be.

Recall, now, that as contextualists we must give a well-motivated answer to the question which taste standards are reflected by naive taste assertions in cases of faultless disagreement. We can determine these standards if we determine the circumstances under which such assertions are true. Suppose that a naive taste assertion has truth conditions of the following form:

$\lceil x$ is *predicate-of-taste* by standard $\theta \rceil$ is true iff x is such-and-such

Then θ is the standard that every x meets iff x meets the condition such-and-such. If we can put such-and-such in terms independent of θ , then we can answer our question.

The question about taste standards thus boils down to a question about truth. Truth, in turn, is connected to usefulness; to wit, assertions are true when they reflect reality so as to contribute (non-accidentally) to the interlocutors' goals. Let us ask, therefore, under what circumstances naive taste assertions are normally, non-accidentally useful. Once we figure this out, we can determine which taste standards naive assertions reflect when they are true.

First, notice that naive taste assertions are seldom useful unless they reflect the tastes of their hearers:

(10) Xenon: The vegetable soup is the tastiest.

If Xenon's assertion does not reflect Xena's tastes then she will not benefit from the advice. If he is trying to help her, he will fail; if hoping to ingratiate himself to her, he will fail; if hoping to receive useful advice in return, he will be less likely to succeed; in the worst case, Xena may retaliate.

We can therefore conclude at least this much: Naive taste assertions can only be expected to be useful in a normal fashion if they reflect the tastes of the hearer.

That, however, is not enough. I maintain that such assertions *also* need to reflect the tastes of the speaker.

Notice an important distinction. What I am saying is not merely that useful naive taste assertions *do* reflect the speaker's tastes. They do, of course, but so do useless ones. In general, naive taste assertions reflect the speakers' tastes for the obvious reason that honest speakers say what they think. But it is not enough, for our purposes, to argue that such assertions reflect the speakers' tastes for this obvious reason. Rather, we must argue that it is *because* they reflect the speakers' tastes that the assertions are useful.

This will not be obvious if we only look at (10) in isolation. However, the great majority of assertions are not one-shot affairs. Rather, they are embedded in the back-and-forth of ongoing linguistic interactions. Assertions have effects, some foreseeable and others less so, on the community's store of beliefs—beliefs that include those of the hearer and those of people who the hearer influences. For illustration take this exchange:

(11) Xenon: Spinach is tasty.

Xena: Spinach is a leafy vegetable, so if it's tasty on its own, it must also be tasty with vinaigrette.

Xenon: Great! Then I'll try it with vinaigrette.

Notice how Xenon's first assertion comes back to him—in this example, quite directly; in real life, in many, less direct ways. If speakers' taste assertions did not reflect the speakers' own preferences, then the assertions' consequences would not feed back toward the speakers' goals, and often they would even accrue against them. Normally, therefore, naive taste assertions will be useful only if they reflect hearer's and speaker's tastes alike. Given the truth–usefulness

connection, we arrive, thus, at the same truth conditions as in Section 3.⁶

4.4 Other Uses of Taste Predicates: Persuasion, Coordination, and Community Building

We have followed the truth–usefulness connection as a guide to the truth conditions of taste assertions involved in faultless disagreement. Such taste assertions are made with an intention to communicate information about what tastes good, etc. It is worth noting, though, that taste predicates can also be used in other ways, sometimes also engendering linguistic expressions of disagreement. Such uses, however, do not affect the semantics of taste assertions.

First, people can make taste assertions in order to *influence* or to *persuade*. When I was very young, people used to tell me that spinach is tasty because they hoped to convince me to like it. Does this kind of use affect the truth conditions of taste assertions? I maintain that it doesn't. The persuasive use is not proprietary to taste assertions—instead, it pertains to assertions in general. For instance, I can try to *persuade* (not simply to inform) my friends that it's going to rain by saying to them (perhaps with a certain intonation) that it is going to rain. The persuasive use is not a particular, conventional, semantic feature of taste assertions, but rather a general, pragmatic feature of assertion itself.

Persuasive uses are a special case of a more general phenomenon, namely, attitude coordination. People can make taste assertions in order to coordinate their preferences (Marques and García-Carpintero, 2014; Stephenson, 2007; MacFarlane, 2007). In such cases, the speakers perform at least two speech acts at once. On one hand, they *assert* that something (x) has a certain taste feature (F) (e. g., that spinach is yucky). On the other hand, they *propose* to establish a consensus, among the interlocutors, that x is F .—The second speech act, notice, is not

⁶These truth conditions—let me reiterate—are for the kind of naive taste assertions that occur in cases of faultless disagreement; for sophisticated assertions such as the ones Lasersohn calls exocentric, the truth conditions may look different.

indicative but imperative.⁷— The first speech act—the assertion—has truth conditions that are quite plausibly independent of the imperative speech act, and must be determined on separate grounds. (Indeed, Marques and García-Carpintero are contextualist about the truth conditions of assertions made for coordination, while Stephenson and MacFarlane are relativists).

Finally, taste assertions can be made for what we may call ‘community building.’ People can communicate about their favorite bands, books, or movies, not so much in order to exchange information, but to build mutual sympathy (Egan, 2010:260–61): ‘XYZ is great!—Oh, you like XYZ too?’ Notice that community-building uses are not specific to taste assertions, but can occur with any controversial subjects, such as, for instance, climate change. Some people believe either that climate change is not happening, or that it is not caused by human activity, and they can assert these beliefs in order to build mutual sympathy with others who share their views. The community-building use thus pertains to the general pragmatics of assertion, just as the persuasive use, and does not affect the semantics of taste assertions in particular.

5 A Mathematical Model for the Semantics of Personal Taste

In the previous section we concluded that when naive taste assertions are true, this is because they reflect both interlocutors’ standards at once. If a taste standard is a disposition, then let us call the interlocutors’ joint dispositions their ‘taste neighborhood.’ A taste neighborhood judges as sharply as its component personal standards converge; as the component standards diverge, the neighborhood grows dull and indiscriminating. To make this notion more perspicuous let me describe a simple mathematical model of one particular taste neighborhood.

(12) Amelie: This skating rink is really fun!

⁷Ruth Millikan discusses imperative uses of the indicative mood in Millikan 1984, Chap. 3.

Bernoulli: It's OK, but it's not really fun.

Let us encode Amelie's and Bernoulli's personal taste verdicts as numbers from 1 to 10; let 'really fun' be a 10, 'average' a 5, etc. Then Amelie and Bernoulli's taste neighborhood yields as its verdicts the two-element sets that contain Amelie's and Bernoulli's personal verdicts. For instance, since Amelie gives the skating rink a 10, while Bernoulli, say, only gives it a 6, their neighborhood assigns to the skating rink the verdict-set $\{6, 10\}$.

A taste assertion is true iff its predicate *rules in* the entire set assigned to the subject by the context-determined taste neighborhood. Otherwise, the assertion is false.⁸ We can look at it this way: If the predicate rules in the entire set, then the predicate is coarse enough to render the interlocutors' taste differences irrelevant.

Since Amelie gives the skating rink a 10, Bernoulli a 6, and therefore their neighborhood a $\{6, 10\}$, it follows that:

1. If Amelie or Bernoulli said, 'The skating rink is above average,' that assertion would be true. 'Above average' means within the interval $[5, 10]$,⁹ and that rules in the entire set $\{6, 10\}$. (That is, $\{6, 10\} \subseteq [5, 10]$.)
2. If Amelie or Bernoulli said, 'The skating rink is below average,' that assertion would be false. 'Below average' means $[1, 5[$, and that does not rule in the entire set $\{6, 10\}$: $\{6, 10\} \not\subseteq [1, 5[$. (Indeed, the two do not even overlap.)
3. It is false to say in this context, as Amelie does, 'The skating rink is really fun!' 'Really fun' means 10, and this rules in only one element of $\{6, 10\}$. ($\{6, 10\} \not\subseteq [10, 10]$)
4. It is also false to say, as Bernoulli does, 'It's OK, but it's not really fun.' 'OK' encodes as the singleton interval $[6, 6]$, and this fails to rule in the neighborhoods' set $\{6, 10\}$.

⁸To simplify the exposition, I am assuming that the asserted sentence has the form $\lceil \text{subject is taste-predicate.} \rceil$. The model is easily generalized to arbitrary sentence forms.

⁹Recall that the interval $[a, b]$ is the set of all real numbers x such that $a \leq x \leq b$.

$$(\{6, 10\} \not\subseteq [6, 6])$$

In dialog (12), therefore, *both* Amelie and Bernoulli make false assertions. This is only possible if their assertions are not contradictory but only contrary, i. e., if they can be false at once, but cannot be true at once. That is indeed the case:

Amelie's assertion is true *iff* Amelie and Bernoulli's taste neighborhood rates the skating rink as a subset of $[10, 10]$; otherwise Amelie's assertion is false.

Bernoulli's assertion is true *iff* Amelie and Bernoulli's taste neighborhood rates the skating rink as a subset of $[6, 6]$; otherwise Bernoulli's assertion is false.

6 Comparisons

People agree or disagree not only about what is tasty, but also about what is tastier than what:

(13) Bertrand: Dark chocolate is tastier than milk chocolate.

Amelie: Yes, it is.

Assume, now, that Bernoulli in fact likes both kinds of chocolate, while Amelie doesn't actually like either. Still, each of the two prefers dark chocolate to milk chocolate:

	Amelie	Bernoulli
dark chocolate	4.5	8.5
milk chocolate	1.5	6.5

As we see in the table, Bernoulli's and Amelie's tastes diverge, in the sense that they disagree about both kinds of chocolate. This may suggest that the assertions in (13) are false. Notice, however, that the assertions are not about what is tasty, but rather, about what is tastier than what. On this latter question, Bernoulli and Amelie agree; and given this contextual information, the assertions in (13) are in fact true.

Let us spell this out in terms of our model from Section 5. If Amelie gives dark chocolate a 4.5 and milk chocolate a 1.5, then we will say that Amelie gives the ordered pair $\langle \text{dark chocolate, milk chocolate} \rangle$ the verdict $\langle 4.5, 1.5 \rangle$. Bernoulli, in turn, gives the same chocolate pair the verdict $\langle 8.5, 6.5 \rangle$. Just as in Section 5 then, Amelie and Bernoulli's neighborhood gives to $\langle \text{dark chocolate, milk chocolate} \rangle$ as a verdict the ordered-pair set $\{ \langle 4.5, 1.5 \rangle, \langle 8.5, 6.5 \rangle \}$.

Once again, as in Section 5, Amelie's and Bernoulli's assertions are true iff their taste neighborhood rates the ordered pair $\langle \text{dark chocolate, milk chocolate} \rangle$ as a subset of the (now two-dimensional) region characteristic to the two-place predicate 'tastier,' as in Figure 1 on page 24. In this figure, we represent dark and milk chocolate's scores on the horizontal and the vertical axis respectively. The points $A_{d,m}$ and $B_{d,m}$ represent Amelie's and Bernoulli's verdicts ($A_{d,m} = \langle 4.5, 1.5 \rangle$, $B_{d,m} = \langle 8.5, 6.5 \rangle$). The entire shaded area constitutes what we may call the predicate's 'target zone.' Notice that points in the target zone have a larger horizontal coordinate (dark chocolate) than vertical (milk).

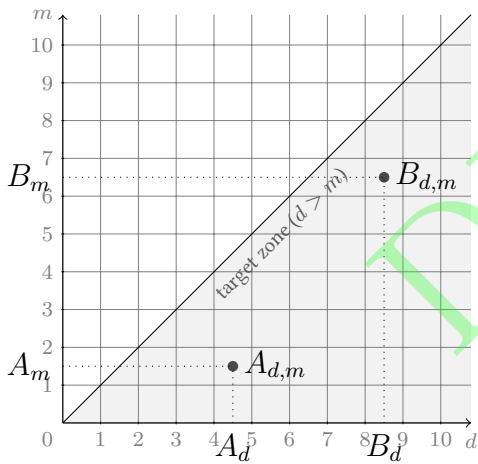


Figure 1: tastier than




Figure 2: disagreement

As we see, the neighborhood verdict $\{ A_{d,m}, B_{d,m} \}$ is contained in the predicate's target zone. Thus our model renders Bernoulli and Amelie's assertions from (13) true.




Figure 3: less tasty than




Figure 4: as tasty as

Other comparisons have similar truth conditions, *mutatis mutandis* (Figure 3 and Figure 4 on page 25). For the predicate ‘less tasty than,’ the target zone is not below but above the angle bisector (i. e., above the line that halves the angle between the axes). For the predicate ‘as tasty as,’ the target zone is the angle bisector itself.

In dialog (13), Amelie and Bernoulli agree. Their assertions therefore reflect both of their preferences, and that is why their assertions are true. But let us also look at case (14), where Cora and Dorian disagree and thus make false assertions:

(14) Cora: Dark chocolate is tastier than milk chocolate.

Dorian: I disagree! *Milk* chocolate is the tastier!

As we see in Figure 2 on page 24, the neighborhoods’ verdict $\{C_{d,m}, D_{d,m}\}$ straddles the angle bisector, neither contained in the target zone for ‘tastier’ as applied to \langle dark chocolate, milk chocolate \rangle , nor as applied to the reverse pair \langle milk chocolate, dark chocolate \rangle . Thus both Cora and Dorian are making false assertions:

Cora’s assertion is true *iff* Cora and Dorian’s taste neighborhood rates \langle dark choco-

late, milk chocolate) as a subset of the area below the angle bisector;¹⁰ otherwise Cora's assertion is false.

Dorian's assertion is true *iff* Cora and Dorian's taste neighborhood rates (dark chocolate, milk chocolate) as a subset of the area above the angle bisector;¹¹ otherwise Dorian's assertion is false.

7 Reliable and Misleading Intuitions

In the literature on relativism, contextualism, and faultless disagreement, intuitions of several kinds are invoked to support one view or another. As John MacFarlane (2011) and Robin McKenna (2014) emphasize, when we choose which intuitions to invoke or to trust, we ought not do so arbitrarily. Both MacFarlane and McKenna consider this a quasi-definitive objection against contextualist views that posit semantic naivety, because these views discard certain intuitions as products of this naivety.¹² One such view is, of course, the one described in this paper. Although I do agree with MacFarlane and McKenna that we cannot pick and choose *arbitrarily* which intuitions to trust, I will argue that we can give good reasons for discarding certain intuitions as semantically naive.

Before I explain, let me sketch a taxonomy of the relevant intuitions.

First, we can distinguish intuitions based on topic; intuitions can concern:

- a) the truth values of assertions,
- b) the assertions' truth conditions,
- c) disagreement, or
- d) retraction.

¹⁰But not below the horizontal axis.

¹¹But not left of the vertical axis.

¹²Note, though, that McKenna is still contextualist, in a different way.

Let me illustrate:

(15) Amanda: The tomato yogurt is cool!

Bernardo: No, it isn't.

Each speaker intuitively feels that his or her own assertion is *true* and the other's false. Moreover, they intuitively feel that the two assertions have incompatible truth *conditions*, and therefore they take the assertions to express *disagreement*. Finally, if Amanda replies to Bernardo:

Amanda: OK, now that I think of it, it's not cool,

then both Amanda and Bernardo will intuitively feel that Amanda has *retracted* what she said before.

Second, intuitions can belong to people differently related to the assertions, namely:

a) speakers, and

b) hearers.

Third, and finally, intuitions can be:

a) intra-context or

b) cross-context.

That is, intuitions can concern either assertions made in one and the same context, or they can concern assertions made in separate contexts. In this article we have looked at cases that involve intra-context intuitions. The literature, however, also gives an important role to cross-context intuitions of truth, disagreement, and retraction, intuitions usually taken as arguments for relativism.¹³

Let us look, now, at speakers' and hearers' intra-context intuitions about truth values.

¹³See, e. g., MacFarlane 2014:8–15.

When people hear assertions, they normally rely on these assertions only if they think the assertions true. In other words, hearers follow their truth-value intuitions as guides to action.

Assume, now, that in situations of a particular type *T*, hearers' truth-value intuitions guide them to *useful* actions. In other words, in *T*-situations truth-value intuitions are a guide to usefulness. Usefulness, however, is a guide to truth. Hence in *T*-situations we can rely on hearers' truth-value intuitions as a guide to the assertions' actual truth values.

In the argument above I am assuming that, *as a rule*, if *A* is a reliable guide to *B* and if *B* is a reliable guide to *C*, then *A* is a reliable guide to *C*. There must, of course, be some exceptions to this rule—namely, when *A* guides us to an unusual subcategory of *B*s that fail to reliably indicate *C*s. In principle, therefore, one could object to the argument above by claiming that taste assertions are one such exception to the rule—that is, by claiming that in matters of taste, usefulness (*B*) is not a guide to truth (*C*). On page 6, however, I have explained why we ought not to treat taste assertions as an exception to the truth–usefulness connection.

Mutatis mutandis, what goes for hearers goes for speakers. Speakers normally make assertions only if they think these assertions true. (There are, of course, exceptions, such as when speakers lie.) Furthermore, I have argued that the truth–usefulness connection applies not only to hearers but also to speakers, i. e., that assertions are normally beneficial to the speakers only when they are true. (Once again, lies are the major exception; however, no one is lying in faultless disagreement cases.) Assume, now, that in situations of type *T*, speakers normally benefit from making sincere assertions. In other words, assume that in *T*-situations, speakers' truth intuitions are a reliable guide to usefulness-to-the-speakers. Given, then, the truth–usefulness connection for speakers, it follows that in *T*-situations speakers' truth intuitions are a reliable guide to the actual truth.

Let us now return to faultless disagreement. Take Amanda and Bernardo, who disagree in (15) over whether tomato yogurt is tasty. Amanda intuits her own assertion to be true and

Bernardo's to be false; Bernardo intuites the other way around. Thus in their quality as speakers, the two intuit their own assertions to be true, while in their quality as hearers, they intuit each other's assertions to be false. Who, then, should we trust—Amanda and Bernardo the speakers, or Amanda and Bernardo the hearers? If we trust the speakers, then both their assertions are true; if we trust the hearers, then both assertions are false.

As I have just argued, we should trust speakers' truth-value intuitions when their intuitions are reliable guides to usefulness for the speakers, and we should trust hearers' truth-value intuitions when their intuitions are reliable guides to usefulness for the hearers. Often this comes down to the same, but in faultless disagreement cases, only *hearers'* intuitions are reliable guides to usefulness. It is useful for Bernardo-the-hearer to find Amanda's assertion false, because this way he avoids eating tomato yogurt, which would displease his palate. It is also useful for Amanda-the-hearer to find Bernardo's reply false, because this way she will stick to her plans and eat tomato yogurt, which she will enjoy. It is *not* useful, however, for Amanda-the-speaker to find her own assertion true (unless something unusual is going on), because this truth intuition leads her to tell Bernardo that tomato yogurt is tasty, and this will either be disbelieved by Bernardo, or lead him astray. For the same reasons, it is useless for Bernardo-the-speaker to find his own assertion true.

Hence in faultless disagreement cases hearers have reliable truth-value intuitions, while speakers do not. Assertions made in such cases are therefore false, because this is what the hearers intuit. This accords with our conclusions in Section 4.

(Notice that I am *not* saying that it is useless for Amanda to believe that tomato yogurt is tasty. Nor am I saying that it is useless for Amanda to intuit that her *belief* is true; what is useless is for Amanda to intuit that her *assertion* toward Bernardo is true. Plausibly enough, the belief and the assertion have different truth conditions.)

So much about truth-value intuitions. As for truth-*conditional* intuitions, I maintain that

we should *generally* trust them, both from speakers and especially from hearers. When hearers accept assertions, the assertions inform not only the hearers' actions, but also their inferences. Inferences, however, are sensitive not just to truth values but also to intensions, and thus to truth conditions. Therefore, in situations where hearers normally make correct inferences, their truth-conditional intuitions are reliable.

At first, what I just conceded looks like an objection to semantic naivety. As long as we trust ordinary speakers' truth-conditional intuitions, we cannot pretend that they are oblivious to certain aspects of the truth conditions. That, however, is just what I maintain in cases of faultless disagreement.

In such cases, however, we have a defeater for the general trustworthiness of truth-conditional intuitions. For the many reasons we have surveyed, when people converse about personal taste, they tend to have similar tastes. When this is so, the interlocutors' semantic mistakes 'cancel out' during inference, the same way as mistakes sometimes cancel out during mathematical derivations. Consider:

Amanda: The lemon frozen yogurt at this place is tasty.

Bernardo: Lemon frozen yogurt is always tastier with mint.

Amanda: Then we should get lemon frozen yogurt and ask for mint on top!

Even if the two are unaware that their assertions' truth conditions involve their taste neighborhood, this will affect neither the validity of their reasoning, nor the usefulness of their actions:

Amanda: [According to our taste neighborhood,] the lemon frozen yogurt at this place is tasty.

Bernardo: [According to our taste neighborhood,] lemon frozen yogurt is always tastier with mint.

Amanda: Then [according to our taste neighborhood,] we should get lemon frozen yogurt and ask for mint on top!

The moral is that we *may* discard truth-conditional intuitions as mistaken, as long as we can explain why these mistakes yield no practical consequences.

All of the above concerns *intra-context* intuitions. What, though, about intuitions across contexts?

Jeremy (on Sunday, to an audience that doesn't contain Karen): Flying kites is fun!

Karen (on Wednesday): What Jeremy said was wrong. Flying kites is boring.

Karen intuits a) that Jeremy's assertion was false, and b) that she disagrees with what Jeremy said. Consider furthermore:

Jeremy (on Thursday): I was wrong. Flying kites is boring.

Jeremy now intuits that he is retracting what he has said on Sunday.

If we wanted to accommodate such cross-context intuitions and still remain contextualists, then we would need to expand assertions' taste neighborhoods to include not only the actual interlocutors, but everyone who could possibly evaluate the assertions' truth values (MacFarlane, 2011). That, though, would be quite absurd: virtually every taste assertion would be rendered false simply because at some possible world someone would disagree with the assertion across contexts. Hence as contextualists we need to *discard* cross-context intuitions.

I maintain, indeed, that cross-context and intra-context intuitions differ greatly in reliability. Intra-context intuitions are central to linguistic communication, hence in general they had better be right. On the other hand, cross-context intuitions are peripheral, plausibly enough a by-product of their intra-context homologues. Since cross-context intuitions are less important for explaining communication, we don't need to accommodate them as much as we need to

accommodate intra-context intuitions. To be sure, this doesn't mean that cross-context intuitions are always unreliable; it does, however, license us to dismiss them when we have specific reasons to do it.

8 Conclusion

I have described and motivated a contextualist solution to the puzzle of faultless disagreement about personal taste. This solution avoids truth relativism, preserves logical bivalence, and still allows us to explain why people arguing about taste seem equally remote from the objective truth.

As a welcome side effect, we have explicated the distinction between what we called 'pulling' and 'pushing' reasons for adopting a theory. We have cast the truth–usefulness connection as a pushing reason for our solution, and we have used the same connection to evaluate the reliability of the linguistic intuitions relevant to the explanation of faultless disagreement.

References

- DeRose, K. (2004). Single Scoreboard Semantics. *Philosophical Studies*, 119:1–21.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context*, volume 1. Oxford University Press.
- Egan, A. (2010). Disputing About Taste. In Feldman, R. and Warfield, T. A., editors, *Disagreement*, pages 247–286. Oxford University Press.
- Kölbel, M. (2002). *Truth Without Objectivity*. Routledge.
- Lasersohn, P. (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy*, 28(6):643–686.

- MacFarlane, J. (2007). Relativism and Disagreement. *Philosophical Studies*, 132:17–31.
- MacFarlane, J. (2011). Epistemic Modals Are Assessment-Sensitive. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*, pages 144–178. Oxford University Press.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press.
- Marques, T. and García-Carpintero, M. (2014). Disagreement About Taste: Commonality Pre-suppositions and Coordination. *Australasian Journal of Philosophy*, 92:701–723.
- McGrath, M. (2005). Review of ‘Truth without Objectivity’ by Max Kölbel. *Philosophy and Phenomenological Research*, 71:491–494.
- McKenna, R. (2014). Shifting Targets and Disagreements. *Australasian Journal of Philosophy*, 92(4):725–742.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. MIT Press.
- Perry, J. (1986). Thought without Representation. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 60:137–151.
- Prior, A. N. (1957). *Time and Modality: Being the John Locke Lectures for 1955–6 delivered in the University of Oxford*. Clarendon Press, Oxford.
- Recanati, F. (2007). *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford University Press.
- Stephenson, T. (2007). Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30:487–525.

Stojanovic, I. (2007). Talking about Taste: Disagreement, Implicit Arguments and Relative Truth. *Linguistics and Philosophy*, 30:691–706.

DRAFT