

Varieties of Disagreement*

John MacFarlane

June 18, 2009

Introduction

The concept of disagreement has been the crux of recent debates between contextualists and relativists about epistemic modals, simple evaluative predicates like “tasty,” and terms of epistemic assessment like “knows.” Let’s focus on a simple example. Sam tastes some Guinness Foreign Export Stout for the first time, loves the flavor, and joyfully declares “That’s tasty.” Sal drinks from the same glass, grimaces, and says, “No, it’s not tasty at all.” The contextualist, impressed by the fact that Sam and Sal both took their assertions to be fully warranted by the (possibly idiosyncratic) effect the drink had on him, concludes that their assertions must concern their own individual reactions, tastes, or standards. The relativist complains that if this is what is going on, Sal’s claim is perfectly compatible with Sam’s, and we lose the sense that the two are *disagreeing*.

Or again: Sally knows that Joe is arriving in Dublin later today, but doesn’t know what airline he is flying. She says, “He might be on the United flight from New York.” Fred replies, “He can’t be on the United flight; they’ve cut back their New York to Dublin route to Wednesdays and Thursdays only.” The contextualist, impressed by the fact that Sally takes herself to be warranted in asserting “Joe might be on the United flight” on the basis of nothing more than her own ignorance of his itinerary, takes her to be claiming, essentially, that what she knows does not rule out Joe’s being on the United flight. The relativist complains that this makes

*This is the lightly edited text of a talk given at University College Dublin and the Institute of Philosophy in London in May 2009.

it impossible to see Fred as disagreeing with her. After all, Fred does not dispute that what Sally knew left it open that Joe might be on the United flight.

These examples are caricatures, but they do give the basic structure of the debate. In each case, the relativist takes there to be a real disagreement between the parties, and argues that the contextualist account cannot explain this.¹ In the face of this challenge, there are two general kinds of response available to the contextualist. The first is to explain the disagreement by finding a less idiosyncratic (but still contextually determined) content that the two parties can be disputing about. In the Guinness case, the contextualist could take Sam to be claiming that the Guinness FES has a taste that will be pleasing to most people, or to the two of them, or to people with the right kind of training. We could then understand Sal to be denying this very proposition. Similarly, in the airline case, the contextualist could take Sally to be claiming that what the *group* (including Fred) knows does not exclude Joe's being on United; Fred could then be understood as rejecting this very proposition.

I have criticized this approach at length elsewhere, so I will not focus on it here. It gives up precisely what seemed most attractive about the contextualist approach: its explanation of why the speakers (Sal and Sally) took themselves to be warranted in making their claims in the first place. It seems to me that (as a first approximation) we take ourselves to be warranted in declaring things tasty whenever we find their tastes pleasing. We do not, at the beginning of an apple tasting class, refrain from making judgements about the tastiness of various apples, even though we may expect that our reactions will mature and develop as the class goes on. Similarly, we do not refrain from making epistemic modal claims when we suspect that someone within earshot might have a relevant bit of knowledge that would rule out the possibility we are putting forward. Simple contextualist theories explain these facts very well. The more we depart from such theories by construing the contents of claims in a less idiosyncratic and more objective way, the less well equipped we are to explain them.

So here I am interested in contextualist responses that retain the idea that the speakers' claims concern their own idiosyncratic tastes or information states. How,

¹For the record, I think it is perfectly possible to argue for relativism without mentioning interpersonal disagreement at all. Instead of considering two people with different tastes, one can consider a single person whose tastes or information has changed over an interval, and ask whether she ought to withdraw an earlier assertion. Perhaps, though, this should be classed as a case of disagreement with one's earlier self. At any rate, disagreement has played such an important part in this debate that it stands in need of closer examination.

keeping that fixed, can the contextualist explain our sense that there is disagreement in these cases? That is the challenge for the contextualist.

There is a related challenge for the relativist. The relativist claims to be able to vindicate our sense that there is real disagreement in a way that the contextualist cannot. However, some critics have urged that relativists about truth are unable to account for disagreement at all. Thus Frege, who favors a contextualist account of taste predicates, writes in his unpublished manuscript “Logic”:

... if something were true only for him who held it to be true, there would be no contradiction between the opinions of different people. So to be consistent, any person holding this view would have no right whatever to contradict the opposite view, he would have to espouse the principle: *non disputandum est*. He would not be able to assert anything at all in the normal sense, and even if his utterances had the form of assertions, they would only have the status of interjections—of expressions of mental states or processes, between which and such states or processes in another person there could be no contradiction. (Frege, 1979, 233)

Does relativism about truth make it possible to understand how there can be disagreements of taste? Or does it make this impossible, as Frege suggests? In order to get clearer about this, we need to ask what disagreement amounts to. We need an account of disagreement that illuminates how it bears on the issues about truth and content that divide contextualists and relativists.

And now we come to the point of this talk. In my earlier efforts to address these issues, I believe I have been asking the wrong question. I have been asking a true/false question when I should have been asking multiple choice. There are several distinct kinds of disagreement, and it is important to look at all of them. If we ask, “What is *real* disagreement?”, instead of “What kinds of disagreement are there?”, our question is unfair to both the contextualist and the relativist. It is unfair to the contextualist because, even if there are kinds of disagreement that contextualist accounts do not capture, there are other kinds that it does capture. And it is unfair to the relativist because it makes it look as if the relativist needs to vindicate the very same kind of disagreement that is secured by objectivist accounts. Even those who are sympathetic to relativism may feel that disagreement about matters of taste is, though genuine, not quite the same kind of thing as disagreement about the age of the earth.

Instead of arguing about what is “real” disagreement, then, I propose to identify several varieties of disagreement. We can then ask, about each dialogue of interest, which of these kinds of disagreement can be found in it, and we can adjudicate between candidate theories of meaning by asking which ones predict that kind of disagreement.

Preliminaries

A couple of preliminary clarifications will be useful, before we begin enumerating varieties of disagreement.

Disagreement as activity and as state

Cappelen and Hawthorne (2009, 60–1) point out that “agree” has both a state and an activity meaning. The same is true of “disagree.” When we characterize two people as disagreeing, we sometimes mean that they are *having a disagreement*—engaging in a kind of activity—and sometimes just that they *disagree*, which is a kind of state.

People can disagree, in the state sense, even if they do not know of each other. In this sense, the ancient Greeks disagreed with the ancient Indians about whether the bodies of the dead should be burned or buried even before Herodotus and other travelers made this disagreement known to them. And I might disagree with nineteenth century Chinese people about the morality of foot-binding. Whether two people disagree is a function of their first-order attitudes, not of their attitudes towards each other.

Whether they are *having a disagreement*, by contrast, depends only on their attitudes and actions towards each other. Two people who agree about all the issues at stake could nonetheless be having a disagreement if, through some misunderstanding, they take their views to differ, or if one is playing devil’s advocate. The question “Why are you disagreeing with me, if we agree about what is at issue?” is perfectly intelligible.

It seems to me that the state senses of disagreement are fundamental, and that any account of the activity sense will have to make reference to a state sense. Having a disagreement requires taking oneself to disagree. In what follows, then, I will only be discussing the state sense.

A note on logical grammar

Disagreement in the state sense is a kind of relation: one can disagree with Sam, but not with Sal, about Guinness FES. One might, then, take the notion we are explicating to be

x disagrees with y .

But this is not sufficiently discriminating. Nobody agrees with anybody about *everything*, so this is a relation everyone will stand in to everyone else. We need a way of saying that Sam and Sal agree in these respects, while disagreeing in those. So we might take our target to be the relation

x disagrees with y about whether p .

But this target is not going to work with all of the varieties of disagreement we will be considering. Some kinds of disagreement involve attitudes without propositional content. In other cases, whether there is a disagreement depends not just on the contents of the relevant attitudes, but on the contexts in which they occur. So a more general target is

x disagrees with y 's ϕ ing in context c .

where ϕ can be replaced by a verb phrase describing an attitude—for example, *believe that Mary is smart*, or *hate the taste of grape jelly*.

We will consider some different ways of explicating this relation; all of them, I think, are genuine kinds of disagreement. I do not think we lose anything by focusing on relations between a person and an attitude or speech act in context. Given this, we can presumably recover a relation between a person and a person, and between a person, a person, and a proposition, in the special cases where this is appropriate.

Non-cotenability

In one sense, I disagree with someone's attitude if I could not coherently adopt that same attitude—an attitude with the same content and force—without changing my mind—that is, without dropping some of my current attitudes. In other words, I disagree with attitudes that are not *cotenable* with my current attitudes.

Many paradigm cases of disagreement are cases of non-cotenability (in addition to being disagreements in other senses). If George believes that all bankers are rich and that McGovern is a banker, and Sally believes that McGovern is poor, then Sally's belief is not cotenable with George's attitudes, since if George came to believe that McGovern is poor while still holding his other attitudes, his beliefs would be logically incoherent.

Asked what disagreement is, I suspect many philosophers' first answer will be what we might call the Simple View: to disagree with someone's belief that p is to have beliefs with whose contents are incompatible with p .² The notion of disagreement captured by the Simple View can be seen as a special case of non-cotenability, where the attitudes are limited to attitudes of full belief. But non-cotenability yields interesting notions of disagreement when applied to other kinds of attitudes as well.

Ned, the weather reporter for Channel 4, has a credence of 0.7 that it will rain tomorrow. Ted, the weather reporter for Channel 5, has a credence of 0.8 that it will rain. Ned could not adopt Ted's attitude without change of mind, so we have a case of non-cotenability, even though both Ned and Ted take it to be pretty likely that it will rain. This is a kind of disagreement, though it is not the first thing one thinks of when one thinks of disagreement. The disagreement between the atheist and the agnostic is also of this kind.

Torfinn Huvenes has called my attention to dialogues like the following:

Bob: The hypothesis is false.

Carol: I disagree, we need to do further testing.³

²Of course, if contents are individuated coarsely, more must be said, since we might not intuitively want to say that Hammurabi disagrees with Sannurabi when the former thinks that *Hesperus* is visible and the latter thinks that *Phosphorus* is not visible. For present purposes, assume a conception of contents, and of compatibility, on which *Hesperus is visible* is compatible with *Phosphorus is not visible*.

³"Varieties of Disagreement and Predicates of Taste"

This is also a case of non-cotenability. In asserting that the hypothesis is false, Bob has expressed his confidence that it is false. This confidence is not cotenable with Carol's attitudes, so in that sense she disagrees.

We can also have non-cotenability of nondoxastic attitudes, like desires, likings, or preferences. Suppose that Jane likes Bob, but Sarah hates him. In a perfectly respectable sense, Jane disagrees with Sarah, even if she believes all the same things about Bob. She does not disagree with Sarah about whether p , for any p , but she disagrees with Sarah about Bob, since Sarah's attitude towards Bob is not cotenable with hers. In this case, the incoherence that would result if she adopted it would not be inconsistency, but a kind of practical incoherence: the incoherence one suffers when one likes and hates the same thing.

In the same sense, two kids might disagree about licorice, one wanting to eat it, the other being repulsed by it. There need not be any proposition they differ about for them to disagree about licorice. It's enough if they just have different attitudes towards licorice.

One mark of disagreement is the capacity to generate disputes, and as C. L. Stevenson points out, even practical non-cotenability of the sort we have been discussing can do that (Stevenson, 1963, ch. 1). The *point* of these arguments may be to change the other person's attitude—for example, to bring Sarah around to liking or at least tolerating Bob—but that doesn't mean that the argument must consist solely of expressing attitudes. To the contrary, as Stevenson emphasizes, changing others' attitudes often involves calling their attention to various sorts of facts. Jane might point out that Bob *did* help build a house for homeless people, and so is presumably not completely self-centered.

So, non-cotenability is a genuine kind of disagreement. As we will see, however, it is not the only kind of disagreement we can make sense of. And it is not the kind of disagreement the relativist should focus on in distinguishing his position from contextualism.⁴

⁴Max Kölbel defines disagreement on behalf of relativists in terms of doxastic non-cotenability (Kölbel, 2004, 305). He says that two parties disagree if one could not rationally accept what the other says without changing her mind. If we understand "accept what the other says" as "come to believe what the other says," as seems natural, then this amounts to something like doxastic non-cotenability. I argue below that there are certain contextualist positions that can secure disagreement in this sense, but fall short of securing the more robust kind of disagreement the relativist aims to capture.

Preclusion of joint satisfaction

Actually, reading Stevenson's account of "disagreement in attitude," it is not entirely clear to me whether he means practical non-cotenability or something a bit different. He first introduces the notion this way:

This occurs when Mr. A has a favorable attitude to something, when Mr. B has an unfavorable or less favorable attitude to it, and when neither is content to let the other's attitude remain unchanged. (Stevenson, 1963, 1)

That sounds exactly like practical non-cotenability. On the next page, though, Stevenson distinguishes disagreement in belief from disagreement in attitude as follows:

The difference between the two senses of "disagreement" is essentially this: the first involves an opposition of beliefs, both of which cannot be true, and the second involves an opposition of attitudes, both of which cannot be satisfied. (2)

Whether two attitudes are cotenable depends only on their forces and their contents. But whether they can both be satisfied depends also on the contexts in which they occur (for example, on who has them and when). As a result, preclusion of joint satisfaction and non-cotenability can come apart.

Here is an example. There is a cupcake on the table. Alvin and Melvin both want to eat it. They both have a desire with the content *to eat that cupcake*. Their desires are the same in force and content, hence cotenable. Yet clearly they cannot be jointly satisfied; the cupcake can only be eaten by one of them.

Meg and Peg are also looking at the cupcake. Meg desires to eat the frosting only. Peg desires to eat the cake part only. Their desires have different contents and are not cotenable. (Desiring to eat the frosting only and to eat the cake part only is practically incoherent.) However, it is perfectly easy for both desires to be satisfied.

I have assumed here a certain view about the content of desires. Desires are usually attributed with infinitival complements: one desires to ϕ , for some ϕ . I take

it, then, that the content of a desire is the kind of thing that is expressed by such a complement: presumably, a property, or perhaps a centered proposition (which has truth values relative to a world, a time, and an agent as “center”).

The counterexamples would not go through if we said that the content of Alvin’s desire is *that Alvin eat the cupcake* and the content of Melvin’s is *that Melvin eat the cupcake*, because now we would have two cotenable attitudes. So if one insisted that the contents of all desires are (uncentered) propositions, the distinction between practical non-cotenable and preclusion of joint satisfaction would become purely notional.

This is not the place to settle a dispute about the contents of desires, but I want to make a methodological point. Our aim here is to get clear about the varieties of disagreement. Philosophers who take the contents of desires to be uncentered propositions should not object to distinguishing between practical non-cotenable and preclusion of joint satisfaction; even if they take this distinction to be purely notional, there is little cost to making it. On the other hand, philosophers who take the contents of desires to be properties or centered propositions will find the distinction essential. So we ought to make it.

Preclusion of joint accuracy

The point made in the last section can be generalized from the practical to the doxastic. As we saw, whether a desire is satisfied depends not just on its content but on its context (for example, on who has it and when). Similarly, whether a belief is accurate depends not just on its content but on its context.

The point can be seen easily if we countenance beliefs with centered propositions as their contents. A centered proposition, recall, is a proposition that has truth values relative to a world and a “center” (a distinguished point of view in the world, usually represented by a time and a location or individual). So, for example, there is a centered proposition *I am eating a sandwich* that is true at a world/time/individual triple (w,t,i) just in case i is eating a sandwich at t in w . Quite a few philosophers have suggested, for various purposes, that we broaden propositional attitude psychology to allow beliefs and other attitudes with centered propositions as their contents. (David Lewis, who originated this approach, talked instead of beliefs as the self-ascriptions of properties, but the distinction seems mostly terminological.)

Suppose, then, that Andy believes the centered proposition *I am eating a sandwich*, and that David believes its complement, the centered proposition *I am not eating a sandwich*. Clearly their beliefs are doxastically non-cotenable; Andy could not come to have David's belief without giving up his own. But for all that, both of their beliefs might be accurate. For Andy's belief is accurate if Andy (the agent of the context in which it occurs) is eating a sandwich (at the time of the context in which it occurs), and David's is accurate if David is not eating a sandwich. If Andy but not David is eating a sandwich, then both beliefs are accurate.

That's a case where non-cotenable beliefs are both accurate. It's also easy to imagine a case where cotenable beliefs preclude each others' accuracy. Suppose that at 2 PM Andy believes the centered proposition *I am eating a sandwich*, while at 3 PM David believes the centered proposition *Nobody was eating a sandwich an hour ago*. Imagine that as far as he knows, he is the first person to eat a sandwich that day. Then David's belief is cotenable with Andy's belief, but clearly the accuracy of Andy's belief precludes the accuracy of David's, and vice versa.

Although we can concede that doxastic non-cotenability is a kind of disagreement, we can now see that it is not going to give us everything we might have wanted in a notion of disagreement. For, in at least one sense of disagreement that we care deeply about, when two people disagree in virtue of having certain beliefs, those beliefs cannot both be accurate. If two people disagree, they can't both be right. Similarly, if they agree, it can't be that one's belief is accurate and the other's inaccurate.⁵

We have, then, a third variety of disagreement. To disagree with someone's attitude, in this sense, is to have attitudes the accuracy of which would preclude its accuracy.

I am not going to try to spell out more precisely what I mean by "preclude"; instead, I'll rely on an intuitive grasp. Certainly "preclude" can't be spelled out modally: saying that it is impossible for A and B both to be accurate falls short of saying that the accuracy of A precludes the accuracy of B, because it may be that A

⁵Some relativists have seemed to want to say that it is possible for both parties in a disagreement of taste to be right. They have characterized relativism as a way of vindicating the idea that there can be "faultless disagreements." If by this they mean only that two disagreeing attitudes can be justified or warranted, the idea is unobjectionable, but hardly specific to the domains they take to be apt for relativist treatment. There can be faultless disagreement, in this sense, about matters such as the age of the earth. If, on the other hand, they mean that two disagreeing attitudes can both be accurate, then I can make sense of the idea only by taking them to mean "doxastically non-cotenable" by "disagreeing."

and B cannot both be accurate because B cannot be accurate, quite independently of any relation to A. It is difficult to say what preclusion amounts to in other terms, but I think we have a tolerable grasp of the notion (otherwise we would not be so confident about the counterexamples we can easily construct to various modal explications).

One might wonder why I talk of “accuracy” instead of just talking of truth. Mostly this is to avoid confusion. An attitude or speech act has a content, and this content can be properly said to be true or false. But the same content can be true relative to one circumstance of evaluation and false relative to another. To say that the attitude or speech act is accurate is, roughly, to say that it is true relative to the circumstance that matters. In the case of attitudes with centered contents, this is the world, time, and agent of the context. So although I now take the content believed by David yesterday—the centered proposition *I am eating a sandwich*—to be false, I take David’s belief yesterday to have been accurate, since its content is true at the triple (@, yesterday, David).

The distinction between truth and accuracy doesn’t matter much when we’re considering whether to assert or believe something. For in that case the assertion or belief will be accurate just in case its content is true (relative to the circumstance we occupy). But it matters a great deal when we are considering the speech acts and attitudes of others, or our own earlier speech acts and attitudes. A past assertion need not be retracted if it was accurate—true relative to the relevant circumstance given the context in which it was made—even if its content is one we now take to be false. Conversely, it ought to be retracted if it was inaccurate, even if its content is one we now take to be true.

In prying apart doxastic non-cotenability and preclusion of joint accuracy, I have appealed to examples involving nonstandard contents, like centered propositions. So someone who held to a steady diet of regular, non-centered possible-worlds propositions might question the need for distinguishing the two varieties of disagreement. To such a question, the proper response is the same as I gave above, in connection with the distinction between practical non-cotenability and preclusion of joint satisfaction. Even if the distinction is merely notional, it seems harmless to recognize it.

But some philosophers seem to want to go beyond questioning the need for the distinction, and rely on a monistic conception of disagreement—perhaps the Simple View, noted above to be a special case of doxastic non-cotenability—as a premise in an argument against centered propositions and other nonstandard con-

tents. This is essentially what Cappelen and Hawthorne do in their recent book (Cappelen and Hawthorne, 2009, 96–8). They argue, more or less, as follows:

1. Two parties agree if they both accept the same proposition.
2. If Bill said two days ago “It is raining in Boston” and Janet said two weeks ago “It is raining in Boston,” it is not correct to say that they agree.
3. So, they must not have been accepting the same proposition.
4. So, they were not asserting (and believing) the tensed proposition *It is raining in Boston*.

Once we have distinguished between disagreement as doxastic non-cotenability and disagreement as preclusion of joint accuracy—and between the corresponding senses of agreement—we can see that this argument equivocates. Premise (1) is only true if “agree” is taken in the first sense, while (2) is only true if “agree” is taken in the second sense. So the argument relies on a refusal to make the distinction. But what could justify such a refusal?

Moreover, it seems to me that a need for the distinction can be seen even if we countenance only eternalist propositions, which have truth values relative to possible worlds.

Consider Jane, in this world, the one we call “actual,” and June, in another possible world. Jane believes that Mars has two moons, and June believes Mars has just one moon. Both of their beliefs are accurate, since in June’s world Mars does have just one moon. Does Jane disagree with this belief of June’s?

In a way, yes. Jane could not adopt the attitude June would have without giving up her own belief. But also, in a way, no. Jane’s belief concerns our world, while June’s concerns hers, and both beliefs are accurate. In at least one important sense, two beliefs that are both accurate cannot be said to disagree. The situation here is analogous to the situation with centered propositions believed by different agents at different times.

One might worry that the argument here hinges on “realist” talk of worlds—talk that makes relations between possible situations look more like relations between times than perhaps they should. But perhaps we do not need the apparatus of worlds. We can ask directly whether June, believing what she actually does, is in disagreement with the belief state June *would have* been in in the imagined

counterfactual situation. Note that the question is *not* whether Jane disagrees with what June would have believed. (I wrongly put it this way in “Relativism and Disagreement,” and Cappelen and Hawthorne rightly called me out on it, Cappelen and Hawthorne 2009, ch. 2, § 17.) That question concerns a relation between June and a content and force, but we can’t settle questions of accuracy unless contexts are also in play. The question, then, is whether Jane disagrees with a counterfactual attitude-in-context June might have had—one that she acknowledges would have been accurate given its context.

It seems to me that the answer should be no (in at least one good sense of disagreement). I concede, though, that it is difficult to have any stable intuitions about the case, so I do not want to rest too much weight on this argument.

Preclusion of joint reflexive accuracy

There is still a further distinction to be made. In order to motivate it, though, I will need to explain the difference between objectivist, nonindexical contextualist, and relativist accounts of taste propositions.

Unlike standard (indexical) contextualism, all three of these accounts are happy to countenance beliefs with the content *that Guinness FES is tasty*—not tasty to Joe, or tasty to most people, but just tasty. But they differ in what they say about the intension of this proposition, and about what it takes for a belief with this content to be accurate.

The objectivist says that the proposition has a standard possible-worlds intension. If we specify a state of the world, then there will be an answer to the question whether the proposition would be true were things that way. And a belief or assertion with this content is accurate just in case it takes place in a world relative to which the proposition is true.

The nonindexical contextualist and the relativist both say that the proposition has a non-standard intension—on one version, it has truth values relative to worlds and *tastes*. So even if we specify a state of the world, there is no saying whether the proposition is true until we specify the relevant taste.

The two views diverge, however, on what they say about the accuracy of beliefs and assertions with such contents. The nonindexical contextualist says that a belief is accurate if its content is true relative to the world in which the belief occurs

and the taste of the believer. So, Sam and Sal may believe incompatible taste propositions, and both their beliefs may be accurate, because they have different tastes.

The relativist, on the other hand, denies that accuracy is an absolute matter. An attitude or assertion can only be said to be accurate relative to a context of assessment. A belief is accurate just in case its content is true at the world in which the belief occurs and the taste of the assessor. Its accuracy (at a context of assessment) does not depend at all on the believer's tastes.

Note that although the relativist allows accuracy to depend on the assessor's tastes, from any one context of assessment, at most one of two incompatible taste propositions will be accurate. So, for the relativist, as for the objectivist, doxastic non-cotenable and preclusion of joint accuracy go together. (They will come apart for the nonindexical contextualist, because the accuracy of a belief will depend not just on its content but on who believes it.)

However, once we relativize the notion of accuracy, there are two interestingly different things we can mean by "preclusion of joint accuracy":

1. The accuracy of my attitudes (as assessed from any context) precludes the accuracy of your attitude (as assessed from that same context).
2. The accuracy of my attitudes (as assessed from my context) precludes the accuracy of your attitude (as assessed from your context).

I will use the term "preclusion of joint accuracy" for (1), and "preclusion of joint reflexive accuracy" for (2).

On a relativist account, when two people disagree about whether something is tasty, joint accuracy is precluded, but joint reflexive accuracy is not. Sam's belief may be accurate as assessed from his context, while Sal's is accurate as assessed from his.

For the relativist, then, preclusion of joint accuracy and preclusion of joint reflexive accuracy come apart. For the objectivist, by contrast, they coincide, because accuracy is absolute. A belief is accurate as assessed from one believer's context just in case it is accurate as assessed from the other's.

The relativist, then, need not claim to be vindicating disagreement in all the same senses as the objectivist is. She can acknowledge that, in some respects, disagreement about taste is less robust than paradigm objective disagreements, which do preclude joint reflexive accuracy.

Disagreement in disputes of taste

Enough distinguishing! Let me remind you of my strategy. Instead of posing the problem in a binary way—is there “real disagreement” between Sam and Sal, and if so can the relativist account capture it?—the idea was to ask which of the varieties of disagreement we have distinguished are present in the dispute between Sam and Sal, and which semantic theories allow for these. So, let’s go to it.

We certainly have practical non-cotenability. Sam has an attitude towards Guinness FES that Sal cannot coherently take on board himself without changing his own attitudes towards Guinness FES. Even if Sal does not disagree with anything Sam *believes*, then, there may be reason for them to argue. Sal may want to change Sam’s attitude about Guinness FES, making it congruent with his own, and to do this he may try to call Sam’s attention to various salient facts about the Guinness. These facts will play a role much like that of premises in an argument, except that their intended effect is not a change of belief but a change in taste.

Explaining how there can be disputes about matters of taste, then, does not seem to require that there are disagreements of taste in any sense stronger than practical non-cotenability. And every theory of meaning for taste predicates predicts that we will have at least this. On expressivist accounts, Sam’s and Sal’s speech acts are nothing more than expressions of their non-cotenable attitudes towards Guinness FES. But contextualists, too, can make use of practical non-cotenability to explain disagreements of taste. For although according to the contextualist, Sam and Sal have *asserted* compatible contents, in doing so they have expressed their non-cotenable attitudes of liking and hating the beer, respectively. Indeed, even if Sam had said “I like Guinness FES” and Sal had said “Well, I hate it,” they could be said to disagree.

However, some of the ways in which Sal might naturally express his disagreement with Sam seem to require something beyond practical non-cotenability. First, there’s the word “No” in “No, it’s not tasty at all.” “No” would be quite infelicitous, I think, with explicit self-avowals of attitude:

Sam: I like this beer.

Sal: No, I don’t like it.

Second, Sal could naturally express his disagreement using devices of propositional anaphora:

I don't believe that!
What you're saying is false!
I can't accept that.

This is hard to explain unless Sal takes himself to disagree with what Sam has asserted, or with a belief Sam thereby expresses. It seems to require not just practical but doxastic non-cotenability. And it is hard to see how standard contextualist or expressivist accounts are going to get that.

One interesting avenue for the contextualist is to suppose that in cases like that of Sam and Sal, one or both speakers is presupposing that they do not have relevantly different tastes.⁶ If Sal is presupposing that Sam's tastes are like his, then the belief expressed by Sam's claim, on the contextualist account—Sam's belief that the Guinness tastes good to him—*is not* doxastically co-tenable with Sal's attitudes: Sal could take it on board only by rejecting his belief that Sam's tastes are like his.

The problem with this approach is that it just isn't plausible to suppose that the presupposition of shared taste is in place in all cases of disagreement about matters of taste. Let it be mutually known by Sam and Sal that their tastes in alcoholic beverages tend to be very different. The dialogue with which I began still sounds natural, and it still looks like a disagreement.

Perhaps a better approach for the contextualist is to retreat to a nonindexical version of contextualism. Such a view would retain the key contextualist idea that the accuracy of Sam's belief about the tastiness of Guinness FES depends on Sam's tastes, while the accuracy of Sal's belief depends on Sal's tastes. But it would secure doxastic non-cotenability, since it would take Sam's and Sal's beliefs to have incompatible intensions. Sam could not accept what Sal says without giving up an attitude of his own.

Can we stop here? Although nonindexical contextualism does predict doxastic non-cotenability, it does not secure preclusion of joint accuracy, since it allows that Sam's and Sal's beliefs, despite their incompatible contents, can both be accurate. Relativism, by contrast, secures preclusion of joint accuracy; indeed, that is the main difference between relativism and nonindexical contextualism. Do we have any reason to suppose that disputes of taste, like the one between Sam and Sal, involve preclusion of joint accuracy?

The matter is delicate. But I think things tip in favor of relativism if the parties to such disagreements think of themselves not just as trying to change the other

⁶For this approach, see de Sa (2008).

party's attitudes, but as trying to *refute* them—where the sign of successful refutation is not just that the other party now holds the content of her original claim to be false, but that she retracts her original assertion as inaccurate.

I think that disputes of taste do have this flavor. If Sam eventually gets Sal to like the taste of Guinness FES, Sal will feel pressure to withdraw his earlier assertion that the beer is not tasty. In this respect, disputes of taste are like disputes about any objective matter—for example, the age of the earth.

In another respect, though, they are not much like disputes about paradigm objective matters. For Sam can only compel Sal to retract his assertion by, so to speak, changing Sal's perspective—bringing it about that Sal occupies a context of assessment that differs in semantically relevant ways from the one he occupied before. For, as long as Sal persists in his aversion to Guinness FES, the relativist account predicts, he is warranted in standing by his original assertion (even if it is inaccurate from Sam's perspective). As long as what he asserted remains true as assessed from her current context, he need not retract. In cases of maximally robust disagreement, by contrast, retraction can be compelled (when it can be compelled at all) without any change of perspective. The very same facts that show a claim to be false as assessed from one perspective will suffice to show it false as assessed from any other.

By distinguishing between preclusion of joint accuracy and preclusion of joint reflexive accuracy, we can mark this difference. I think that in disputes of taste we can find the former but not the latter.

Conclusion

Disagreement is the crux of debates between relativists, objectivists, and contextualists. Objectivism accounts for the disagreement we feel in disputes of taste, at the cost of imputing implausible kinds of error and chauvinism to speakers; contextualism avoids chauvinism at the cost of losing the disagreement. Relativism, it is alleged, does better than objectivism because it avoids imputing error and chauvinism, and better than contextualism because it vindicates our intuitions of disagreement.

But if the question is posed in a binary, all-or-nothing way—does relativism allow that disputes of tastes are genuine disagreements, or does it not?—it tends to generate conflicting answers. It is common for objectivists to balk at accepting the

relativist’s claim to vindicate genuine disagreement about matters of taste. After all, on the relativist views, aren’t both parties right from their own perspectives? And doesn’t that show that it isn’t *really* disagreement at all? On the other hand, it is common for contextualists to balk at the relativist’s claim that there *is* genuine disagreement about matters of taste.

By distinguishing varieties of disagreement, we can sharpen up the question and explain why the original question provokes such disparate answers. The question is not whether there is “genuine” disagreement about matters of taste, but rather which of the varieties of disagreement we have distinguished characterizes disagreements of taste. And the main kinds of account we have considered can be defined by the answers they give to this question:⁷

Type of account	Type of disagreement
Standard contextualism, Expressivism	Practical non-contentability
Nonindexical contextualism	Doxastic non-cotenability
Relativism	Preclusion of joint accuracy
Objectivism	Preclusion of joint reflexive accuracy

Evaluating the case for relativism about predicates of taste, then, does not require settling what kinds of disagreement are “genuine,” an issue that seems merely terminological. It just requires determining whether disputes of taste are characterized by preclusion of joint accuracy, for example, or just by doxastic non-cotenability. And we can do this by considering the diagnostics outlined above for these varieties of disagreement.

What this chart shows very clearly is that the relativist can use disagreement as the crux of an argument against the contextualist, while still conceding to the objectivist that there are ways in which the kind of disagreement vindicated by the relativist account falls short of the kind of disagreement one finds about paradigm matters of objective fact. Indeed, the relativist can claim to have found a comfortable middle ground between the objectivist position, which attributes to disputes

⁷Here it is assumed that the parties to the agreement have different tastes, and that it is not plausible to interpret the debate as concerning some shared standard of taste. I do not mean to claim that the indexical contextualist can *never* vindicate more than merely verbal disagreement, or that the nonindexical contextualist can never vindicate more than doxastic non-cotenability. In some cases the two parties may share the same taste, or the same taste may be contextually relevant.

of taste more robust disagreement than there actually is, and the contextualist position, which does not find enough disagreement.

References

Herman Cappelen and John Hawthorne. *Relativism and Monadic Truth*. Oxford University Press, Oxford, 2009.

Dan López de Sa. Presuppositions of commonality: An indexical relativist account of disagreement. In Manuel García-Carpintero and Max Kölbel, editors, *Relative Truth*, pages 297–310. Oxford University Press, Oxford, 2008.

Gottlob Frege. *Posthumous Writings*. University of Chicago Press, Chicago, 1979.

Max Kölbel. Indexical relativism versus genuine relativism. *International Journal of Philosophical Studies*, 12:297–313, 2004.

Charles L. Stevenson. *Facts and Values*. Yale University Press, New Haven, 1963.