Assessment Sensitivity

Relative Truth and Its Applications

Chapters 1–7

John MacFarlane
Department of Philosophy
University of California, Berkeley
E-mail: jgm@berkeley.edu

Draft of December 13, 2009.

Draft: December 13, 2009

Contents

Draft: December 13, 2009

1 Motivating Relativism

THIS book is about how we might make sense of the idea that truth is relative, and how we might use this idea to give satisfying accounts of parts of our thought and talk that have resisted traditional methods of analysis. Although there is a substantial philosophical literature on relativism about truth, going back to Plato's *Theaetetus*, this literature (both pro and con) has tended to focus on refutations of the doctrine, or refutations of these refutations, at the expense of saying clearly what the doctrine *is*. The approach here will be to start by giving a clear account of the view, and then to use the view to solve some problems that have concerned philosophers and semanticists. The main aim is to put relativist solutions to these problems on the table, so that they may be compared with non-relativist solutions and accepted or rejected on their merits. Comparatively little space will be devoted to blanket objections to the coherence of relativism, because these will largely be dispelled by a clear statement of the view.

Saying what relative truth amounts to will take some careful concept-mongering. Before we step into that, it will be useful to see why one might want to get clearer about relative truth. Let us start with a story.

> Sam and Sal are hiking in the woods, and they come upon a clump of Hen-of-the-Woods (*Grifola frondosa*) at the base of a tree.
>
> "Should we take some of this home?" asks Sal.
>
> "Are you crazy?" replies Sam. "Hen-of-the-Woods is poisonous!"
>
> "No, it's not poisonous," says Sal. "I've eaten it many times before with no ill effects."

Nearly everyone, I think, will accept that Sam and Sal are making objective claims about the world; that they are disagreeing in what

they say; and that at most one of them is getting it right. It is true that *poisonous* is in some sense a relative concept; what is poisonous to a human may not be poisonous to a moose or a lizard. But that does not tempt us in the least to deny that there is a substantive disagreement here about a matter of plain fact: whether Hen-of-the-Woods is poisonous to humans. Sam and Sal could have made their claims more explicit by adding the phrase "to humans," but here there was no reason to.

However, let us continue the story.

> Sal reaches down and breaks off two chunks of the fungus, putting one in his own mouth and offering the other to Sam.
>
> "It's tasty, isn't it?"
>
> Sam chews for a while and frowns. "It may be edible," he says, "but tasty it is *not!*"

In its surface features, this dialogue is much like the previous one. But many reflective people will balk at accepting the characterizations they accepted previously: that Sam and Sal are making objective claims about the world; that they are disagreeing in what they say; and that at most one of them is getting it right. If pressed, they may say that there just isn't a "fact of the matter" about whether Hen-of-the-Woods is tasty, in the sense that there is a fact of the matter about whether it is poisonous. Or they may say that questions about tastiness are "subjective" in a way that questions about poisonousness are not. Let's label this syndrome of reactions the *intuition of deficient objectivity*.

1.1 OBJECTIVISM

What underlies this intuition, I suggest, is a realization that if "tasty," like "poisonous," expresses an objective property of things, then our ordinary methods for deciding which things to call "tasty" are radically defective. What methods are these? To a pretty good first approximation, we go by the following principle:

> TP Call something "tasty" just in case its taste is pleasing to you.

No doubt a few extra qualifications are needed. We don't change our views about which foods are tasty just after brushing our teeth, even though the toothpaste can make foods taste very different, so

perhaps we should add "in normal conditions." But even without these qualifications, TP is pretty close to an accurate description of the conditions under which we take ourselves to be warranted in describing something as "tasty." TP certainly gives a natural explanation of what guides Sam's and Sal's application of the predicate, in the story above. Sal calls Hen-of-the-Woods "tasty" because its taste pleases him; Sam refuses to do so because its taste disgusts him.

If you are skeptical that TP guides our use of "tasty," consider how odd it would sound to say:

(1) I'm not sure whether espresso is tasty, but I hate how it tastes.

(2) I've never been able to stand the taste of durian. Might it be tasty?

(3) I love orange juice and hate tomato juice. But who knows? Perhaps tomato juice is tastier.

These speeches sound bizarre. In each case there is a strong tension between the definiteness of the affective reaction and the unwillingness to make a tastiness judgement.[1] But to reject TP is to allow that claims like these can be warranted.

Indeed, it is not clear that our practices in using "tasty" could change in such a way that (1–3) became natural, without losing their point and purpose entirely. We classify things as tasty or not tasty in order to help guide our gustatory deliberations. We eat things we regard as tasty because we expect them to taste good to us. Conversely, we may avoid eating things we don't know are tasty, because they might taste bad to us. But these explanations presuppose something like TP.

By itself, TP is not inconsistent with a robust objectivism about "tasty." If all of us took pleasure in the same foods (in normal conditions), then it would not be unreasonable to regard this pleasure as a natural indicator of some shared objective property of the foods. But in fact, there are large differences in the foods different people find pleasant tasting. A strongly spiced pickle that delights the taste of an Indian may be disgusting to an Eskimo, while the Eskimo's favored

[1] Of course, there are cases in which no tension is felt, for example: "I hate how this stuff tastes right after I've brushed my teeth, but it may well be tasty." So consider a case in which none of the usual impediments to tastiness judgements obtain.

breakfast of raw whale blubber may be disgusting to the Indian. Nor do we need to cut across cultures to find examples: even siblings brought up in the same way can find different foods pleasant. Moreover, we are all well aware of these facts. So if we take "tasty" to express an objective property, we must regard TP as a very unreliable principle for applying it.

Perhaps, the objectivist might reply, each of us believes that our own propensities to take pleasure in food are sensitive to the property of tastiness, even if others' are not. We all think we have won the lottery and acquired a sense of taste that tracks objective tastiness. That would explain our adherence to TP in the face of widespread and evident disagreement in taste. But to say this would be to attribute an unreflective chauvinism to every competent speaker. What basis do we have for taking our own gustatory pleasure to be better correlated with tastiness than anyone else's?

It is useful to compare the case of colors. As with tastes, people do not universally agree about colors. But when there is disagreement, people do not blithely continue to maintain their own views without hesitation. The fact that someone else perceives a different color makes one hesitate in one's own color judgements. It makes one suspect that the lighting is funny, or that one is ill or under the influence of a drug. To insist in such a case that one's own judgement is right, and that the other's is wrong, would be rash and unwarranted. But when it comes to disagreement about whether something is "tasty," we find no comparable hesitation. Why should speakers be chauvinistic in one case but not in the other?

Perhaps there is something the objectivist can say here. Psychologists have shown that those who have low levels of skill in an area tend to wildly overestimate their own abilities (Kruger and Dunning, 1999). In one study, students were given a test of standard English grammar and asked to estimate their percentile rank among the other students taking the test. Students scoring in the bottom quartile on the test rated themselves, on average, in the 60th percentile. Surprisingly, this overestimation persisted, and even became worse, after the students became aware of the discrepancies between their answers and their peers', by being shown others' ungraded exams (1126–7). The researchers explained this by positing that "the same knowledge that

underlies the ability to produce correct judgment is also the knowledge that underlies the ability to recognize correct judgment" (1122).[2]

Perhaps, then, we are all chauvinistic when it comes to taste because we are all very bad at recognizing when something is tasty. Our lack of ability makes us overconfident in our own judgements, even in the face of disagreement with our peers. The question would remain why people who are bad at recognizing colors—color-blind people—do not exhibit a similar overconfidence. But perhaps it is because they routinely receive negative feedback that helps even those with low ability calibrate their own accuracy (Kruger and Dunning, 1999, p. 1131), whereas it is rare for people to chastise others for their judgements of tastiness.

However, the package deal the objectivism is now offering—wholesale attribution of chauvinism, made more palatable by wholesale attribution of cluelessness—is rather hard to swallow. First, it is hard to accept the idea that most of us are highly unreliable in our judgements of tastiness. We all learned the concept "tasty," I suppose, by being exposed to foods that caused pleasure and having mom or dad say "tasty!" It is difficult to believe that the concept we acquired through this procedure expresses an obscure property that we are not very reliable in picking out. How did our word get the meaning the objectivist says it has? Second, even in the face of these experimental findings, it is hard to see why reflection on the facts of disagreement, and on the similarities in our respective trainings with "tasty," shouldn't make at least some of us less chauvinistic and more prone to refrain from judgements of tastiness. But it is difficult if not impossible to find people who suspend judgement about which foods are "tasty" in the way exhibited in (3), above.

There is a further, even more devastating consideration against the chauvinism hypothesis: we use "tasty" in conformance with TP even when we expect our tastes to become better educated. Suppose that,

[2]Interestingly, an earlier version of the study had subjects making judgements about how funny different jokes were. After observing results similar to the ones described above, the researchers speculated that "…it may have been the tendency to define humor idiosyncratically, and in ways favorable to one's tastes and sensibilities, that produced the miscalibration we observed—not the tendency of the incompetent to miss their own failings" (Kruger and Dunning, 1999, p. 1124). The fact that the same results could be obtained in a paradigmatically objective domain is striking.

having grown up tasting only grocery-store Red Delicious apples, Sam enrolls in an apple tasting course. During the four-week course, the students will taste heirloom apples from all over the country. Sam is assured by the instructor that by the end of the course, his tastes in apples will be completely changed. On the first day, the instructor gives Sam four apples to try and asks him which is tastiest. Will he shrug his shoulders and remain agnostic? No—he will answer confidently, on the basis of his present tastes. If we explain this by positing a chauvinistic belief that his tastes track objective tastiness, then we have to suppose that he is taking the course in the belief that it will make his tastiness judgements *less* accurate. And that is odd, if not downright irrational. (And I hesitate to call it irrational only because it is unclear what the *point* of making tastiness judgements would be on the objectivist view.)

1.2 Contextualism

The basic problem with objectivism, as we have seen, is its inability to explain why people tend to use "tasty" in accord with the TP. What account of the meaning of "tasty" would most naturally explain this? An obvious candidate is a simple form of contextualism. According to contextualism, "tasty" is a context-sensitive word. When Sam utters "Hen-of-the-Woods is tasty," he asserts that the taste of Hen-of-the-Woods is pleasing to *Sam*. When Sal utters the same sentence, he asserts that the taste of Hen-of-the-Woods is pleasing to *Sal*. So both Sam's and Sal's claims can be true; neither is led astray by his use of TP.

The kind of contextualism at issue here is the idea that "tasty" needs to be contextually completed by an *experiencer* (or experiencers), or perhaps by a standard of taste. There may be other respects in which "tasty" is contextually sensitive. Indeed, because "tasty" is a gradable adjective, one would expect it to be contextually sensitive at least with respect to the threshold: how high on the tastiness scale something has to be in order to count as "tasty."[3] This kind of contextual variation is not going to help make sense of the dispute between Sam and Sal: it's not as if they agree about *how* tasty Hen-of-the-Woods is, but disagree

[3]See Glanzberg (2007, pp. 8–9), drawing on Kennedy (2007).

about whether this is sufficiently tasty to count as tasty. The question at issue here is whether there is a *further* contextual sensitivity to an experiencer, or to a standard of tastiness that depends on experiencers.

A plausible model for the (putative) context sensitivity of "tasty" would be words like "local," "ready," and "tall." These are generally thought to express *relational* notions whose relata are supplied by context if not explicitly specified. The same bar can be local to Berkeley but not local to San Diego. Alice can be ready to run a mile but not ready to go fishing or take her exam. Sam can be tall for a graduate student but not tall for a basketball player. When one says simply that a bar is "local," or that Alice is "ready," or that Sam is "tall," one intends to predicate one of these more determinate properties. This much, I think, is uncontroversial, though there is a lot of controversy about just how to explain what is happening with "bare" uses of "local," "ready," or "tall."[4] Some writers hold that such words are associated with variables in the logical form that, when not bound by quantifiers or supplied a value explicitly, are given values by context. Some hold that the completion or enrichment does not require any syntactic trigger. Some hold that these words express simple, nonrelational properties, but that the full communicative content of a speech act is richer and more determinate than the minimal "official" content of its sentential vehicle. We won't need to sort out these issues about semantic content here. We will focus instead on *asserted* content, as it is common ground between all of these writers that in saying "Alice is ready," one is asserting that she is ready for X, for some X. In this sense, at least, words like "ready" are clearly context-sensitive.

Is "tasty" context-sensitive in the same sense? Are utterances of "Hen-of-the-Woods is tasty" generally understood as assertions that Hen-of-the-Woods tastes good to the speaker? In favor of this hypothesis, it might be noted that "tasty," like "local" and "ready," can occur in explicitly relativized form. I can characterize a food as "tasty for teenagers" or as "tasty for me." It is natural to suppose that when such relativizations are not explicitly supplied, the relevant experiencer or standard is supplied by context. Indeed, as Lasersohn (2005, p. 656) points out, it is difficult to see how the objectivist can explain the

[4] See for example Stanley (2007), Cappelen and Lepore (2005), and the essays in Preyer and Peter (2007).

explicitly relativized forms. With paradigm objective predicates, like "five feet tall," we have no similar explicitly relativized forms; we do not say that someone in "five feet tall to me," or "five feet tall for a teenager." So it is difficult to see how these relativizing phrases could be dealt with on an objectivist account.

However, there are some serious problems for contextualism about "tasty." It cannot well account for our intuitions of agreement and disagreement, and it cannot explain why speakers are willing to retract earlier assertions made using "tasty" when their tastes have changed.

1.2.1 Agreement and disagreement

When, in our dialogue, Sal says that Hen-of-the-Woods is tasty and Sam says that it is not tasty, they seem to be disagreeing with each other, not making compatible claims. This is not expected on the hypothesis that each is asserting that Hen-of-the-Woods tastes good to *him*, because the claim that Hen-of-the-Woods tastes good to Sal is perfectly compatible with the claim that it does not taste good to Sam.

Here we do see a clear difference between "tasty" and our paradigm context-sensitive words. Suppose Abe says that Sarah's favorite bar is a local bar (meaning local to Anchorage), and Sam says that her favorite bar is not a local bar (meaning local to Savannah). Although we can truly say,

(4) Abe said that Sarah's favorite bar was a local bar, and Sam replied that it was not a local bar,

we would not describe this as a case of disagreement. Again, if Abe says that Sarah is ready (meaning *ready to go swimming*), and Sam says that she is not ready (meaning *not ready for her bar exam*), then there is no real disagreement, only the surface appearance of it. Abe and Sam might excusably think that they disagree, but only because they have misunderstood each other.

The contextualist might resist the intuition that Sam and Sal are disagreeing, or try to explain it away. She might take some support here from the saying, "there's no disputing taste." So it is worth recounting some reasons for thinking that there really is disagreement in such cases, and considering some ways in which the data might be reinterpreted.

First, it is natural to use explicit marks of disagreement, such as "No," "I disagree," "you're mistaken," or "that's false." These responses would be inappropriate if the two parties were simply making claims about what tastes good to *them*:

(5) A: Hen-of-the-Woods is tasty.
B: No/I disagree/You're mistaken/That's false, it's not tasty.

(6) A: Hen-of-the-Woods tastes good to me.
B: No/I disagree/You're mistaken/That's false, it doesn't taste good to me.

Faced with this data, contextualists sometimes note that words like "No," "You're mistaken," and "That's false" sometimes target something other than the asserted proposition. For example, they can target the content of the reported speech or attitude:

(7) A: Sahin said that you had a car.
B: No/That's false. I don't have a car.

They can also target a presupposition of the assertion:

(8) A: Your wife is very beautiful.
B: No/You're mistaken. We're not married.

Grice (1989, pp. 64–5) observes that disagreement markers can also target the result of "factoring out" a shared assumption from the asserted content. He gives this nice example:

(9) A: Either Wilson or Heath will be the next Prime Minister.
B: I disagree, it will be either Wilson or Thorpe.

Here there is disagreement even thought the two disjunctions are compatible. The explanation is that it is "accepted as common ground that Wilson is a serious possibility" (65). So what is being rejected is just that Heath is the other serious contender.

But it is difficult to see how any of these models would apply to the contextualist's proposal about "tasty." Moreover, a striking difference between (5) and (7) and (8) is that the asserted proposition is explicitly negated in the reply ("…it's not tasty"). So the real parallels would be these:

(7′) A: Sahin said that you had a car.
B: No/That's false. He didn't say that I had a car.

(8′) A: Your wife is very beautiful.
B: No/You're mistaken. She's not very beautiful.

And here "No," "You're mistaken," and "That's false" clearly target the whole asserted proposition.[5]

The contextualist might try claiming that the marks of disagreement express attitudes towards the *words* used, not the propositions they express. So, "No" in (5) would mean: "No, I wouldn't use that sentence to make an assertion." And "That's false" would mean "That sentence, as used by me now, would express a falsehood."[6] But, in the absence of data supporting these alternative uses of "No" and "That's false" in other contexts, this just seems like special pleading. Moreover, to explain the data, the contextualist would have to hold that in disputes of this kind, "No" and "That's false" *always* get the nonstandard reading. Otherwise there ought to be a reading of the following dialogue in which B is not contradicting herself:

(10) A: Apples are tasty.
B: That's not true. But apples are tasty.

The contextualist needs to explain why such readings do not seem available.

Note, also, that the phenomenon persists even when the demonstrative "That" is replaced with a term that explicitly denotes the proposition expressed. Instead of replying to Sal by saying "That's false," Sam might have said (somewhat pedantically) "The proposition you expressed is false." Here the nonstandard reading is explicitly blocked. Moreover, there are plenty of ways of disagreeing with what someone has said besides saying "No" or "That's not true." One might

[5](9) is trickier. Here it seems okay for the objector to say: "I disagree, it's not the case that Wilson or Heath will be the next Prime Minister," but only if some emphasis is given to "Heath." The fact that special emphasis is needed suggests that this is a case of *metalinguistic negation* (Horn, 1989, ch. 6), as does the fact that the negation cannot be incorporated into the disjunction: "#I disagree, neither Wilson nor Heath will be the next Prime Minister" (Horn, 1989, §6.4.1).

[6]I have not seen this argument in print, but I have heard it in conversation, and Kölbel (2002, p. 39) finds it worth criticizing.

say, for example, "I disagree" or "You're mistaken." The contextualist would have to find a reading for "disagree" on which the assignment of different truth values to the same indexical sentence at different contexts of use could count as a "disagreement."

A second indication that we take ourselves to be disagreeing about matters of tastiness, besides the explicit disagreement markers, is that we sometimes *argue* about them: "Brussels sprouts, tasty? They taste like grass! Do you also say that grass is tasty?" We do not generally argue with others' claims about what tastes good to *them*, so the fact that we argue about what is "tasty" speaks against the contextualist analysis.

It is of course open to the contextualist to say that our tendency to argue about claims of taste, and our perception that we are disagreeing with each other in making them, is just a delusion. But if the contextualist is willing to attribute this much systematic error to speakers, it is unclear what reason remains to prefer contextualism to a simple objectivist view. After all, what seemed unattractive about objectivism was precisely that it forced us to attribute systematic error to speakers. Indeed, it seems that the contextualist will have to attribute the same kind of chauvinism that the objectivist does, *plus* a semantic error that the objectivist does not attribute. For in order to explain why we take ourselves to be disagreeing in our claims of taste, the contextualist will have to take us to have an inchoate objectivist theory of the semantics of these statements. But if that is how we think of them, then our habits of asserting that things are tasty on the basis of our own affective reactions, in the face of abundant evidence of the diversity of such reactions, must be explained by the same unreflective chauvinism we found objectionable in our discussion of objectivism.

We have focused here on disagreement, but the analogous points can be made about *agreement*. Suppose both Sam and Sal like the taste of raisins. Both might say,

(11) Raisins are tasty,

and we will naturally report them as having agreed:

(12) Sal and Sam agree that raisins are tasty.

On the contextualist analysis, (12) must be interpreted as

(12) a. Sal and Sam agree that raisins are tasty to Sal, or
b. Sal and Sam agree that raisins are tasty to Sam.[7]

But we can easily construct a case in which (12) seems true even though none of these readings are true. Just imagine that Sam and Sal both like the taste of raisins, but neither thinks the other does. They seem to agree, not about whether raisins taste good to some party or parties, but about whether they are tasty—where that is something different.[8]

1.2.2 Retraction

When our own tastes change, so that a food we used to find pleasant to the taste now tastes bad, we may say that we were mistaken in saying that the food was "tasty." When I was a kid, I once told my mother, "Fish sticks are tasty." Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I've changed my mind about the tastiness of fish sticks. So, if someone said, "But you said years ago that fish sticks were tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty any more," since that would imply that their taste changed. Nor would I say, "When I said that, I only meant that they were tasty to me then." I *didn't* mean that. Indeed, at the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty.

The contextualist cannot easily explain why I would retract my earlier assertion. On the contextualist account, the content I expressed then by "fish sticks are tasty" is perfectly compatible with the content I express now by "fish sticks are not tasty." So retraction should not be required. Indeed, it should seem as odd as it does in this dialogue:

SAM: [*in Phoenix*] You can get a swamp cooler at any local hardware store.
SAM: [*in Boston*] Nobody sells swamp coolers around here.
JANE: But you said you can get one at any local hardware store!
SAM: I take that back.

[7] Or even "Sal and Sam agree that raisins are tasty to Sam and Sal both." See §1.2.4, below.

[8] Cappelen and Hawthorne (2009, ch. 2) accept that "agrees" is a good test for invariance of content.

The contextualist might try to explain retraction of claims of taste by taking them to be claims about what would taste pleasant to a suitably idealized version of the agent.[9] If, in saying that fish sticks are tasty, I was predicting that they would be pleasing to my more educated palate, then I ought to retract my claim in light of what my more educated palate tells me.

But is it really plausible that in calling things "tasty," we are making claims about how they will strike idealized versions of ourselves? Consider the story of Sam and the apple tasting course, recounted in §1.1. On the first day, the instructor gives Sam four apples to try and asks him which is tastiest. If he were really being asked to say which would best please his future, educated palate, shouldn't he shrug his shoulders and remain agnostic? But that is not how we seem to use the word "tasty." He will answer confidently, despite his belief that within a month his tastes may be very different.

Like objectivism, then, contextualism fails to capture what is distinctive about words like "tasty." To be sure, there is something that seems right about it: it respects the idea—enshrined in TP—that the proper criterion for applying the word "tasty" is one's own affective reactions. But there is also something wrong about it: it doesn't respect our sense that the dialogue with which we started, Sam and Sal are genuinely disagreeing with each other (making contradictory claims). And it can't explain why it should be appropriate to retract a tastiness claim made in light of one's earlier tastes.

1.2.3 Clarifying the challenge

It is important to be clear about the nature of the challenge to contextualism. The challenger can concede that there are many uses of "tasty" for which a contextualist analysis of what is asserted is correct. All that is needed are *some* uses that cannot be accounted for on contextualist lines.

In general, the path from a speaker's words to the proposition she asserts can be a crooked one, capable of being followed only by a listener with appropriate knowledge of the speaker's expectations. Suppose we're all waiting for a visiting speaker to appear. The door opens and three men walk in. One of them is somewhat taller and

[9] For the idealization move, see Egan (forthcoming).

heavier than the other. By saying, "the big guy," I can assert that the largest of the three is the visiting speaker. I don't need to put all of that in words, because I know that my hearers will divine what I'm trying to tell them.

More to the point: being color-blind, I will sometimes use "That's green" to assert that the demonstrated object looks green *to me*. If I am met with the objection, "No, it's not, it's red," I will respond "I only meant that it looked green to me." In such cases, I rely on my hearers to supply the relativization to me; if they fail to do so, they have misunderstood me (perhaps excusably—it may be that *I* am at fault for assuming too much ability to divine my own intentions). Importantly, I don't *always* use "That's green" this way. Usually I use it to assert that something is green, and when I do this, I will regard myself as contradicted by someone else who says, "No, it's red."

In the same way, it may be that speakers sometimes use "That's tasty" to assert that things taste good *to them*. A sign of such uses is that the speakers won't regard themselves as disagreeing with others who say of the same thing that "it's not tasty." If the others signal disagreement, they will make it clear that they only *meant* that the thing tasted good to them. It would be surprising if people did not sometimes use "tasty" this way; after all, we tend to use language as economically as possible.

The point against the contextualist is not that "That's tasty" can *never* be used to say that something is tasty to the speaker—that would be implausible, given the flexibility of language—but that it is sometimes (and, I would suggest, normally) used to say something else, and the contextualist has given no account of what this something else might be. Even if speakers sometimes fall back and say "I only meant that it tasted good to me," they do not *normally* do this.

The uses of "tasty" Lasersohn (2005) calls "exocentric" can also be explained on contextualist lines. When I say of a small child's baby food, "This brand is really tasty," I am asserting that this brand tastes good to the child. If you reply by trying some, scowling, and saying that it isn't tasty at all, you've simply misunderstood what I meant to assert. I would be perfectly happy to paraphrase what I asserted with an explicit relativization to the child, and I don't regard myself as

joining issue with others who assert that the baby food isn't tasty.[10]

Defenders of contextualism often point out that the intuitions of disagreement that challenge contextualism become weak or nonexistent when the speakers are of different species. For example, Cappelen and Hawthorne (2009, ch. 4) observe that "[t]hose of us who are disgusted at the thought of drinking milk that has hair floating around in it were shaky on disagreement verdicts when imagining talking cats drinking from saucers full of milk that had floating hair in it." This suggests, perhaps, that when people use "tasty" they mean "tasty for humans," or perhaps "tasty for creatures like us." But even this could be conceded without damaging the central critique of contextualism given above. The argument is not that what is asserted by predicating "tasty" of a thing never varies with the context, but that one cannot fully explain our perceptions of agreement and disagreement by adverting to contextual shifts in the proposition asserted. Even if we suppose that in uttering "That's tasty," we are normally asserting only that the thing is tasty for humans, this concession to contextualism hardly helps explain the disagreements we can find among human speakers.[11]

The challenge to contextualism focuses on paradigm cases of disagreement, for which no contextualist (or explicitly relational) account of what is asserted seems plausible. These cases, if compelling, show that some account of "tasty" that goes beyond the contextualist approach is needed. But the challenge is consistent with the very plausible idea that what is asserted by predications of "tasty" does depend to some extent on contextual factors.

[10]Lasersohn takes a different approach, holding that in this case one has asserted the same proposition that one's interlocutor has denied, but that one's interlocutor has erred in evaluating one's assertion "autocentrically" instead of "exocentrically." Some difficulties with this approach will be discussed in chapter 7. I do not see any motivation for resisting the contextualist analysis in this case, once one sees that accepting the contextualist account of some uses of "tasty" doesn't commit one to accepting it for all uses.

[11]Compare Richard (2004), who notes that the truth of predications of "richness" depend both on a reference class ("rich for a Manhattan resident" vs. "rich for a Queens resident") and on a cutoff point (exactly how rich do you have to be to be rich for a Manhattan resident?). Richard suggests a contextualist treatment of the former, while giving reasons for thinking that the latter is not amenable to a contextualist treatment.

1.2.4 Collective contextualism

Perhaps the most obvious strategy for accounting for a disagreement in a contextualist framework is to construe predications of "tasty" as concerning what tastes good to a contextually relevant *group* that includes all parties to the disagreement. This move secures something for the disagreement to be about: whether the food in question tastes good to all (or some, or most) of the members of the group. We might call this *collective* contextualism, and contrast it with the *solipsistic* contextualism we have been considering up to now.

Although collective contextualism can explain the data about disagreement and agreement that proved troublesome for solipsistic contextualism, it has a hard time explaining why people are not more hesitant in predicating "tasty" of things. Consider again Sal's assertion, of Hen-of-the-Woods, that "It's tasty." According to collective contextualism, what Sal has asserted is that Hen-of-the-Woods tastes good (under normal conditions) to both Sal *and* Sam. This is something that Sal should not assert unless he is fairly confident that Sam shares his own tastes. But surely Sal can legitimately assert "It's tasty" on the basis of his own reactions, without having any idea whether Sam is prone to agree with him—and even if he thinks that Sam is likely to disagree with him. Collective contextualism also seems to predict, even more implausibly, that on hearing Sam's response, Sal should immediately retract his assertion, since he now has excellent evidence that it is false that Hen-of-the-Woods tastes good to both of them. Thus collective contextualism makes false normative predictions.[12]

The contextualist may urge that we (along with Sal) are making a mistake in thinking that it is okay for him to assert what he does without being confident that Sam's tastes are relevantly similar, or in thinking that he need not retract his assertion when he hears Sam's response. But then some explanation is needed of why we should systematically make such a mistake. It is unclear to me what such an explanation could look like. Nor is it clear why such an explanation, if it were produced, could not be appropriated in defense of objectivism. As we have already observed, the case for contextualism over objectivism is undermined if massive systematic error in speakers' judgements of the truth of claims is conceded.

[12]For similar objections, see Lasersohn (2005, pp. 651–2).

Things get much worse for collective contextualism when one considers the range of speakers (and thinkers) who might be said to disagree with the Sal. Suppose Sarah, who has been hiding behind the bushes, jumps up after Sal's avowal and says "Hen-of-the-woods is definitely not tasty." Surely the case for saying that Sarah and Sal disagree is as strong as the case for saying that Sam and Sal disagree. So the collective contextualist must include Sarah in the contextually relevant group. Or suppose Jim later watches a videotape of the proceedings and exclaims, "Hen-of-the-woods is not tasty." He seems to be disagreeing with Sal, too. So Jim must also be in the contextually relevant group. It is hard to see how any specification of the relevant group short of "anyone who hears, or will ever hear of, this assertion" could account for all the potential disagreements. And now the problems scouted earlier are multiplied, since the contextualist has Sal asserting something about the tastes of an unknown and indefinitely large group of people. How could he ever take himself to be warranted in making such a claim?

A more plausible version of collective contextualism is suggested by David Lewis's metaphor of conversation as a game with an evolving "scoreboard" (Lewis, 1979*b*). Lewis suggests, for example, that one component of conversational score is a setting for "standards of precision." This can be set explicitly—"let's speak strictly for a while"—but more often it changes through *accomodation*; that is, it evolves "in such a way as is required in order to make whatever occurs count as correct play" (240):

> Taking standards of precision as a component of conversational score, we once more find a rule of accommodation at work. One way to change the standards is to say something that would be unacceptable if the standards remained unchanged. If you say "Italy is boot-shaped" and get away with it, low standards are required and the standards fall if need be; thereafter "France is hexagonal" is true enough. But if you deny that Italy is boot-shaped, pointing out the differences, what you have said requires high standards under which "France is hexagonal" is far from true enough. (245)

Along similar lines, the collective contextualist might posit a *standard of taste* as a component of conversational score. This shared standard

would fix what counts as "tasty" within the conversation. This standard would be set and adjusted by negotiation and accommodation. If a speaker asserted that spinach is tasty, for example, and this assertion were accepted, then the shared standard would become one that counts spinach as tasty. Because the standard is shared and determined collectively, it need not coincide with any of the conversationalists' idiosyncratic tastes.[13]

This form of collective contextualism can explain why a speaker might assert that something is "tasty" even when there are doubts about whether it tastes good to the other parties to the conversation, and even when it is known not to taste good to them. Such assertions can be seen as attempts to "push" the shared standard of taste in a particular direction. Arguments and disagreements about what is "tasty" can, on this view, have two sources. It may be that the parties agree about the shared standard and disagree about what the food in question tastes like. (This is easiest to imagine when they don't have the food right in front of them.) Or it may be that the parties both have accurate knowledge of the taste of the food (described non-evaluatively) but disagree about what the standard should be. The disagreement in this case would be like the disagreement between people who know what France looks like on a map but dispute whether it should be accepted that France is "hexagonal."[14]

One might worry that the proposal disconnects the shared standard of taste too much from individual affective reactions. A strong connection could be forged by insisting that individuals not accomodate unless the new standard of taste accords with their own tastes. Accommodation would then require an actual shift in tastes, and an attempt to "push" the shared standard would, in effect, be an attempt to change others' tastes. But it is not clear that such a radical move is needed. Given that conversational conclusions have effects on things like what mushrooms the group buys for dinner, it is easy to see why speakers should have reason to push the shared standard to resemble their own idiosyncratic tastes as far as possible, but also to acquiesce in a standard that departs from this ideal in various ways.

[13]See Lasersohn (2005, pp. 659–662) and de Sa (2008, pp. 302–3) for relevant discussion.

[14]For a similar strategy in defense of a form of epistemic contextualism, see DeRose (2004).

One might also worry about whether speakers have sufficient grasp of the notion of a standard of taste for it to play the role proposed. We can understand Lewis's example because we can think of a standard of precision as a cutoff point on a roughly linear scale. In accepting that Italy is "boot-shaped," we push the cutoff for acceptable sloppiness a ways down this line. In refusing to accept that France is "hexagonal," we push it up the line. But we have no comparable grasp of what a standard of *taste* would look like. A standard of taste would certainly not be a cutoff point on a linear scale. It would be something much more complex. But then it is unclear how accommodation would work. Suppose we accept an assertion that spinach is "tasty." We need to move the shared standard of taste to a setting on which spinach counts as "tasty." But how do we do that? There are many different standards of taste that would count spinach as tasty, but diverge in other dimensions. So it is hard to see how the moves in a conversation could establish an even moderately determinate shared standard of taste, the way they can establish a shared standard of precision.[15]

A more serious worry about the shared scoreboard version of collective contextualism is its inability to make good sense of continued, clear-eyed disagreement about matters of taste. (By "clear-eyed" disagreement, I mean disagreement in which each party is in a good position to make the judgement, and each is fully aware of the others' position and views.) It is easy to imagine Sal and Sam continuing to disagree—perhaps indefinitely—about whether Hen-of-the-Woods is "tasty." They may find it profitless to continue the discussion for much longer, but they will not abandon their own views on discovering the other's intractability. This is not what we should expect on the shared scoreboard proposal. For it says that in such cases, no shared standard is established, and so (presumably) all predications of "tasty" lack truth value. Realizing this, and seeing that their attempts to push the shared standard in their own favored directions have failed, Sal and Sam should stop predicating "tasty" of things. But that is not what they will do.

[15]Of course, even standards of precision can have several dimensions. One might be very precise about the application of shape terms, for example, but not about the application of color terms. Still, it is easier to see how conversationalists might represent to themselves a shared standard of precision, even a multidimensional one, than a shared standard of taste.

It is useful here to compare "tasty" with our paradigm, "hexagonal." It is difficult to imagine a discussion between two people who have a map of France right in front of them continuing like this:

A: France is hexagonal.
B: No, it's not. Look at these wavy lines.
A: I see the wavy lines, but that's irrelevant. France is hexagonal.
B: It most certainly isn't.
A: You haven't convinced me.
B: Nor you me.

The dialogue is absurd because both parties would immediately abandon their attempts to use "hexagonal" according to their preferred standards when they saw that the other would not accept this standard. But a parallel discussion involving "tasty" is quite imaginable. That suggests, I think, that "tasty" is not well understood on the model of "hexagonal," as the collective contextualist has suggested.

Another problem with the proposal is that it makes sense of disagreement only within the confines of a "conversation," something for which it makes sense to imagine a shared scoreboard. But we perceive disagreement about what is tasty outside of such contexts, too. I am sure there are people in China who are disgusted by foods I find quite pleasing, and vice versa. It seems to me that we disagree about whether these foods are tasty. But we are not involved in any kind of conversation; not only do we not exchange words or have mutual expectations, we don't even know each other. The problem with the single scoreboard approach is that it explains only *intra*-conversational disagreement, leaving *inter*-conversational disagreement unaccounted for. This is not a stable resting point. Once the importance of accounting for disagreement has been conceded, one cannot limit oneself to disagreement within conversations.[16] And it is hopeless to widen the bounds of "conversations" as needed to make *all* disagreement intra-conversational. For it is only if conversations are bounded and relatively self-contained that we can really make sense of the idea of a shared scoreboard.

[16]For a similar point in response to DeRose (2004) on epistemic contextualism, see Feldman (2001, p. 29).

1.3 Expressivism

Given the failure of objectivism and contextualism to account for the facts about our use of "tasty," it is natural to question what both take for granted: that in deploying "tasty" we are making genuine assertions, taking a stand on how things are. Compare what A. J. Ayer says about moral vocabulary:

> The presence of an ethical symbol in a proposition adds nothing to its factual content. Thus if I say to someone, 'You acted wrongly in stealing that money,' I am not stating anything more than if I had simply said, 'You stole that money.' In adding that this action is wrong I am not making any further statement about it. I am simply evincing my moral disapproval of it. It is as if I had said, 'You stole that money,' in a peculiar tone of horror, or written it with the addition of some special exclamation marks. (Ayer, 1959, p. 107)

Applying Ayer's thought to "tasty," we get what I will call *classical expressivism*: the view that in saying "It's tasty" one is not making an assertion, but simply expressing one's liking for a food.

It is crucial to mark the distinction between *expressing* one's liking for a food and *asserting* that one like the food. One does the former, but not the latter, when one smacks one's lips in delight after a good meal. One does the latter, but perhaps not the former, when one tells one's host, with an unconcealed expression of dutiful weariness, that one liked his cooking.[17]

The classical expressivist agrees with the objectivist, against the contextualist, that in saying "That's tasty" one is not asserting that the food tastes pleasant to one (or to a larger group). But that is because,

[17] Things are complicated by the fact that, according to some speech-act theorists, assertion is the expression of belief (for discussion, see MacFarlane, forthcoming). The expressivist who takes this view of assertion can still draw a principled line between assertions and the expression of "non-cognitive" attitudes like desires and preferences. However, she must tread delicately in saying what it is to express an attitude. As Jackson and Pettit (1998) note, one might naturally take "That's good" to express not just the speaker's approval, but also her *belief* that she has this attitude, in the absence of which she would not have uttered the sentence. The expressivist will need to explicate "expression" in a way that distinguishes between the first-order attitude of approval and the second-order belief about this attitude.

unlike the objectivist, the expressivist doesn't think one is asserting anything at all. For the expressivist, saying "That's tasty" is just a verbal way of smacking one's lips, just as "Drat!" is a verbal way of expressing disappointment. In this way, the expressivist avoids the fundamental problem that dogs objectivism: explaining how we can persist in making assertions that (unless we are chauvinistic) we can only regard as highly prone to be mistaken. One makes no mistake (except, occasionally, a mistake of etiquette) in expressing one's liking for a food.

1.3.1 Disagreement and retraction

The objectivist might legitimately wonder, though, whether the expressivist does any better than the contextualist in accounting for the apparent *disagreement* we express using "tasty." Ayer himself notes that expressivism vindicates disagreement only in a relatively weak sense:

> Another man may disagree with me about the wrongness of stealing, in the sense that he may not have the same feeling about stealing as I have, and he may quarrel with me on account of my moral sentiments. But he cannot, strictly speaking, contradict me. For in saying that a certain type of action is right or wrong, I am not making any factual statement, not even a statement about my own state of mind. I am merely expressing certain moral sentiments. And the man who is ostensibly contradicting me is merely expressing his moral sentiments. (Ayer, 1959, p. 107)

We might usefully distinguish *disagreement in attitude* from *disagreement in claim*. The classical expressivist can make sense of the first in disputes involving "tasty," but not the second. Perhaps, though, disagreement in attitude is all that is implicit in our ordinary use of "tasty." So argues Maudlin (2007):

> They [the ordinary folk] might well maintain that although Wright's and Williamson's attitudes are genuinely incompatible, in the sense that no single person can simultaneously have them both, they are not attitudes to any proposition at all: they are rather attitudes toward rhubarb, or towards eating rhubarb, or towards how rhubarb tastes (to the given individual). (In this

> sense, sitting and standing are genuinely incompatible postures, postures that involve no propositional attitudes.)

Although Ayer himself concluded that an expressivist account of aesthetic terms implies that there is "no possibility of arguing about questions of value in aesthetics" (Ayer, 1959, p. 113), the kind of disagreement in attitudes that expressivists accept is enough to make such arguments intelligible. Suppose Lizzie likes Sam, while Sal hates him. Lizzie and Sal might try to induce each other to share their attitude towards Sam, and they might do so by offering considerations and counter-considerations, in just the way we do when we are arguing about the truth of a proposition. The fact that in this case the dispute does not concern any particular fact about Sam does not prevent it from having the shape of an argument.

But if Maudlin is right that the only disagreement that exists in disputes involving "tasty" is disagreement in attitude, it is hard to see what would favor expressivism over contextualism. For the contextualist, too, can concede that in such disputes the two parties have attitudes to the food that are incompatible in the way that sitting and standing are incompatible postures—they cannot belong to a single person at a single time. We earlier rejected contextualism for its inability to account for disagreement, but if this is all the disagreement there is, contextualism can account for it handily.[18]

The expressivist might still object that contextualism predicts the possibility of *agreement* when it isn't possible intuitively:

> A: That's tasty (*tastes-good-to-A*).
> B: Yes, I agree. But it isn't tasty (*tastes-good-to-B*).

But the contextualist can explain what is wrong with this dialogue on pragmatic grounds. Compare:

> NYU STUDENT: The statue of liberty is a local attraction (*local-to-New-York*).
> PARENT, ON THE PHONE FROM MILWAUKEE: Yes, I agree. But it isn't a local attraction (*local-to-Milwaukee*).[19]

[18] The point is made in Jackson and Pettit (1998, p. 251) and Dreier (1999, p. 569). I do not mean to suggest here that Maudlin is defending expressivism; he may have in mind a contextualist view.

[19] I owe the example to Niko Kolodny.

Even when a contextualist treatment is obviously appropriate (as it is with "local"), there is a strong expectation that the contextual completion will remain throughout a sentence unless some indication is given to the contrary. The infelicity of the first dialogue does not, then, count against a contextualist treatment of "tasty."

So, even if Maudlin were right that parties to a disagrement about the "tastiness" of rhubarb disagree only in the sense that they have conflicting attitudes towards rhubarb, this would simply undermine our objections to a contextualist treatment of "tasty," and would leave us no reason for preferring expressivism.

However, there is good reason to think that disagreements involving "tasty" go beyond the kind of disagreement in attitude that Maudlin finds. For one thing, mere disagreement in attitude would not motivate retraction. On becoming convinced through experience that peaty whiskeys are tasty, one might say:

(13) Last year I said that they weren't very tasty, but I take that back. I was wrong.

An assertion can be retracted, but it doesn't make much sense to "take back" or retract the expression of an attitude. (Imagine a dirty old man attempting to "take back" a lecherous leer, on finding that its object is an employee of his.)

It is not clear, then, that expressivism does any better than contextualism in explaining disagreement about taste. On the other hand, it faces a number of difficult problems that contextualism avoids. These will occupy us for the rest of the chapter.

1.3.2 Force and content

Frege taught us to analyze speech acts by factoring them into two components—*force* and *content*. Consider, for example, Tom's assertion that there is fresh powder in the mountains. Its content—what he has asserted—is that there is fresh powder in the mountains. Its force is that of an assertion. He could have asserted that there is black ice in the mountains; in that case, his speech act would have had the same force but different content. Or he could have asked whether there is fresh powder in the mountains; in that case, his speech act would have had a different force but the same content. The same distinction

can be applied to mental states. Wondering whether there is fresh powder in the mountains and desiring that there be fresh powder in the mountains share a content but differ in force; believing that there is fresh powder in the mountains and believing that there is black ice in the mountains share a force but differ in content.

The force/content analysis makes the study of language and thought more systematic. One part of our study can concern itself with the possible contents of thoughts and speech acts, and another with the possible forces. Combining these, we can account for the significance of acts with any of these possible contents and possible forces.

Classical expressivism gives up the force/content analysis in the domains to which it applies. It denies that there are propositions characterizing foods as tasty. (If there were, the job of "That's tasty" would presumably be to assert such propositions, and we wouldn't need to talk separately of expressing attitudes.) Instead of letting the significance of utterances of "That's tasty" emerge from separate accounts of assertoric force and the content of tastiness-ascribing propositions, the expressivist explains their significance *directly*, by saying what they are used to do. And that is problematic for at least four reasons.

NON-DECLARATIVES

The first reason is that "tasty" occurs not just in declarative sentences like "That's tasty," but in interrogative, imperative, and optative ones:

(14) Is that tasty?

(15) Make it tasty!

(16) If only that were tasty!

None of these sentences are used to express the speaker's liking for the demonstrated food. So even if we accept the expressivist's account of the meaning of "That's tasty," we are left without any account of the meaning of very similar non-declarative sentences.

On the classical approach, by contrast, we give an account of the propositional content expressed by "That is tasty" (at a context), and our existing accounts of interrogative, imperatival, and optative force

combine with this to give us an account of the meanings of these sentences.[20]

MENTAL ATTITUDES

Second, in addition to *saying* "That's tasty," we can also *think* it. One can believe that a certain food is tasty, suppose that it is tasty, wonder whether it is tasty, and desire that it be tasty. In so doing, one is not expressing an attitude, because one need not be expressing anything at all. One might just keep one's thoughts to oneself. So the classical expressivist account does not extend in any obvious way to an account of these attitudes. But it does rule out the standard kind of account in terms of content and attitudinal force, because it denies that there *is* a content of the sort that would be needed (the proposition that *that is tasty*).

A hard-line expressivist response would bite the bullet and deny that there are genuine attitudes of believing that a food is tasty, or wondering whether a food is tasty. This bullet-biting response would have to be coupled with an expressivist account of what we are doing when we say, for example,

(17) He believes that licorice is tasty.

(18) He wishes that licorice were tasty.

But again, it is unclear what such an account could look like. What are *we* expressing when we utter sentences like (17) or (18)? Certainly not any attitude of our own towards licorice.

PROPOSITIONAL ANAPHORA

A third problem is that it is natural to use propositional anaphora in connection with uses of "tasty":

(19) This fish is tasty!

(20) a. Yes, that's true.

[20] The most classical modern version of this can be found in John Searle's work on speech-act theory (Searle, 1969, 1979). Searle's account is perhaps too rigidly classical to account for the phenomena—most linguists now take the contents of questions to be something other than propositions, for example (Groenendijk and Stokhof, 1997; Hamblin, 1973; Karttunen, 1977)—but virtually everyone who does systematic semantics accepts some version of the force/content distinction.

b. No, that's not true.

c. Sam said that too.

d. That's just what Sarah promised.

The uses of "that" in (20a–20d) are most naturally understood as pronouns referring back to the proposition expressed by (19). But the classical expressivist can't explain them this way, having denied that (19) *does* express a proposition. And it is unclear how the classical expressivist *can* explain them. Presumably (20a) will be understood as an expression of agreement in attitude with the first speaker, and (20b) as an expression of disagreement in attitude. But the fact that expressions of agreement and disagreement should take this form—with the surface appearance of propositional anaphora—needs explaining. Surely the simplest hypothesis is that there really is propositional anaphora is these cases. But that requires that "This fish is tasty" expresses a proposition, which the expressivist denies.

EMBEDDINGS

A more general problem for classical expressivism is how to extend its account of standalone sentences predicating "tasty" of some subject to an account of *arbitrary* sentences involving "tasty," including, for example,

(21) If that's tasty, he'll eat it.

(22) It will be tasty or the cook will give you your money back.

(23) That might be tasty.

(24) There are no tasty cookies in that jar.

All of these sentences employ "tasty" as a predicate, but in none of them is anything being *called* "tasty." What classical expressivists have done is to give an account of what one is doing in *calling* something "tasty." But this account does not extend to the uses of "tasty" in (21–24), in which nothing is being called tasty, and in which the speaker need not be expressing liking for anything at all.

The point is made forcefully in Geach (1960), who accuses expressivists of losing sight of Frege's distinction between *predicating F* of *a* and *asserting F* of *a*:

> In order that the use of a sentence in which "P" is predicated of a thing may count as an act of *calling* the thing "P," the sentence must be used assertively; and this is something quite distinct from the predication, for, as we have remarked, "P" may stil be predicated of the thing even in a sentence used nonassertively as a clause within another sentence. Hence, calling a thing "P" has to be explained in terms of predicating "P" of the thing, not the other way round. (Geach, 1960, p. 223)[21]

As Geach observes, the expressivist cannot meet the objection by saying that merely predicative uses of "tasty" have a different meaning than assertive uses, for then simple instances of modus ponens will be guilty of equivocation:

(25) a. If that is tasty$_{predicative}$, he will eat it.
b. That is tasty$_{assertive}$.
c. So, he will eat it.

The solution, Geach thinks, is to recognize that the two occurrences of "that is tasty" have a common *content* (the same truth conditions), though only the first is put forth with assertoric force. But that solution is not available to the expressivist, who does not think that "that is tasty" *has* a content or truth conditions.

Blackburn (1984) makes a valiant attempt to meet the objection head-on, by trying to show a classical expressivist account of ethical terms can be extended to an account of compound sentences like

(26) If gambling is wrong, then helping others gamble is wrong.[22]

If we analyze "gambling is wrong" as

(27) $B!(g)$
Boo for gambling!

and "helping others gamble is wrong" as

(28) $B!(hog)$
Boo for helping others gamble!

then we can analyze the conditional (26) as

[21] See also Geach (1965), Searle (1962), and Searle (1969, §6.2).

[22] For a fuller discussion, to which this one is indebted, see Kölbel (2002, ch. 4).

(29) $H!(|B!(g)|;|B!(hog)|)$
Hurray for the involvement of the attitude *boo for helping others gamble!* in the attitude *boo for gambling!*

(Here, as the gloss indicates, "$|B!(g)|$" denotes the attitude of being "boo for gambling" and $\ulcorner A;B\urcorner$ signifies the "involvement" of B in A.) This example gives the template for a compositional account on which all sentences including "tasty" would express attitudes—sometimes towards complexes of other attitudes, or the "involvement" of one attitude in another.

It is worth noting how ambitious a project this is. In order to handle conditional sentences, Blackburn had to find a relation between attitudes, "involvement," such that the endorsement of a conditional could be construed as the *hurraying* of an involvement relation. To handle disjunctions, he would need to find another relation between attitudes. To handle embeddings under modals, he would need still more machinery. Blackburn is thus committed to reconstructing virtually all of truth-conditional semantics along expressivist lines. This is a gigantic project, whose direction is just hinted at in Blackburn (1984).

Two further problems make things much worse for the project. The first problem is how to deal with *mixed cases*—compounds of evaluative and non-evaluative sentences (Hale, 1986). For example:

(30) If gambling is wrong, then Jim will avoid gambling.

(31) If gambling causes poverty, then gambling is wrong.

Our template construes conditionals containing "tasty" as expressions of attitudes about the involvement of one attitude in another. If we are to keep to this template here, we need to think of the non-evaluative sentences in the compound as designating attitudes. One might reasonably take them as designating beliefs. Then we could analyze (30) and (31) as

(32) $H!(|B!(g)|;|Bel!(jag)|)$
Hurray for the involvement of the attitude *believing that Jim will avoid gambling* in the attitude *boo for gambling!*

(33) $H!(|Bel!(gcp)|;|B!(g)|)$
Hurray for the involvement of the attitude *boo for gambling!* in

the attitude *believing gambling causes poverty*.[23]

But it is hard to see how we could accept these analyses and refuse to analyze

(34) If gambling causes poverty, then Jim will avoid gambling

as

(35) $H!(|Bel!(gcp)|;|Bel!(jag)|)$
Hurray for the involvement of the attitude *believing Jim will avoid gambling* in the attitude *believing gambling causes poverty*.[24]

And now we have an account on which *all* conditionals (and indeed, all other kinds of compound sentences) are to be understood as expressions of attitudes.[25] What started as an attempt to *supplement* truth-conditional semantics now threatens to supplant it entirely.

A second problem arises if we want expressivist accounts of different kinds of evaluative vocabulary. Presumably the kind of approval one expresses using "tasty" is very different from the kind of approval one expresses using "good" or "right" or "beautiful." A proper expressivist account ought to mark the difference between these attitudes, perhaps by subscripting the signs for "boo!" and "hurray!". We could then distinguish between the bishop's proclamation:

(36) $H!_{ethical}(m)$
Missionaries are good

and the cannibal's:

[23] See Blackburn (1984, p. 193) for an example like this. Blackburn analyzes it in terms of *knowing* rather than believing, but the difference won't matter for our purposes here.

[24] If one takes utterances of (34) to be ordinary assertions, and not expressions of approval, one will have no explanation of what would be wrong with accepting (30) and (31) while rejecting (34). But if one analyzes (34) as (35), one can give the same kind of explanation as Blackburn gives of the cogency of modus ponens inferences: there is a kind of practical incoherence in approving of the involvement of an attitude *B* in another attitude *A*, and of *C* in *B*, but not approving of the involvement of *C* in *A*.

[25] The expressivist might say that conditionals containing no evaluative vocabulary could be analyzed either as expressions of attitudes towards the involvement of attitudes or in standard truth-conditional fashion. But this would be to posit an ambiguity in conditionals for which there is no independent evidence.

(37) $H!_{gustatory}(m)$
Missionaries are tasty.

But which of these distinguished varieties of approbation is expressed by a conditional like

(38) If carrots are tasty, apples are tasty?

It would certainly be odd to analyze (38) as

(39) $H_{gustatory}!(|H!(c)|;|H!(a)|)$.

After all, combinations of attitudes don't taste like anything, so one can't really like their taste. But it also seems odd to analyze it as

(40) $H_{ethical}!(|H!(c)|;|H!(a)|)$.

(38) does not seem like an ethical judgement at all. The expressivist seems forced to posit a distinct kind of pro-attitude for this role. This attitude would have to be explained.

There is much more to be said about the prospects of working out classical expressivism along these lines.[26] But it should be clear that the project faces very serious difficulties, which can be solved, if at all, only at the price of a frightful complexity. All of these difficulties vanish if we treat "that is tasty" as expressing a content and having truth conditions. For then we get, "for free," an understanding of non-declarative sentences involving "tasty," of various mental attitudes that would be reported using "tasty," and of propositional anaphora. And we could appeal to existing truth-conditional accounts of disjunction, negation, conditionals, tenses, modals, quantifiers, and other forms of combination to understand the contribution "tasty" makes to the meaning of sentences in which it occurs embedded. Clearly, then, the only motivation for taking the hard road is a deep conviction that truth-conditional semantics cannot give an adequate account of the meanings of "tasty" or other evaluative terms.

I think that the expressivist's conviction is not wholly unwarranted. Standard paradigms for doing truth-conditional semantics *do* lack the resources for dealing with "tasty." But the solution, I will urge, is not to abandon the whole project of truth-conditional semantics, but to broaden it. Relativism, I will suggest, is expressivism done right.

[26] See Kölbel (2002, ch. 4), Hale (1986), Blackburn (1988).

1.3.3 Contemporary expressivism

Expressivists after Blackburn (1984) have tended to take a different tack. Instead of developing an entirely new kind of non-truth-conditional compositional meaning theory, they have sought to co-opt the machinery of truth-conditional semantics for their own purposes. Perhaps the clearest example of this trend is Gibbard (1990, ch. 5).[27]

Gibbard describes his semantics as a slight variation on standard possible-worlds semantics. Each sentence receives as its content (at a context) the set of factual-normative worlds at which it holds (Gibbard, 1990, p. 97). A factual-normative world $\langle w, n\rangle$ is, as Gibbard puts it, a "completely opinionated credal-normative state" (95): a combination of a complete representation of all the (non-normative) facts (w) and a complete normative system (n) that classes every action as permitted, obligatory, or forbidden in all possible factual circumstances.[28] A sentence *holds* at a factual-normative world $\langle w, n\rangle$ if the result of replacing "permitted," "obligatory," and "forbidden" with "permitted according to n," "obligatory according to n," "forbidden according to n" would be true if the facts were as w describes (96).

This apparatus gives a uniform treatment to normative, non-normative, and mixed sentences. They all get assigned sets of factual-normative worlds. The treatment of truth-functional compounds of these sentences is just as in standard possible-worlds semantics:[29]

$$\begin{aligned} [\![\phi \wedge \psi]\!] &= [\![\phi]\!] \cap [\![\psi]\!] \\ [\![\phi \vee \psi]\!] &= [\![\phi]\!] \cup [\![\psi]\!] \\ [\![\neg\phi]\!] &= U \setminus [\![\phi]\!] \end{aligned}$$

Quantifiers, modal operators, and other constructions can also be dealt with in standard ways that do not distinguish between evaluative and non-evaluative contents. So Gibbard has no trouble explaining

[27] See also Blackburn (1988), which adopts a version of Hintikka's semantics for deontic modals. I focus on Gibbard's version, because it is more elegant and closer to standard truth-conditional approaches. Blackburn and Gibbard agree that their approaches are "in essence the same" (Gibbard (2003, p. 83), Blackburn (1993, 195 n. 15)).

[28] In Gibbard (2003), Gibbard talks instead of "fact-prac" worlds, which are "maximally specific combinations of decisional and factual content" (47). The basic idea, though, is the same.

[29] Here $[\![\phi]\!]$ to denotes the semantic value of ϕ, and U denotes the set containing all factual-normative worlds.

the validity of Geach's modus ponens argument (25). The argument is valid for the same reason any modus ponens argument is valid: the conclusion holds at every factual-normative world at which the premises all hold.

This is a major departure from classical expressivism. Instead of rejecting the force/content distinction and undertaking the arduous task of giving a systematic account of meaning without it, Gibbard enriches the possible-worlds semanticist's formal representation of contents. In this framework, we can think of evaluative statements as assertions with contents that are *normatively variable* and factual statements as assertions with contents that are *factually variable*.[30]

> NORMATIVELY/FACTUALLY VARIABLE A set S of factual-normative worlds is *normatively variable* if there is some world w and systems of norms n_1, n_2 such that $\langle w, n_1\rangle \in S$ but $\langle w, n_2\rangle \notin S$, and *factually variable* if there is some system of norms n and worlds w_1, w_2 such that $\langle w_1, n\rangle \in S$ but $\langle w_2, n\rangle \notin S$.

Of course, many statements will be mixed—their contents will be both normatively and factually variable. An example is "Antony ought to attack," whose aptness depends both on the disposition of the troops and on norms of military strategy; even simpler examples can be derived by conjoining a normatively variable and a factually variable content.

Unlike the classical expressivist, who takes evaluative utterances to have a categorically different significance from factual assertions, Gibbard gives a unified account of the significance of asserting a content, be it purely factual, purely normative, or mixed. To assert that p is to express the state of mind of accepting or judging that p; and to judge that p is to "rule out" certain "combinations of normative systems with factual possibilities" (Gibbard, 1990, p. 99)—all those combinations not contained in the content of p.

Indeed, although Gibbard (1990, p. 8) clings to the classical expressivist tenet that normative talk is not "factual," even that is given up in Gibbard (2003, p. x):

[30] The terminology is mine, but the thought is Gibbard's: "I speak of a single attitude 'accepting' that one can take towards distinct items of content," yielding in one case a belief and in another a decision (Gibbard, 2003, p. 47).

> Does this mean that there are no facts of what I ought to do, no truths and falsehoods? Previously I thought so, but other philosophers challenged me to say what this denial could mean. In this book, I withdraw the denial and turn non-commital. In one sense there clearly are 'facts' of what a person ought to do, and in a sense of the word 'true' there is a fact of the matter. That's a minimalist sense, in which 'It's *true* that pain is to be avoided' just amounts to saying that pain is to be avoided—and likewise for 'It's a fact that'. Perhaps, as I used to think, there are senses too in which we can sensibly debate whether *ought* conclusions are true or false. Nothing in this book, though, depends on whether there is any such sense.

Gibbard is even willing to entertain talk of *beliefs* with evaluative contents: "I genuinely believe that pain is bad, and my expressivistic theory, filled out, explains what believing this consists in" (Gibbard, 2003, p. 183). What, then, is left of the classical expressivist's program, which started from a fundamental distinction between factual assertions and expressions of attitudes? Why shouldn't we see Gibbard's view as a form of contextualism (Dreier, 1999) or relativism (Kölbel, 2002, pp. 113–14), which relativizes the truth of sentences not just to possible states of affairs but to normative systems? What makes his view a form of *expressivism*?[31]

When Gibbard addresses this question, he tends to emphasize a point about order of explanation: that getting clear about the meaning of evaluative terms is a matter of understanding the states of mind these terms are used to express, not of understanding what things have to be like for these terms to apply to them:[32]

[31] It is remarkable that Gibbard now seems to reject *all* of the "clustered characteristics" by which Dreier (1999, p. 563) defines expressivism: "Notoriously, expressivists deny that normative judgments are true or false. They deny that normative judgments state propositions. And they deny that the function of normative judgments is to express the speaker's beliefs; instead, they express some non-cognitive attitude. Most important, I think, expressivists deny that normative judgments report or represent facts, or say how the world is."

[32] See also Dreier (1996, p. 29), who characterizes views as "norm expressivist" if they "attempt to explain the meaning of a normative predicate, *P*, not by saying what property it denotes, but indirectly, by saying what someone *does* by calling something *P*."

> The analysis is not directly of what it is for something to *be* rational, but of what it is for someone to *judge* that something is rational. We explain the term by saying what state of mind it expresses. (Gibbard, 1990, p. 8)
>
> The term "expressivism" I mean to cover any account of meanings that follows this indirect path: to explain the meaning of a term, explain what states of mind the term can be used to express. (Gibbard, 2003, p. 7)

The role of the machinery of factual-normative worlds, as Gibbard conceives it, is not to say say what things normative predicates are *true of*, but to give a systematic account of what states of minds are expressed by uses of arbitrarily complex sentences. (This is perhaps why Gibbard talks of sentences *holding at* factual-normative worlds, not *being true at* them.) To assert that p, Gibbard says, is to express the state of mind of *ruling out* certain combinations of normative systems and factual possibilities (those that are not consistent with the content of p).

The trouble with this way of characterizing expressivism is that it does not discriminate between evaluative discourse and anything else. As we observed above, Gibbard's account applies uniformly to evaluative, non-evaulative, and mixed discourse. Gibbard does not highlight this feature of his view, and often talks as if he is giving an account of "normative judgements" only. However, as he recognizes, very few of our everyday normative judgements are *purely* normative; most are mixed cases with contents that are both normatively and factually variable. Any account of normative judgement must be capable of applying to these as well. But then such an account must handle judgements with purely factual contents, too, as a limiting case. Gibbard's way of dealing with logical compounding leaves him with no clean way to give separate accounts of normative and purely factual judgements, since the conjunction of two mixed contents can be a purely factual content.

If an expressivist account is an account that explains the meaning of a term by telling us what states of mind arbitrarily complex sentences containing that term can be used to express, then Gibbard has given us an expressivist account not just of "rational" and "ought" but of "dog," "red," "or," and everything else. The account explains the meaning of

every word by telling us what mental state of "ruling out" combinations of normative systems and factual possibilities we would be expressing if we asserted arbitrarily complex sentences containing it.

One might concede this point and cast Gibbard as an expressivist about *all* discourse. This kind of global expressivism might still seem distinct from ordinary truth-conditional semantics, conceived as an attempt to explain the meanings of terms by saying what they denote or represent. In standard truth-conditional semantics, it might be thought, the assignment of denotations or satisfaction conditions *is* the explanation of meaning, whereas in Gibbard's account, the assignment of factual-normative world contents is just *machinery* for giving a systematic account of the *use* of arbitrarily complex sentences.

But the contrast being drawn here rests on an overly simplistic conception of truth-conditional semantics. Neither Donald Davidson nor David Lewis—perhaps the two most prominent promoters of truth-conditional semantics—thought of the assignment of denotations to expressions that figures in compositional semantics as *itself* the explanation of meaning. Both thought of these assignments as nothing more than machinery for systematizing truth conditions for whole sentences.[33] Moreover, both thought that assignments of truth conditions to whole sentences could count as an explanation of meaning only against a background theory that connected truth conditions with the use of sentences.[34] Gibbard's approach shares its structure, if not its detail, with theirs.

Thus, I suggest, there is room for a truth-conditional treatment of evaluative discourse that coincides with Gibbard's in its explanatory upshots, but emphasizes its continuity with standard semantic approaches rather than its differences from them. Such an approach would use a notion of truth relative to a world and evaluative standard. In this sense (and perhaps in others) it would be "relativist." It would eschew expressivist slogans that *identify* believing that Hen-of-the-Woods is tasty with liking its taste. Instead, it would give an account of the truth conditions of the proposition that Hen-of-the-Woods is

[33]See especially Davidson (1984, Essay 15) and Lewis (1980). The point will be discussed further in chapter 3, below.

[34]See Davidson (1984, Essays 9–10) and Lewis (1983). The point is made earlier in Dummett (1959), and will be discussed further in chapter 5, below.

tasty that explains why those who seeks to believe the truth will tend to believe this proposition just when they like the taste of Hen-of-the-Woods. How, exactly, such a view would differ in substance from Gibbard's is a topic to which we will return in chapter **??**.

1.4 Truth relativism

We have looked at three approaches to understanding the meaning of "tasty" and kindred words: objectivism, contextualism, and expressivism. Each approach is motivated by a genuine insight. The objectivist's insight is that there can be real disagreement about what is "tasty." The contextualist's insight is that "tasty" is context-sensitive in some way. And the expressivist's insight is that we can best understand "tasty" by looking at what one *does* in characterizing something as "tasty."

Relativism about truth is attractive because it offers to accommodate all three insights. The relativist says that Sam and Sal are disagreeing about the truth of a single proposition—that Hen-of-the-Woods is tasty. They are not talking past each other, as they would be on a simple contextualist account. Yet the truth of this proposition is not wholly independent of context. For its truth, according to the relativist, is assessment-dependent: as assessed from Sal's perspective, it is true, but as assessed from Sam's, it is false. It is true "for Sal," because Sal likes the taste of Hen-of-the-Woods, but not true "for Sam," because Sam doesn't like the taste. Thus, in accepting this proposition, Sal expresses the attitude of liking the taste of Hen-of-the-Woods, while in rejecting it, Sam expresses the attitude of not liking it.

Relativism also avoids the main problems that beset the three orthodox positions. The objectivist is committed to saying that at least one of the parties in a dispute of taste is making a mistake, and that forces us to attribute an implausible degree of unreflective chauvinism to users of "tasty." On the relativist view, by contrast, each party accepts a proposition that is true relative to his own perspective, and each is fully justified in doing so. The contextualist is committed to saying that the parties in disputes of taste are often talking past one another, disagreeing merely verbally. According to the relativist, the disagreement is real: there is a single proposition about which each

party takes the other to have the wrong view. The expressivist faces difficulties explaining the meaning of "tasty" in complex expressions. The relativist avoids this problem by operating inside the framework of truth-conditional semantics, with a standard distinction between force and content.

The difficulties facing the three standard approaches to "tasty" and other deficiently objective discourse, then, provide a strong motivation for seeing whether a coherent relativist view can be articulated (cf. Wright, 2001, p. 53).

However, truth relativism is widely regarded as an incoherent position, and so not even a contender. In chapter 2, we will survey some of the main objections to relativism about truth. It will be seen that the classical concerns about self-refutation are no threat to the deployment of relative truth envisaged here. But substantive concerns remain about the coherence of relative truth—and even about what it could *mean* to say that truth is relative.

Chapters 3–5 are devoted to answering these concerns. In chapter 3 I will argue that relativism about truth should be understood as the view that truth is *assessment-sensitive*. (Though there are other things one might mean by "relativism about truth," this is the one that is philosophically interesting and relevant to our motivating concerns.) Assessment sensitivity is understood by analogy with ordinary context sensitivity, or, as I call it, *use-sensitivity*. Just as the truth of uses of ordinary context-sensitive sentences and depends on features of the context in which they are used, so the truth of uses of assessment-sensitive sentences depends on features of the context in which they are assessed. Building on ideas of Lewis and Kaplan, I develop a framework that makes room for assessment-sensitivity.

In chapter 4, I show how propositions can fit into this framework, and I extend the notion of assessment sensitivity from sentences to propositions. This allows us to draw an important distinction between relativism about truth (which involves a commitment to assessment sensitivity) and nonindexical contextualism (which does not), and shows that taking propositional truth to be relevant to parameters besides possible worlds (and possibly times) is neither necessary nor sufficient for relativism about truth, in the sense articulated here.

In chapter 5, I address the substantive philosophical question that

remains: what does it *mean* to talk of truth relative to a context of assessment? I do this by explaining the theoretical role of assessment-relative truth in a truth-conditional semantic theory. The combined theory of chapters 3–5 allows us to formulate "relativist" semantic theories and derive from them substantive predictions about language use, so that they can be compared and evaluated against non-relativist alternatives. By doing this, I claim, we have "made sense of relative truth" and warded off apriori objections to its intelligibility.

Chapter 6 is devoted to the topic of disagreement, a central crux in the debate between relativists, indexical and nonindexical contextualists, classical expressivists, and objectivists. I distinguish several varieties or "levels" of disagreement and show how the issue between the different semantic approaches we have considered can be reduced to an issue about what kind of disagreement there is about matters of taste (or any other domain under discussion).

Once the framework of chapters 3–6 is in place, it becomes a broadly empirical question whether any of our thought and talk is best understood in terms of a relativist semantics. Chapters 7–**??** make the case for an affirmative answer. In chapter 7 I develop a simple relativist semantics for "tasty" and consider what the relativist can and should say about the interaction of "tasty" with tense, modality, quantifiers, attitude verbs, and factive verbs and adjectives. In chapter **??**, I argue that a relativist treatment of future contingent statements provides a satisfactory resolution of a two millenia-old deadlock between competing accounts. In chapter **??**, I consider how a relativist account of knowledge attributions might steer a middle course between contextualist and invariantist accounts. In chapter **??**, I make a case for a relativist treatment of epistemic modal claims, like *Joe might be in Boston*, over the standard contextualist and expressivist alternatives. In chapter **??** (based on joint work with Niko Kolodny), I criticize the view, widespread in the literature on ethics and practical reason, that "ought" is ambiguous between objective and subjective senses, arguing that a relativist treatment better accounts for our use of "ought," and better fits the purposes for which we use "ought." In chapter **??** (also based on work with Kolodny), I show how a relativist account of indicative conditionals emerges naturally out of existing views about indicatives and the relativist account of epistemic modals, and I show how this

account solves a paradox involving conditional obligation.

Finally, in chapter **??**, I compare the approach recommended here to several alternative approaches that resemble it in various ways, including Herman Cappelen's and Ernie Lepore's speech-act pluralism, the view (explored by Kit Fine and Iris Einheuser) that the facts themselves are relative, Allan Gibbard's sophisticated version of expressivism, and alternative versions of relativist semantics due to Peter Lasersohn, Andy Egan, François Recanati, Max Kölbel, Brian Weatherson, and Tamina Stephenson.

Part I

Foundations

2 *The standard objections*

The consensus among analytic philosophers is that relativism about truth is incoherent, or, at best, hopelessly confused. Here is a representative sampling of attitudes:

> That (total) relativism is inconsistent is a truism among philosophers. After all, is it not *obviously* contradictory to *hold* a point of view while at the same time holding that *no* point of view is more justified or right than any other? (Putnam, 1981, p. 119)
>
> . . . the contemporary consensus among analytic philosophers is that relativism is not just wrong, but too confused a position to be worth taking seriously. (Bennigson, 1999, p. 211)
>
> . . . the label 'relativistic' is widely regarded as pejorative, and few philosophers have been willing to mount an explicit defense of relativism. (Swoyer, 1982, p. 84)
>
> Relativism is even sillier than it at first appears. Indeed, if relativism were not so popular, it wouldn't be worth discussing at all. And even given its popularity, it isn't worth discussing for long. (Whyte, 1993, p. 112)
>
> . . . of all the conceptual options that have ever crossed the mind of the philosophical tribe, none has attracted quite the scorn and ridicule of the relativist. (Margolis, 1991, p. xiv)

Even Richard Rorty, who is often taken by analytic philosophers to be a relativist about truth (Boghossian, 2006), repudiates the doctrine:

> Truth is, to be sure, an absolute notion, in the following sense: 'true for me but not for you' and 'true in my culture but not in yours' are weird, pointless locutions. So is 'true then, but not now.' Whereas we often say 'good for this purpose, but not for that' and 'right in this situation, but not in that,' it seems pointlessly paradoxical to relativize truth to purposes or situations. (Rorty, 1998, p. 2)

Before we continue, then, let us look at some of the reasons philosophers have given for dismissing truth relativism out of hand, with an eye to establishing criteria of adequacy for a defensible relativism about truth.

2.1 WHAT ARE THE BEARERS OF RELATIVE TRUTH?

The relativist's thesis must be that the truth of *something* is relative: but what? Newton-Smith (1981, p. 35) argues that the relativist has no good answer. For the thesis that a *sentence* could be true in one social group or theory Ψ and false in another Θ is trivial, since the sentence could have different meanings in Ψ and Θ. Nobody would deny that sentences with different meanings can have different truth values. An interesting relativism, then, must "focus not on sentences but on what is expressed by a sentence"—a *proposition*. But the thesis that a single proposition can be true in Ψ and false in Θ is incoherent:[1]

> Let p be the proposition expressed by sentence 'S_1' in Ψ and by sentence 'S_2' in Θ. Could it be the case that p is true in Ψ and false in Θ? No, for it is a necessary condition for the sentence 'S_1' to express the same proposition as the sentence 'S_2' that the sentences have the same truth-conditions. To specify the truth-conditions of a sentence is to specify what would make it true and to specify what would make it false. If in fact 'S_1' and 'S_2' differ in truth-value, their truth-conditions must be different. If their truth-conditions differ they say different things—they say that different conditions obtain—and hence they do not express the same proposition. Thus if we focus on propositions we cannot find a proposition expressed by a sentence 'S_1' in Θ and by a sentence 'S_2' in Ψ which is true in the one case and false in the other.[2] Newton-Smith (1981, p. 35)

The thought Newton-Smith is expressing is surprisingly common, and it has led some relativists to put their position as a thesis about the truth of *utterances* or *assertions*. This, I think, is misguided (as

[1] For similar arguments, see also Husserl (2001, p. 79), Newton-Smith (1982, pp. 107–8), Swoyer (1982, p. 105), Burke (1979, p. 204), L. Stevenson (1988, pp. 282–3). For critical discussion, see White (1986, p. 332), Hales (1997*a*, p. 39), and Kölbel (2002, pp. 119–122).

[2] Newton-Smith seems to have inadvertently reversed "Θ" and "Ψ" here.

I will argue in chapter 3). It is also unnecessary, because Newton-Smith's objection is weak. But a proper discussion will be easier once we have defined some concepts that will allow us to disambiguate Newton-Smith's talk of "truth-conditions" and "in Θ" (§ 4.8.1). For now let us just register that a relativist about truth owes an account of propositions that makes it possible to say that the truth of a *proposition* is relative.

2.2 Self-refutation

Perhaps the most famous charge against relativism about truth is that it is self-refuting. This charge is leveled against a very strong kind of global relativism: the view that *all* truths are true merely relatively, and that nothing is true absolutely. In its simplest form, the refutation takes the form of a dilemma. If the global relativist says that relativism is true for everyone, then she is acknowledging that there is at least one non-relative truth, and this contradicts her thesis of global relativism. On the other hand, if she concedes that relativism is false for someone, or equivalently that absolutism is true for someone, then...

Then what? It is usually conceded that there is no real contradiction in the relativist's holding that relativism is true for herself, although there are others for whom it is not true. Plato's Socrates is sometimes read as finding a real contradiction in a relativist position he attributes to Protagoras, but if so his argument cheats by dropping the crucial "for x" qualifiers at the final stage.[3] Hales (1997*a*) shows how global relativism is self-refuting if we assume a principle analogous to the S5 axiom of modal logic: if it is relatively true that it is absolutely true that p, then it is absolutely true that p.[4] But it is not clear why

[3] See M. F. Burnyeat (1976*b*, pp. 174–5), who cites Grote, Runciman, Sayre, and Vlastos for the charge. Burnyeat tries ingeniously to find a more subtle argument in the text, but Fine (1998*a*) is probably right that Plato, like Sextus and all other ancient commentators (cf. M. F. Burnyeat, 1976*a*), takes Protagoras to be a subjectivist rather than a relativist. Subjectivism, the view that everything that appears to be the case is true (absolutely), is just the sort of view that would require the radical Heraclitean metaphysics of temporary person-relative appearance-objects attributed to Protagoras in the first part of the dialogue. And it is cogently refuted by the Socratic argument that seems feeble as a response to relativism.

[4] The S5 axiom says that if it is possible that it is necessary that p, then it is necessary that p.

the global relativist should accept this principle.[5] Most commentators agree that the problem with the second horn of the relativist's dilemma is more subtle than outright inconsistency.

2.2.1 Pragmatic inconsistency?

It is often suggested that the global relativist's claim is pragmatically self-refuting, in the sense that the content of the relativist's claim is inconsistent with a true description of what the relativist would be doing in asserting it. (A noncontroversial example of a pragmatically self-refuting claim would be *that I am not asserting anything*.) The charge is spelled out clearly by John Passmore:

> [E]ven if we can make some sense of the description of p as 'being true for x' ... Protagoras is still asserting that 'p is true for x' and 'p is not true for y'; these propositions he is taking to be true. It has to be true not only for x but for everybody that 'p is true for x' since this exactly what is involved in asserting that 'man is the measure of all things.'
>
> The fundamental criticism of Protagoras can now be put thus: to engage in discourse at all he has to assert that something is the case. (Passmore, 1961, p. 67)

Passmore's idea is that to assert something is to put it forward as true, not just for oneself, but for everyone—true absolutely. So, while the relativist's thesis entails that it is not true absolutely, in asserting it the relativist is putting it forward as true absolutely. The thesis is therefore pragmatically self-refuting. As Myles Burnyeat puts the point: "No amount of maneuvering with his relativizing qualifiers will extricate Protagoras from the commitment to truth absolute which is bound up with the very act of assertion" (M. F. Burnyeat, 1976*b*, p. 195).

But why should the relativist concede that an assertion is a "commitment to truth absolute"? Why shouldn't the relativist say that in asserting that p, one is putting forward p as *relatively* true—perhaps

[5] See Shogenji (1997) for this criticism and Hales (1997*b*) for a reply. Bennigson (1999) defends the consistency of global relativism by describing a model in which relative truth is truth in some accessible framework, absolute truth is truth in every accessible framework, and the absolutist's framework is accessible from the relativist's framework, but not vice versa. This amounts to a rejection of Hales' S5-like premise, since the S5 axiom requires that accessibility be transitive.

as true for oneself?[6] To be sure, there is a substantive task here for the relativist, who must reassure us that we can make good sense of assertion (claim-making) without appealing to "absolute truth." But it is certainly not obvious that the position to which the relativist has been driven here is an untenable one.

2.2.2 Regress of formulation?

Indeed, many who pursue a self-refutation argument have conceded that it must take a more subtle form (M. F. Burnyeat (1976*b*, pp. 192–3), Putnam (1981, pp. 120–1), Vallicella (1984, pp. 462–3), Lockie (2003, p. 331), Boghossian (2006, p. 54)). The real problem, they think, is that the relativist faces a kind of regress in formulating her own position. When the relativist says,

(41) I'm only putting my thesis forward as true *for me*,

the objector can ask whether *this* claim is being put forward as true absolutely. If the relativist says yes, then she has conceded that there is at least one absolute truth, and stands refuted. If she says no,

(42) I'm only putting (41) forward as true *for me*,

then the procedure can be iterated, and so on indefinitely. And this is supposed to spell doom for the relativist's position. As Putnam explains:

> A *total* relativist would have to say that whether or not X is true *relative* to P is *itself* relative. At this point our grasp on what the position even means begins to wobble, as Plato observed. (Putnam, 1981, p. 121)

But what, exactly, is problematic about the relativist's willingness to say of *all* of her assertions—even metatheoretic ones like (41)—that

[6] Cf. Kölbel (2002, p. 123): "The relativist might concede that asserting something does constitute certain commitments, such as the obligation to state reasons for what one has asserted if asked to do so, to defend what one has asserted if challenged, and to retract one's assertion if one is unable to defend it against challenge. But he or she will will deny that commitment to the absolute truth of what has been asserted is among the commitments constituted by an assertion." As Bennigson (1999, p. 215) notes, the relativist could even hang on to the slogan that to assert something is to put it forward as true, adding that "to say that a sentence is true is to say only that it is true in the framework in which the discussion is occurring."

she is only putting them forth as true for herself, or true relative to her own framework or perspective?

M. F. Burnyeat (1976*b*, p. 193) suggests that the problem lies in the complexity of the propositions to which we will be led if we iterate the move from (41) to (42):

> Protagoras, as Socrates keeps saying, is a clever fellow, but he is not so clever that there is no limit to the complexity of the propositions he can understand and so judge to be true. Therefore, the relativist prefix 'It is true for Protagoras that…,' unlike the absolute prefix, admits of only limited reiteration.

But it is hard to see how this objection hits home. The relativist need only move from stage k of the regress to stage $k + 1$ if an intelligible question has been raised about whether stage k has been put forth as true absolutely or merely as true relatively. But surely this question is intelligible if and only if its possible answers are; they are of equal complexity. So, if stage $k + 1$ is unintelligible because of its complexity, so is the question that would require the relativist to produce it, and we can rest content at stage k (cf. Bennigson, 1999, pp. 224–6).

Boghossian (2006, p. 56) gives a somewhat different diagnosis. According to Boghossian, the relativist holds that

> [i]f our factual judgements are to have any prospect of being true, we must not construe utterances of the form "p" as expressing the claim p but rather as expressing the claim *According to a theory, T, that we accept, p*. (52)

But of course it would be odd for the relativist to hold that there are absolute facts about what theories say (and hence, presumably, about the contents of minds), but about nothing else. So claims about what theories say also have to be understood as merely claims about what theories say, and a regress ensues:

> The upshot is that the fact-relativist is committed to the view that the only facts there are, are infinitary facts of the form:
>
> > According to a theory that we accept, there is a theory that we accept and according to this latter theory, there is a theory we accept and … there have been dinosaurs.
>
> But it is absurd to propose that, in order for our utterances to have any prospect of being true, what we must mean by them

> are infinitary propositions that we could neither express nor understand.

This is indeed absurd, but Boghossian's argument that the relativist is committed to it depends on a tendentious characterization of the relativist's position. Boghossian's relativist is really a kind of contextualist: she takes a speaker who utters "snow is white" to have asserted that according to her world-theory, snow is white.[7] But the relativist need not, and should not, hold that to put p forward as true for oneself is to put forward the claim *that p is true for oneself*. The point of "for oneself" is not to characterize the *content* that is asserted, but to characterize what the relativist is *doing* in making her assertion: putting its content forward as *true for herself*.

Putnam (1981) sees that the infinite regress argument cannot bear much weight (120), but he thinks that Plato's argument points dimly towards an argument he finds in Wittgenstein:

> The argument is that the relativist cannot, in the end, make any sense of the distinction between *being right* and *thinking he is right*; and that means that there is, in the end, no difference between *asserting* or *thinking*, on the one hand, and *making noises (or producing mental images)* on the other. But this means that (on this conception) I am not a *thinker* at all but a *mere* animal. To hold such a view is to commit a sort of mental suicide. (122)

But this argument works (if it works at all) only against an extreme subjectivist relativist—one who holds that "p is true for X" is equivalent to "X believes that p." If X can be wrong about what is true for X—if it can be false for X that p is true for X, even though X believes p—then it seems we *do* have a distinction between X's being right and X's thinking she is right.

[7] Boghossian models his version of truth relativism on Gilbert Harman's version of moral relativism (Harman, 1975), which is essentially a form of contextualism about terms of moral evaluation. As Kölbel (2002, p. 119) observes, "It can be shown that no global relativist can accept Harman's view that relativity is always a matter of logical form and empty argument places. For if he accepted that, any predicate would have an indefinite number of argument places. . . . So, global relativists must have a different view of what is involved in being relative to some parameter whether some x is F." See also Wright (forthcoming).

It is not clear, then, that any of the "self-refutation" arguments against global relativism are compelling, though they do raise questions that the relativist must answer. What is it, exactly, to put something forward as true relatively? By what right does such an act count as an assertion or judgement? What is it for a proposition to be true for X, if not for X to believe it? These are important questions, but the objections have given no principled reason to think that the relativist cannot answer them satisfactorily.

2.2.3 Is local relativism immune?

A common relativist response to the self-refutation argument is to point out that it targets only a radical *global* relativism according to which nothing is true absolutely (Nozick (2001, p. 15), MacFarlane (2005*a*, 338 n. 19)). However, many relativists about truth are *local* relativists. They think that relative truth is characteristic only of certain kinds of discourse—for example, discourse about taste. Indeed, the sorts of motivations for relativism about truth that we have been examining *require* that relativism not be global. For the point is to explain the differences between "deficiently objective" discourse and "robustly objective" discourse by taking the former, but not the latter, to be true only "relatively."

Global relativism about truth tends to be motivated by an entirely different set of considerations. It comes out of the historicized and relativized versions of Kantianism that flourished in the nineteenth century (Meiland, 1977, p. 568). Kant's idea was that all objects of our cognition are partly constituted by the innate and *a priori* modes in which they must be organized in order to be intelligible to thinkers. Relativism emerges when the object-constituting modes of organization are allowed to be contingent and *a posteriori*. Thus, according to Swoyer (1982, p. 86),

> We can think of the general argument for relativism as involving two stages: first, a defense of a constructivistic epistemology according to which the knower somehow organizes or constitutes what is known and, second, an argument that there is no uniquely correct way of doing this.
>
> The first stage owes an obvious debt to Kant.

Global relativism, so motivated, tends to see truth as relativized to such things as "conceptual schemes, conceptual frameworks, linguistic

frameworks, forms of life, modes of discourse or thought, Weltanschaungen, disciplinary matrices, paradigms, constellations of absolute presuppositions, points of view, perspectives, or worlds" (Meiland and Krausz, 1982, p. 84).

Historically, self-refutation arguments have been targeted at global relativisms. And it is not clear how a self-refutation argument against local relativism would even get going: the local relativist can simply say that she is putting her thesis forward as true absolutely, grasping the horn of the dilemma that was not available to the global relativist. There is no inconsistency or even pragmatic incoherence in saying, for example, that it is absolutely true that claims of taste are true only relative to judges or standards of taste.

However, it would be a mistake for the local relativist to be too complacent. For the real problem the self-refutation argument raises for the global relativist—explaining how putting something forward as true not absolutely, but only relatively, could count as an assertion—is equally pressing for the local relativist. The local relativist could say that there are two kinds of assertion—putting forward as absolutely true and putting foward as relatively true—and that the relativist thesis itself is being asserted in the first way. But there is no evident advantage to saying this. If we can understand assertion as putting forward as relatively true at all, we can understand an assertion of the relativist thesis in this way. And to distinguish two kinds of assertion, one appropriate for "robustly objective" claims, one for "deficiently objective" claims, would be to surrender one of the relativist's chief advantages over expressivism: the promise of explaining the differences between these kinds of claims by appealing to differences in their contents, rather than their illocutionary forces. The local relativist should concede, then, that even the thesis of relativism is being put forth as relatively true. (This does not rule out saying that the thesis *is* true absolutely; the point is just that to assert it is not to present it as true absolutely.)

The upshot of the self-refutation argument, then, is that the relativist about truth, whether global or local, owes an account of assertion that does not assume that the asserted contents have absolute truth values. We will take up this task in chapter 5.

2.3 THE EQUIVALENCE SCHEMA

It is sometimes argued that relativism about truth is incompatible with the

> EQUIVALENCE SCHEMA The proposition that Φ is true iff Φ.

Although the Equivalence Schema is implicated in the Liar Paradox, and may need to be qualified or restricted in some way, it is generally regarded as fundamental to our use of the truth predicate. The point can be made intuitively: it would be incoherent to say, for example, that it is true that dogs bark, while denying that dogs bark, or to say that dogs bark while denying that it is true that they do. But the Equivalence Schema can also be motivated on logical and expressive grounds. English and other natural languages do not allow quantification into sentence position: the grammatical position occupied by 'Φ' in the Equivalence Schema. So, we cannot express our agreement with everything Billy asserted by saying,

(43) For all P, if Billy asserted that P, P.

(Note that 'P' here occupies a place that could be filled by a sentence, and could not be filled by a proper name.) Natural languages get around this apparent limitation in expressive power by providing a way to simulate quantification into sentence position using ordinary quantification over objects:

(44) For all propositions x, if Billy asserted x, x is true.

(Note that 'x' here occupies a place that could be filled by a proper name.) But if (44) is to do the work that (43) would do, we must be able to move from (44) and

(45) Billy asserted that snow is white

to

(46) Snow is white.

And what we need for that is precisely an instance of the Equivalence Schema:

(47) (The proposition) that snow is white is true iff snow is white.

The idea that the Equivalence Schema precludes relativism about truth is put forcefully by Frank Ramsey in his posthumously published article "The Nature of Truth" (Ramsey, 2001). Ramsey criticizes philosophers who "produce definitions of truth according to which the earth can be round without its being true that it is round" (441), noting specifically that "... according to William James a pragmatist could think both that Shakespeare's plays were written by Bacon and that someone else's opinion that Shakespeare wrote them might be perfectly true 'for him.'" (445–5 n. 12, citing James (1909, p. 274) = James (1978, p. 313)).[8]

It is not clear, though, why saying that a proposition is true "for him" should debar the pragmatist from saying that the same proposition is "not true," as the Equivalence Schema demands. A relativist could acknowledge that the truth predicate in ordinary use is a monadic predicate that behaves as the Equivalence Schema describes, while promoting the theoretical utility of a relativized (dyadic) truth predicate (cf. Kölbel, 2008*b*; Unwin, 1987, pp. 304–5). Indeed, it would be natural for the relativist to say that, just as

(48) Hen-of-the-Woods is tasty

is true relative to some judges or standards and false relative to others, so

(49) The proposition that Hen-of-the-Woods is tasty is true

is relatively true in exactly the same way. In fact, the relativist *must* say this if she is to preserve the Equivalence Schema for the monadic truth predicate.

Of course, this concession raises serious questions about how the monadic truth predicate and the dyadic truth predicate are related, and why the latter deserves to be called a truth predicate at all. Suppose we use "true" to express the monadic property which obeys the Equivalence Schema and "True for x" to express the relativist's relational property. Then, as Fox (1994, p. 73) notes, either the Equivalence Schema holds for "True for x" no matter what the value of x, or it does not. If it does—that is, if

[8] Interestingly, even some relativists about truth agree that their view is incompatible with the equivalence schema (Nozick, 2001, p. 41).

(50) $\forall x$(the proposition that Φ is True for x iff Φ)

is a valid schema, then the relativization to x looks like an idle wheel. But if it does not, "the most cogent of arguments that Truth needs relativizing could not carry over to an argument that truth does."

Thus it is crucial for the relativist to give an adequate account of the relation between her relativized notion of truth and the monadic predicate we use in ordinary talk. We will return to this issue in § 4.7.

2.4 WHAT DOES IT MEAN?

Perhaps the most pressing objection to relativism about truth is that it is not clear what it could *mean* to call a proposition "true for Sal" or "true relative to Sal's tastes."

The problem is not that "true for x" doesn't have a use in non-philosophical English. It is often used to specify the domain for which a generalization holds. For example, in

(51) ...while the doors to high civil, military and academic office have been opened to merit for members of other communities, this has not been *true for Muslims*. (Shissler, 2003, p. 153)

the generalization

(52) the doors to high civil, military and academic office have been opened to merit for members of their communities

is being said not to hold for Muslims. Interestingly, "this" in (51) denotes not the proposition expressed immediately before it,

(53) the doors to high civil, military and academic office have been opened to merit for members of other [than Muslim] communities,

but rather the generalization expressed by (52)—a generalization with unspecified domain that has been abstracted from (53).[9] Whether this generalization deserves to be called a proposition is not something we will need to settle here, because I think it can be agreed that its

[9]This abstraction must be triggered by the predicate "true for Muslims," since if we substitute (say) "widely recognized," "this" denotes the proposition expressed by (53).

"relative truth" is not relative truth in any very interesting sense: what is expressed by (51) is simply that the doors to high civil, mililitary and academic offices have been opened to merit for members of non-Muslim communities, but not for Muslims.

Here are some further examples of the same phenomenon:

(54) Some of those religious leaders commute between homelands and host countries, not only providing religious continuity but also maintaining political contacts. That has been particularly *true for Turks, Pakistanis, Iranians, Sikhs, Jews, and Palestinians*. (Sheffer, 2003, p. 167)

(55) About half of all fatal heart attacks occur in women, but a woman who has a heart attack has a twofold risk of dying within the first two weeks, compared to her male counterpart. This is especially *true for women under the age of forty-five*. (Zaret and Shubak-Sharpe, 2007, p. 170)

(56) It is well known that certain human subjects are especially resistant to the gas and I have frequently found this to be *true for dogs*. (Jackson, 1917, p. 70)

(57) "The pianos used by Mozart, Beethoven or Chopin were radically different from the large, loud black instruments found in all modern concert halls," Moroney said. "The same is *true for violins*." (Maclay, 2001)

(58) Not only was it useless to try to derive the satisfaction she needed from another individual through subordinating her needs to that person's, but even attempting to do so was destroying her. Generalizing her insights, she concluded that what was *true for her* was *true for all people*. (Fellman, 2008, p. 62)

"True for" also seems to have what we might call an *intentional* use: to call something "true for x" is sometimes just to say that x *takes* it to be true, or that it is "true in x's book." The work here is being done by "for," not "true," since in a similar spirit we can say (of Sarah the creationist)

(59) For Sarah, that fossil is less than 5000 years old.

or (of John the color-blind man)

(60) For John, those socks are the same color.

or (of Elroy the imaginative child)

(61) For Elroy, ant mounds are space stations.

This use of "true for" is no help to the relativist, as Meiland and Krausz observe:

> If all that were meant by saying a belief is true for Jones is that Jones holds that belief, then every belief that Jones holds would be true for Jones. But the relativist rejects this notion of relative truth; he or she takes the notion of relative truth more seriously than this. "Relative truth" is a form of truth; the expression "relative truth" is not a name for something bearing little relation to our ordinary conception of truth. And just as our ordinary conception of truth allows a person to hold beliefs which are false, so too the notion of relative truth must allow an individual to hold beliefs which are false *for him* or *her*. If it were not possible for a person to hold beliefs which were false for him or her, then the notion of relative truth would be superfluous . . . (Meiland and Krausz (1982, p. 4); cf. Fox (1994, pp. 70–1), Vallicella (1984, p. 454), Swoyer (1982, p. 94))

Thus the relativist cannot claim to be explicating a relational truth locution that is already in use in natural speech. If she uses a relativized truth predicate, she must explain what it means. And she must do so in a way that makes it clear why the relativized predicate she is explaining is a relativized *truth* predicate, and not something else entirely. This is, I think, the *principal* challenge for truth relativism, and the one that the existing literature has made least progress in answering.

Extant answers typically start with one of the traditional "theories of truth": correspondence, pragmatic, or epistemic. (Deflationary theories, according to which there is nothing to be said about truth beyond the Equivalence Schema, are no help to the relativist, for reasons canvassed in the previous section.) They then attempt to show that these theories, when properly understood, lead to the idea that truth is a relative property. The point is perhaps easiest to see with a pragmatic theory of truth. If truth is, as James says, "only the expedient in the

way of our thinking" (James, 1978, p. 106), then relativism about truth is just the plausible thesis that what is expedient for one person to think need not be expedient for another.[10] James seems to accept this thesis: "... in any concrete account of what is denoted by 'truth' in human life, the word can only be used relatively to some particular trower" (James, 1978, p. 313).

Epistemic theories of truth, which call true what a community of idealized enquirers would be justified in believing, also give a clear sense to the idea that truth might be relative. Bennigson (1999, p. 213) motivates his relativism in this way:

> Begin with an epistemic account of truth as some sort of idealization of rational acceptability: true sentences are those which disinterested inquirers would assent to under ideal conditions, or at some idealized 'limit of inquiry.' The relativist simply adds that different communities of inquirers, starting from different sets of assumptions about what is plausible, noteworthy, explanatory, etc., might approach different limits. Thus, on the appropriate epistemic conception of truth, conflicting conclusions could be true for different communities.

Even if we suppose (rather implausibly) that any two communities of idealized enquirers would have access to the same observations and experimental results, an epistemic conception of truth tends toward relativism. Most philosophers have abandoned Carnap's idea that the relation *evidence e confirms proposition p* can be spelled out in formal logical terms. Whether a given proposition is supported by a given body of evidence, and how strongly, depends on facts about the background of inquiry—for Goodman (1979), the relative "entrenchment" of predicates, for Bayesians, prior probabilities and a background corpus.[11] If two communities of idealized enquirers differed in these factors, then even if they went on to have all the same observations

[10] How plausible this is will depend on how, exactly, one spells out the pragmatist's slogan. James himself does not limit the kinds of expediency that might be at issue (he says "Expedient in almost any fashion"), and he acknowledges that idiosyncracies of taste may play a role: "Truth in science is what gives us the maximum possible sum of satisfactions, taste included, but consistency both with previous truth and with novel fact is always the most imperious claimant" (104).

[11] See Fitelson (2005) for a nice survey.

and perform all the same experiments, they might diverge in their justified beliefs at "the end of inquiry."

Although they give comfort to the relativist, though, pragmatic and epistemic theories of truth are not very plausible. Surely it is coherent to suppose that there are truths that even idealized enquirers could not come to know, and truths that it would not be expedient (in James's very broad sense) to believe.[12] Moreover, pragmatic and epistemic theories would support a diffuse *global* relativism, not the kind of targeted local relativism we might use to explain the characteristic features of "deficiently objective" discourse.

It is not surprising, then, that some relativists have sought to make sense of their doctrine in the framework of a correspondence theory of truth. Jack Meiland proposes that "ϕ is true for Jones" means "ϕ corresponds to reality for Jones." To the obvious objection—what does "corresponds to reality for Jones" mean?—Meiland has this reply:

> . . . although this question is embarrassing in the sense that it is difficult for the relativist to give any useful answer to it, nevertheless the relativist is in no worse a position than the absolutist at this point. . . . relativism is not to be faulted for being unable to give an account of that which the absolutist cannot give an account of in his own position either. (Meiland, 1977, p. 580)

Meiland is right to insist that the relativist not be held to a higher standard in explicating truth than the absolutist. And he is right that "corresponds to reality for Jones" is no less intelligible than "corresponds to reality" (assuming nothing more is said to explicate *that*). However, to say that truth is "correspondence to the facts" is, at best, only to give the schema for an explication of truth. By itself it provides no illumination. And neither does saying that truth for Jones is correspondence to reality for Jones.

Though Meiland's specific account is unilluminating, I think the strategy he pursues is a promising one: look at the best non-relativist explication of truth, and explicate relative truth in a similar way, using similar materials. We will return to this project in chapter 5.

* * *

[12] For an in-depth discussion, see Künne (2003, ch. 7).

Let us survey the ground we have covered so far. In chapters 1, we examined three standard ways of explaining "deficiently objective" discourse, and found them wanting. Objectivism is forced to attribute a high degree of unreflective chauvinism to speakers. Contextualism fails to explain our perceptions of disagreement among taste claims. Expressivism has difficulty explaining the meanings of taste predicates when they occur embedded in other contexts. Truth relativism promises a more adequate account, one that integrates the insights behind the other three while avoiding their weaknesses.

However, truth relativism is often dismissed as confused or incoherent. In this chapter, we have looked at some of the standard objections to truth relativism. Under closer examination, none of these objections look like knock-down arguments. Still, they do point to real problems that a relativist must address. The relativist needs an account of propositions that allows them to be "merely relatively true." She needs to explain how the ordinary monadic truth predicate is related to her relativized truth predicate. She needs to explain what one is doing in asserting propositions, if one is not putting them forward as true absolutely. And she needs to say more about what her relativized "true for" or "true at" *means*.

The aim of the next three chapters is to state the truth relativist's position clearly enough that all of these questions, and more, can be answered. In chapters 3 and 4, I argue that truth relativism is best understood as the view that sentences or propositions may be *assessment-sensitive*—they may have different truth values relative to different contexts of assessment. Building on David Lewis's and David Kaplan's work on the semantics of indexicals, I articulate a framework that allows us to distinguish assessment sensitivity from use sensitivity (that is, ordinary context sensitivity), and from the relativity of truth to indices or circumstances of evaluation. It is only when these distinctions have been made that we can isolate truth relativism and see what is genuinely distinctive of it.

In chapter 5, I turn to the philosophical task of understanding what it *means* to talk of truth relative to a context of assessment. Adopting Meiland's general strategy, I start with an explication of truth relative to a context of use—the target notion of a formal truth-conditional semantics—and show how with a few small tweaks it may

be transformed into an explication of truth relative to a context of use and context of assessment. The resulting account allows us to see the *practical* difference between assessment-sensitive and non-assessment-sensitive semantic theories, and tells us precisely what to look for in adjudicating between them. It thus renders relativism about truth intelligible and gives particular relativist proposals determinate empirical content.

3 Assessment Sensitivity

Most of the literature on truth relativism concerns either motivations for relativizing truth or arguments against the coherence of truth relativism. Comparatively little attention has been given to saying with precision what it *is* to be a truth relativist. Discussions have typically contented themselves with sloppy formulations, or precise ones that do not distinguish relativist theories from contextualist ones. The aim of this chapter is to say precisely what kinds of views count as forms of truth relativism. As Meiland (1977, p. 568) rightly says:

> Perhaps truth is relative; perhaps not. But I think that we cannot decide whether or not truth is relative until we first determine what "relative truth" might be.

The characterization of truth relativism proferred below will count as relativist some views that others would not, and fail to count as relativist some views that others would. The project is not to give a general account of the meaning of "relative truth" as that phrase is used in philosophical discourse. It is used in many ways. Nor is it to insist that there is only one legitimate or useful thing to mean by this phrase. The characterization is offered as an *explication* (in Rudolf Carnap's sense) of philosophical talk of "relative truth." If, after my explication, some readers prefer to continue using the phrase "relative truth" in some other way, that is fine. Not much hangs on the words, so long as the concept to which I would prefer to attach them—assessment sensitivity—is clearly grasped.

3.1 Characterizing relativism

One might think that being a relativist about truth is just a matter of relativizing truth to some parameter. But it is not that simple. Many

relativizations of truth are entirely orthodox.

3.1.1 Sentences

Considered by itself, in abstraction from any particular context of use, the sentence

(62) I have been to China.

cannot be said to be true or false. It can be used to say something with either truth value. For certain purposes, we might find it useful to assign (62) truth values relative to possible *contexts of use*, which determine a denotation for "I" (the speaker) and a reference time (the time of use). This way of relativizing truth is familiar from David Kaplan's pioneering work on indexicals (Kaplan, 1989). But nobody would say that Kaplan is a "relativist about truth" in any philosophically interesting sense. This relativization simply registers a fact obvious to everyone—that in general, whether sentences express truths or falsehoods depends on the settings in which they are used.

Compositional semantics demands other kinds of relativization of sentence truth, as well. Consider the problem of giving systematic truth conditions for quantified sentences, like

(63) For all integers x, there exists an integer y such that $x + y = 0$.

If we think of this sentence in the standard way, as the result of combining a quantifier "For all integers x" with an open sentence

(64) There exists an integer y such that $x + y = 0$,

then a compositional semantics ought to give truth conditions for (63) as a function of the truth conditions of (64). But (64) is an open sentence; there are no "conditions" under which (64) is true simpliciter, only conditions under which it is true for some value of x or another. Tarski's solution to this problem was to recursively define *truth on an assignment of values to the variables* rather than truth simpliciter (Tarski, 1935, 1983).[1] Thus, for example, the clause governing the universal quantifier looks like this:

[1] Tarski encoded his assignments as infinite sequences of values, and so talked of "truth on a sequence," but the decision to use a sequence rather than a function is just a technical one.

(65) $\ulcorner \forall \alpha \phi \urcorner$ is true on assignment a iff for every assignment a' that differs from a at most in the value it gives to α, ϕ is true on a'.

Note that even the clauses for the truth-functional connectives must be stated in terms of truth on an assignment:

(66) $\ulcorner \neg \phi \urcorner$ is true on assignment a iff ϕ is not true on a.

For these connectives may operate on open formulas, as is the case in

(67) $\forall x \neg (x < 0)$.

So truth is relativized to assignments for all formulas. But nobody would call Tarski a "relativist about truth" on this account.

Why aren't these relativizations of truth philosophically problematic? For relativization to an assignment, the answer is clear: this relativization is just a technical device, not something we need to make sense of independently of its role in systematizing absolute truth values. At the end of the day, what we care about is truth, not truth on an assignment. So our recursive definition of truth on an assignment for arbitrary formulas is of interest to us only because truth simpliciter can be defined in terms of truth on an assignment:

(68) If ϕ is a sentence, then ϕ is true iff ϕ is true on every assignment.[2]

Because the role played by truth on an assignment is a purely technical one, we could use different terminology without changing the theory in any important way. Instead of talking of "truth on an assignment," for example, we could define a valuation function v that maps sentence/assignment pairs to 0 or 1:

(65′)

$$v(\ulcorner \forall \alpha \phi \urcorner, a) = \begin{cases} 1 & \text{if for every assignment } a' \text{ that differs} \\ & \text{from } a \text{ at most in the value it gives to } \alpha, \\ & v(\phi, a') = 1 \\ 0 & \text{otherwise} \end{cases}$$

[2]Equivalently, "on *some* assignment," or even "on assignment a_0," since a sentence—a formula with no free variables—will have the same truth value on every assignment.

(66′)

$$v(\ulcorner\neg\phi\urcorner, a) = \begin{cases} 1 & \text{if } v(\phi, a) = 0 \\ 0 & \text{otherwise} \end{cases}$$

And then, at the end:

(68′) If ϕ is a sentence, then ϕ is true iff for every assignment a, $v(\phi, a) = 1$.

The recursive definition of v does exactly the same work as the recursive definition of truth on an assignment. Talk of truth relative an assignment, then, is consistent with holding that truth in the philosophically interesting sense is absolute.

What about truth relative to a context of use? One might try a parallel strategy here, arguing that we talk of truth at a context for sentences only as a technical device for systematizing truth simpliciter for *utterances*—particular acts of uttering or using sentences.[3] Since an utterance always occurs at a context, we can define utterance truth in terms of sentence truth at a context as follows:

(69) An utterance u is true iff there is a sentence S and a context c such that u is an utterance of S at c and S is true at c.

Utterance truth is standardly thought to be "absolute."

3.1.2 Utterances

This suggests that we might characterize truth relativism as the view that *utterance truth* is relative: one and the same utterance (of a declarative sentence) can be true, relative to X, and false, relative to Y. This is in the right ballpark, I think.[4] But there are some reasons to be dissatisfied with it as a characterization of truth relativism.

[3] Sometimes it is thought that although truth for sentence *types* is context-relative, truth for sentence *tokens*—particular sounds or acoustic blasts—is absolute. But even a sentence *token* can have different truth values on different occasions of use. When I leave my office for a quick errand, I put an old yellow post-it note with a token of "I'll be back in a minute" on my door. Sometimes this sentence token expresses a truth, sometimes a falsehood. For relevant discussion, see Percival (1994, pp. 204–5), Perry (2001, pp. 37–9).

[4] In MacFarlane (2003), I characterize the relativist as someone who rejects the Absoluteness of Utterance Truth.

First, it is linguistically odd to talk of utterances—in the sense of utterance *acts*, not the things uttered (sentences)—as being true or false. In general, we characterize actions as correct or incorrect, but not as true or false. It might be suggested that although "true" and "false" do not apply to all kinds of actions, they do apply to certain speech acts. However, it sounds strange to say "That speech act was true" or "What he did in asserting that sentence was true." This suggests that when we say "His assertion was false" or "That was a true utterance," we are using "assertion" and "utterance" to refer to what is asserted, and not to the *act* of asserting it (Strawson (1950, p. 130), Bar-Hillel (1973, p. 304)).

By itself, this is not a compelling reason for rejecting talk of utterance truth in a theoretical context. Donald Davidson, acknowledging the oddity of characterizing utterance acts as true, says: "Verbal felicity apart, there is no reason not to call the utterance of a sentence, under conditions that make the sentence true, a true utterance" (Davidson, 1990, p. 310). Utterance truth plays a theoretical role in his system, and is assumed to be absolute. Since the relativist is making a theoretical claim, it is okay if it consists in disputing a theoretical claim about a theoretical notion. The ordinary use of "true" is, as we will see, another topic.

However, there are further reasons for avoiding talk of utterance truth. A second reason is given by David Kaplan:

> . . . it is important to distinguish an *utterance* from a *sentence-in-context*. The former notion is from the theory of speech acts, the latter from semantics. Utterances take time, and utterances of distinct sentences cannot be simultaneous (i.e., in the same context). But to develop a logic of demonstratives it seems most natural to be able to evaluate several premises and a conclusion all in the same context. Thus the notion of ϕ *being true in* c *and* $\mathfrak{A}$ does not require an utterance of ϕ. In particular, c_A need not be uttering ϕ in c_W at c_T. (Kaplan, 1989, p. 563)

It would be odd if whether a view counted as relativist about truth could only be discerned from within the theory of speech acts. One might expect there to be a *semantic* difference between relativist and nonrelativist views. So, assuming Kaplan is right that the notion of an utterance belongs to the theory of speech acts and not semantics,

there is good reason to look for a characterization of relativism in terms of sentence truth or proposition truth.

A third reason not to talk of utterance truth that a single utterance act might count as the assertion of two distinct propositions. Aaron Zimmerman provides this example:

> Suppose that after Max drinks himself into unconsciousness, his buddies cover him in plaster and let it harden. The next day, Max relates the night's events to his father and concludes his account by saying, "I really got plastered," intending his statement to express a double entendre. What has Max said? One option is that he expresses the conjunctive proposition that he drank to the point of intoxication and was covered in plaster. Another option is that Max's utterance is elliptical; what he really said was that he was plastered in more than one sense of 'plastered'. But a third option is that Max produces a single token utterance that expresses two propositions: (1) that he got intoxicated, and (2) that he was covered in plaster. Suppose that (contrary to fact) Max has misremembered the incident, and though right in thinking he was covered in plaster, he is mistaken in thinking he got intoxicated. A common intuition is that what Max says is neither entirely false nor entirely true in such a circumstance. Rather, one of the things he says is true while the other is false. Of the three options here considered, only the third renders this verdict. (Zimmerman, 2007, pp. 315–16)

Assuming Zimmerman's third option is correct, we seem to have one utterance here with two truth values. We might deal with this by saying that the utterance has a truth value only relative to a "construal," but as Zimmerman notes, this would not be an interesting kind of relativism about truth.

A more promising response would be to talk not of utterances but of *assertions*—again in the "act" sense of assert*ings*, rather than the "object" sense of things assert*ed*. It is plausible that in Zimmerman's story, Max makes two distinct assertions by making one utterance (just as one might kill two birds by casting one stone). So we might try characterizing relativism about truth as the view that *assertions* have truth values only relative to some parameter.

This characterization may be satisfactory for some purposes, but it is not sufficiently fundamental. Indeed, it immediately raises questions about the truth of propositions—the contents of assertions—that demand answers. Suppose Jim asserts that p by uttering sentence S at context c, and suppose that his assertion is true relative to X but false relative to Y. Is this relativity attributable to a relativity of the truth of p? If so, what is this relativity, and why can't relativism about truth be characterized directly in terms of it? If not, how can an assertion of a proposition whose truth is not relative to X and Y have truth values only relative to X and Y?[5]

Finally, it seems odd to define relativism about truth in terms of assertion. Surely, if truth is relative in some kind of discourse, that has ramifications not just for assertion but for other speech acts (commands, questions, bets, conjectures, and so on), and also for mental states. Couldn't there be non-language-using creatures with *beliefs* whose truth is relative? If truth relativism has something special to do with assertion, we should at least expect some kind of explanation. The proffered account does not provide one.[6]

3.1.3 Propositions

All of this suggests that the relativist doctrine should be stated not as a claim about the truth of assert*ings*, but as a claim about the things that are assert*ed*, which, following tradition, I will call *propositions*.[7] Propositions are usually thought of as the "primary bearers of truth value." What this means is that other things that have truth values (sentences, beliefs, assertions, etc.) have them by virtue of standing in an appropriate relation to propositions that have those truth values. It is natural to think, then, that if all of these other things have their truth values only relatively, it is because propositions do.

Accordingly, Max Kölbel has characterized "non-tame" relativism about truth as the view that "the truth of propositions (or contents) of some kind can be relative" (Kölbel, 2002, p. 119). By this criterion, though, just about everyone who uses propositions in formal semantics

[5] *Exercise to the reader:* Return to these questions after reading chapter 4.

[6] *Exercise to the reader:* Return to this question after reading chapter 5.

[7] For some arguments for the theoretical utility of propositions, and the point that propositions are not to be identified with sentence meanings, see Cartwright (1962).

would count as a non-tame relativist. For it is orthodox practice to relativize truth of propositions to possible worlds—and in some frameworks worlds and times.[8] For example, the proposition that dodos are extinct in 2004 is true in the actual world, but there are possible worlds relative to which the very same proposition is false. Surely this much relativism does not constitute "relativism about truth" in the sense we are trying to capture. It can be motivated by considerations that have nothing to do with deficient objectivity: for example, by the need to understand claims like

(70) That's true, but it wouldn't have been true if you'd called her.

And it is perfectly compatible with the idea that particular assertions can be assigned absolute truth values. We only need to add that an assertion that p is true (or false) simpliciter just in case p is true (or false) at the actual world (or, in some versions, the world in which the assertion is made).

Faced with this fact, most writers who seek to characterize truth relativism at the propositional level resort to discrimination. Relativizing propositional truth to possible worlds, they say, is just a formal way of registering the fact that the truth of a proposition depends on how things are. It is relativizing propositional truth to factors *beyond* just worlds that makes one a relativist about truth (Nozick (2001, p. 19), Stanley (2005, p. 137), Zimmerman (2007, p. 316), Kölbel (2008*a*, p. 4)).

This characterization is problematic in several ways. First, it casts *temporalism*—the view that propositions have truth values relative to times, in addition to possible worlds—as a kind of truth relativism. And this seems to put the line between relativist and nonrelativist views in the wrong place. Every reason against letting relativization of propositional truth to possible worlds count as "truth relativism" applies also to relativization of propositional truth to times:

1. Both relativizations have a claim to be orthodox. This is not always acknowledged by opponents of temporalism. But the view that propositional truth is time-relative was widespread in ancient, medieval, and modern philosophy, and only began to wane in the twentieth century (for an illuminating account of the

[8] See, for example, Stalnaker (1987), Kripke (1972), Lewis (1986), and Kaplan (1989).

history, see Prior (1957, Appendix A)). More recently, Prior (1957, 2003) and Kaplan (1989, pp. 502–9) have taken propositions to have truth values relative to (worlds and) times.[9]

2. Both relativizations are motivated in ways that do not connect with the traditional motivations for relativism—for example, by attention to our ordinary talk about when two people (or one person at two different times) believe "the same thing." Just as (70) gives us some reason to relativize propositional truth to worlds, so

   (71) That was true during the Clinton years, but it isn't true now.

   gives us some reason to relativize propositional truth to times. The debates between temporalists and eternalists turn on the significance of these ordinary language considerations, on general philosophical issues, and on technical issues concerning tense and semantic value.[10] But they do not seem to turn on any of the issues that are at stake in traditional debates about relative truth.

3. Most importantly, both relativizations are consistent with the absoluteness of utterance (or assertion) truth. Just as the eternalist will say that an assertion that p is true simpliciter if p is true at the world in which the assertion occurs, so the temporalist will say that an assertion that p is true simpliciter if p is true at the world and time at which the assertion occurs. The temporalist and the eternalist will agree on all questions about the truth of various dateable utterances, and they will take all such questions to have "absolute" answers. If the eternalists' commitment to the absoluteness of utterance truth is what keeps them from being counted as truth relativists, temporalists should not be counted as truth relativists either.

[9] Kaplan talks of "contents" and notes that his usage departs from "the traditional notion of a proposition" (503 n. 28; cf. 546). However, he is explicit that he intends his contents to be "what is said" by utterances of declarative sentences, and notes that "[t]he content of a sentence in a given context is what has traditionally been called a proposition" (500).

[10] For relevant discussion, see King (2003); Richard (1980, 1982, 2003); Salmon (1986, 2003), and especially Kaplan (1989, 503 n. 28).

It is not important here whether one accepts or rejects temporalism. The question is whether the temporalist position—even if it is wrong or misguided—should count as a form of relativism about truth. What I am arguing is that there is no good reason to count it as relativist that would not apply equally to eternalism.

One might try saying that one is a relativist about truth if one relativizes propositional truth to something besides worlds and times. But now the characterization begins to look unprincipled. Are worlds and times the only innocuous parameters? What about the proposal—also considered by Kaplan—to countenance "locationally neutral" propositions, like the proposition *that it is raining*, that have truth values relative to worlds, times, and *locations* (Kaplan, 1989, p. 504)? This proposal does not seem different in kind from the proposal to relativize propositional truth to times. Again, there may be reasons for countenancing temporally neutral propositions but not locationally neutral ones, but the question here is just whether a commitment to locationally neutral propositions makes one a relativist about truth, and if so why. What is the principle by which parameters are sorted into those relativization to which makes one a relativist about truth, and all the others?[11]

Putting aside worries about the unprincipled nature of the division between "innocuous" and "suspicious" parameters of propositional truth, there is a more serious problem with the proposed characterization of relativism. The problem is that one can describe views that make utterance (or assertion) truth relative without countenancing any nonstandard parameters of propositional truth at all. We can describe a simple example using the temporalist's tensed (time-relative) propositions. Reasonable temporalists, like Kaplan, will say that an assertion of a tensed proposition is true just in case the proposition is true at the world and time of utterance. So, if at 2 PM I assert that Socrates is sitting, then I have asserted truly just in case the proposition I have asserted—*that Socrates is sitting*—is true at 2 PM. But instead of taking this reasonable view, one could instead say that such an assertion has no truth value simpliciter, but only time-relative truth values: as

[11]Nozick (2001, 307 n. 7) admits (and regrets) that he has no principled basis for demarcating "the harmless factors, relativity to which does not constitute relativism, from the factors that make for relativism."

assessed from time t_1, an assertion of p at t_0 is true just in case p is true at (the world of utterance and) t_1.[12] Such a view would be silly, of course, but the question is whether it should count as a form of relativism about truth. Clearly it should, since it denies that dated assertions have absolute truth values. Despite that, it does not relativize propositional truth to any "nonstandard" parameters. So relativization to nonstandard parameters is not necessary for relativism about truth. (This point will be made more forcefully in chapter **??**, where we will examine a relativist view that is not at all silly, and that relativizes propositional truth to nothing besides possible worlds.)

I will argue in chapter 4 that relativization of propositional truth to nonstandard parameters is not *sufficient* for relativism about truth, either. As we will see, it is not the *kind* of parameters to which one relativizes propositional truth that makes one a relativist, but rather what one does with them.

3.2 Assessment Sensitivity

I am going to suggest that what makes one a relativist about truth is a commitment to the *assessment sensitivity* of some sentences or propositions. The primary task of this section, then, is to explain what assessment sensitivity is. For simplicity, I will work in a semantic framework, due to Lewis (1980), that works only with sentences (and open formulas), not propositions. Then, in chapter 4, we will see what assessment sensitivity looks like in a semantic framework, like that of Kaplan (1989), that makes use of the notion of a proposition.

3.2.1 Truth at a context of use

The goal of a semantic theory for a language L, as Lewis (1980) conceives it, is to define truth at a context of use for arbitrary sentences of L. That is, given any sentence S of L, the semantic theory must tell us what a context must be like in order for an utterance of S at that context to express a truth. For example, a semantic theory for English will tell us that

(72) I am six feet tall

[12] This is essentially the view (6) criticized by Evans (1985, p. 347), substituting assertion truth for "correctness."

is true at a context if the agent of that context is six feet in height at the time and world of the context. It will tell us that

(73) Snow is white and grass is green

is true at a context just in case snow is white and grass is green at the time and world of the context. And so on. In short, it will give us "truth conditions" for all the sentences in the language.

We will think of a context as a possible occasion of use of a sentence (Kaplan, 1989, p. 494).[13] Formally, we might model a context as a sequence of parameters (agent, world, time, location, and so on) or as a "centered possible world" (a world with a designated time and location as "center"). We will assume that however a context is represented, it determines a unique agent, time, world, and location. (By the "agent" of a context, I mean the user or potential user of the sentence. One might talk more loosely of the "speaker," but we might want to consider contexts at which the agent is not speaking.) We will assume also that the agent of a context exists at the world and time of the context, and is at the location of the context at the time of the context. This assumption is required by the idea that a context is a possible occasion of use.[14]

Why is truth-at-a-context the target notion of a semantic theory?[15] Because truth-at-a-context has direct pragmatic relevance. When we speak, in the normal case, we try to use sentences that are true at our contexts, and we expect others to be doing the same:

> The foremost thing we do with words is to impart information, and this is how we do it. Suppose (1) that you do not know whether *A* or *B* or ...; and (2) that I do know; and (3) that I

[13] This is an objective concept of context. Contrast Stalnaker (1978), who thinks of a context as the set of propositions that are taken for granted as common ground in a conversation.

[14] See Lewis (1980, pp. 28–9), Kaplan (1989, 512 n. 37). For an argument that this restriction on contexts should be dropped, see Predelli (2005, ch. 2).

[15] Lewis talks at first of "truth-in-English" as the target notion, but as he notes in §3, for non-mathematical languages truth will depend on "features of the situation in which the words are said," so the target notion becomes "truth-in-English at a context." "To do their first job of determining whether truth-in-English would be achieved if a given sentence were uttered in a given context, the semantic values of sentences must provide information about the dependence of truth on context" (Lewis, 1998, p. 31).

> want you to know; and (4) that no extraneous reasons much constrain my choice of words; and (5) that we both know that the conditions (1)–(5) obtain. Then I will be truthful and you will be trusting and thereby you will come to share my knowledge. I will find something to say that depends for its truth on whether *A* or *B* or ... and that I take to be true. I will say it and you will hear it. You, trusting me to be willing and able to tell the truth, will then be in a position to infer whether *A* or *B* or ... (Lewis (1998, p. 22), cf. Lewis (1983, §III))

So the condition for a sentence to be true at a context is the central semantic fact we need to know if we are to *use* the sentence and understand others' uses of it.[16] Truth at a context is the point at which semantics makes contact with pragmatics, in the broad sense—the study of the use of language.

3.2.2 Truth at an index and context

It is a simple enough matter to state the condition for a particular sentence to be true at a context. But a semantic theory for a language needs to encode truth conditions for *all* sentences of the language. Since natural languages (and most useful artificial ones) allow the formation of arbitrarily complex sentences, there will be infinitely many of them. Obviously, we can't just *list* them together with their truth conditions. We need some way of *computing* the truth conditions of a sentence from a structural description of it.

For some simple languages, we can do this by direct recursion. Suppose our language contains just two atomic sentences, "I am happy" and "Grass is green," together with a unary connective "It is not the case that" and a binary connective "and." Then we can specify truth conditions for all of its sentences with the following clauses:[17]

[16] This might be disputed, on the grounds that one could know that "2 + 2 = 4" is true at *every* context without having any idea how to use it. However, to know that "2 + 2 = 4" is true at every context is not to know the *condition* for it to be true at a context. To know this, one must know that the sentence is true at a context iff the sum of 2 and 2 is 4 at the world of that context. It may be that this condition is satisfied by every context, but the extra bit of mathematical knowledge it takes to see this is not part of the condition itself.

[17] Usually one would give a rule for generating truth conditions for atomic sentences from specifications of the extensions of terms and predicates, but for simplicity, and because we have only two atomic sentences, I have avoided this here.

(74) "I am happy" is true at c iff the agent of c is happy at the world of c.

(75) "Grass is green" is true at c iff grass is green at the world of c.

(76) ⌜It is not the case that ϕ⌝ is true at c iff ϕ is not true at c.

(77) ⌜ϕ and ψ⌝ is true at c iff ϕ is true at c and ψ is true at c.

The same technique will work for any language with a finite number of atomic sentences and truth-functional connectives. But it will not work for languages with quantifiers or non-truth-functional operators. We have already seen why it will not work for quantifiers: quantified sentences are constructed out of open formulas, and these do not have truth values at contexts (§3.1.1, above). To see why it won't work for non-truth-functional operators, suppose we add to our language a unary operator "It has always been the case that." The natural clause would be:

(78) ⌜It has always been the case that ϕ⌝ is true at c iff it has always been the case that ϕ is true at c.

But this doesn't give us what we want: "true at c" is a timeless predicate; if ϕ is true at c, then it has always been and will always be true at c. So putting "it has always been the case that" in front of "ϕ is true at c" in our clause has no effect. We might try:

(79) ⌜It has always been the case that ϕ⌝ is true at c iff for every context c' that differs from c at most in taking place at an earlier time, ϕ is true at c'.

But this won't work, either, for two reasons.

The first problem is that if ϕ contains a time-sensitive indexical like "now" or "yesterday," its denotation will shift as we evaluate ϕ relative to the time-shifted contexts. And this will get the truth conditions wrong: "now" and "yesterday" should not shift their denotations when embedded under "it has always been the case that…" (Kamp, 1971).[18]

[18] Kaplan (1989, pp. 510–12) argues, further, that natural languages do not contain any operators that shift contexts, as "It has always been the case that" does on the semantics of (79). The claim has been widely accepted, though Schlenker (2003) has questioned it.

The second problem is that if we look only at contexts that differ from c *only* in the time of the context—agreeing with c on the agent and world of the context—we won't be looking at *all* the times prior to the time of c. Since the agent (speaker) of a context must *exist* at the time and world of the context, we will not be looking at any times prior to the birth of the speaker of c in the world of c. There just *aren't* any possible contexts c' such that the agent of c' = the agent of c, the world of c' = the world of c, and the time of c' is earlier than the birth of the agent of c in the world of c. Surely, though, "It has always been the case that" must quantify over *all* times prior to the time of the context. Lewis summarizes the problem succinctly:

> Unless our grammar explains away all seeming cases of shiftiness, we need to know what happens to the truth values of constituent sentences when one feature of context is shifted and the rest are held fixed. But features of context do not vary independently. No two contexts differ by only one feature. Shift one feature only, and the result of the shift is not a context at all. (Lewis (1998, p. 29); cf. Kaplan (1989, p. 509))[19]

The solution, Lewis suggests, is to relativize truth not just to contexts but to *indices*: "packages of features of context so combined that they *can* vary independently":

> An index is an n-tuple of features of context of various sorts; call these features the *coordinates* of the index. We impose no requirement that the coordinates of an index should all be features of any one context. For instance, an index might have among its coordinates a speaker, a time before his birth, and a world where he never lived at all. Any n-tuple of things of the right kinds is an index. So, although we can never go from one context to another by shifting only one feature, we can always go from one index to another by shifting only one coordinate. (Lewis, 1998, pp. 29–30)

[19] Lewis might seem to be overstating things. Couldn't two contexts differ only in the time of the context? Certainly the agent and world of the context could be the same. But, as Lewis notes, there are "countless other features" of contexts that might, in theory, be semantically significant: for example, the temperature of the context, the conversationally salient objects of the context, and so on. In any case, the argument against (79) does not depend on the strong claim that *no* two contexts differ in just the time of the context.

Since our operator "It has always been the case that" shifts the time, handling it will require indices containing at least a time coordinate. If we also have operators like "Possibly," which shift the world, we will need indices that contain both time and world coordinates. Instead of defining "true at c" directly, then, we will give a recursive definition of "true at $c, \langle w, t \rangle$." The clause for "It has always been the case that" will look like this:

(80) ⌜It has always been the case that ϕ⌝ is true at $c, \langle w, t \rangle$ iff for every time $t' \leq t$, ϕ is true at $c, \langle w, t' \rangle$.

Using clauses of this kind, we can define truth at a context and an index for arbitrary sentences. However, our project was to define truth at a context. Truth at a context has direct pragmatic relevance; truth at a context and index is a technical notion. How can we turn our definition of the latter into a definition of the former? In Lewis's framework the trick is easy. We have assumed that coordinates of indices will be "features of context." So, by "initializing" the index with the values provided by context, we can define truth at a context in terms of truth at a context and index:

> Let us say that sentence s is true at context c iff s is true at c at the index of the context c. (Lewis, 1998, p. 31)

By the *index of the context*, Lewis means the index whose coordinates are given the values determined by the context. In our example, the index of a context c would be $\langle w_c, t_c \rangle$, where w_c is the world of c and t_c is the time of c. So, although we cannot define truth at a context directly, we can define it indirectly, by recursively defining truth at a context and index, and then defining truth at a context in terms of this more technical notion. (In MacFarlane (2003, §V), I call the definition of truth at a context and index the *semantics proper* and the definition of truth at a context in terms of this the *postsemantics*, and I will sometimes use this terminology in what follows.)

Although Lewis requires that each coordinate of the index be a "feature of context," this requirement can be relaxed. What is essential is that we have some way of moving from truth-at-a-context-and-index to truth-at-a-context. Lewis's requirement gives us a particularly straightforward way of making this move, since it guarantees that the context of use will supply an initial value for every coordinate of the index:

that is, for any coordinate X, we can always talk of "the X of the context." But the requirement does not make sense for some shiftable coordinates. Consider assignments of values to the variables, for example, which are shifted by quantifiers. Assignments are not features of contexts; contexts determine places, times, worlds, and many other things, but not assignments of values to variables (cf. Kaplan (1989, pp. 592–3), Belnap, Perloff, and Xu (2001, pp. 150–1)). Or consider the way *histories* are treated in "branching time" semantics in the style of Thomason (1970) and Belnap, Perloff, and Xu (2001). Here indices are moment/history pairs, where moments are concrete time slices of the universe, and histories are maximal chains of moments ordered by a causal ordering relation—complete possible histories of the universe. Temporal operators ("It will be the case that", "It has always been the case that", and so on) shift the moment coordinate, while historical modals ("It is settled that", "It is still possible that") shift the history coordinate. But while contexts of use determine a moment, they do not determine a history (Belnap and Green (1994), Belnap, Perloff, and Xu (2001, pp. 151–2)). Lewis's requirement would therefore rule out both assignments and histories as coordinates of indices.[20]

Given that context does not initialize the assignment parameter, how do we eliminate the relativization to assignment in a definition of truth at a context? By quantifying over all assignments (see (68), above). Similarly, we might eliminate the relativization to histories by quantifying over all histories that pass through the moment of the context.[21] Thus, if our indices consist of a moment, a history, and an

[20] Although Lewis (1980) says nothing about assignments or quantifiers, Lewis (1970*b*) does talk of an assignment coordinate of indices. Perhaps Lewis (1980) would do semantics for quantified languages by relativizing truth to a context, an index, *and* an assignment. But there is no good reason, other than the requirement being discussed here, not to count the assignment as a coordinate of indices. The motivation for assignments is exactly the same as the motivation for other coordinates of indices: the proper treatment of shiftiness. Lewis also does not discuss histories, but his discussion of "overlapping worlds" in Lewis (1986, pp. 199–209) suggests that he would reject the idea that contexts do not determine a unique history (for criticism, see MacFarlane (2008)).

[21] This yields the supervaluational semantics of Thomason (1970); as we will see in chapter **??**, it is not quite right, but the problem does not lie with the idea of quantifying over histories in definiting truth at a context.

assignment, we might define truth at a context in terms of truth at a context and index as follows:

(81) A sentence S is true at context c iff for all indices $\langle m_c, h, a\rangle$, where m_c is the moment of c, $m_c \in h$, and a is any assignment, S is true at $c, \langle m_c, h, a\rangle$.

(Note that this definition allows that some sentences may be neither true nor false at a context—where a sentence is false iff its negation is true. This will happen whenever there are histories h_1 and h_2 containing m_c such that S is true at $\langle m_c, h_1, a\rangle$ and false at $\langle m_c, h_2, a\rangle$ for some assignment a. Thus bivalence holds for truth at a context and index, but not for truth at a context.)

It makes sense, then, to relax Lewis's requirement that indices be *features* of context, as long as we can still define truth at a context in terms of truth at a context and index.

Neither of the relativizations of truth we have considered so far involves us in any philosophically controversial kind of "relative truth." The relativization to contexts is required because the same sentence can be used to make true or false claims, depending on the context. The relativization to indices is required as a technical expedient for systematizing truth at a context. Since indices have no theoretical role beyond their role in defining truth at a context, the only motivation for positing a coordinate of indices is the presence of an operator that shifts it; conversely, the only grounds for objecting to a coordinate of indices is the absence of such operators. General considerations about truth and reality simply aren't relevant here.

Indeed, as we have already observed for the special case of assignments, we could dispense with talk of truth at a context and index in favor of a function from sentence, context, index triples onto $\{0, 1\}$, and then define truth at a context directly in terms of this function:

(81′) A sentence S is true at context c iff for all indices $\langle m_c, h, a\rangle$, where m_c is the moment of c, $m_c \in h$, and a is any assignment, $v(S, c, \langle m_c, h, a\rangle) = 1$.

So from a philosophical point of view, no eyebrows should be raised even at "wild" coordinates of indices like standards of precision or aesthetic standards. These are merely technical devices for systematizing truth at a context, to be justified (or not) on *linguistic* grounds.

For example, if "strictly speaking" is best understood as a sentential operator that shifts standards of precision, we will need a standards of precision coordinate; if not, not. In any case the debate is not a distinctively philosophical one.

3.2.3 Contexts of assessment

I now want to suggest that the philosophically interesting line between truth absolutism and truth relativism is crossed when we relativize truth not just to a context of use and an index, but also to a *context of assessment*.

We are already comfortable with the notion of a "context of use," understood as a possible situation in which a sentence might be used. So we ought to be able to make good sense of the notion of a "context of assessment"—a possible situation in which a use of a sentence is being *assessed*. There shouldn't be anything controversial about contexts of assessment: if there can be assessments of uses of sentences, then surely we can talk of the contexts in which these assessments would occur.

To move from Lewis's framework to a framework in which relativist proposals can be described, we need only give contexts of assessment a role in our semantics parallel to that of contexts of use. Our target notion, then—the one with direct pragmatic relevance—will be not "true as used at c," but "true as used at c_1 and assessed from c_2."[22]

Ontologically speaking, contexts of use and contexts of assessment can be thought of as the same kind of thing. They might both, for example, be modeled as "centered possible worlds" (possible worlds with a designated time and agent or location). The qualifiers "of use" and "of assessment" distinguish two different uses to which a context can be put in semantics. We can think of a context as a possible situation of *use* of a sentence, or as a possible situation of *assessment* of a use of a sentence. In the former case, the agent of the context is the user of the sentence—the speaker, when the use is a speaking—while in the latter, the agent of the context is the assessor of a use of the sentence. Since the same contexts can serve either as contexts

[22]We will return to the question *how* this doubly-relativized truth predicate is pragmatically relevant in chapter 5, below.

of use or as contexts of assessment, it makes sense to ask whether a sentence is true as used at and assessed from the very same context.

A particular dated use of a sentence may be assessed from indefinitely many possible contexts. Thus, although we may talk of "*the* context of use" for such a use, we may not talk in the same way of "*the* context of assessment." The definite article will be appropriate only when we have in mind not just a particular use, but a particular assessment. It is important that the context of assessment is not fixed in any way by facts about the context of use, including the speaker's intentions; there is no "correct" context from which to assess a particular speech act.[23]

How might a context of assessment enter into the semantics for an expression? In just the same ways as a context of use. In general, there are two ways for a feature of a context of use to be semantically relevant: it can be *locally* relevant, by playing a role in the semantics proper—that is, in the recursive clause for a particular linguistic construction—or it can be *globally* relevant, by playing a role in the postsemantics—the definition of truth at a context in terms of truth at a context and index. For example, in Kaplan's semantics for indexicals (Kaplan, 1989), the world of the context of use (w_c) is locally relevant because of the role it plays in the clause for the operator 'A' ("it is actually the case that"):

(82) $\ulcorner A\phi \urcorner$ is true at $c, \langle w, t, a\rangle$ iff ϕ is true at $c, \langle w_c, t, a\rangle$ (545, with notational changes)

and globally relevant through its role in the definition of truth at a context:

[23]This distinguishes the proposal being made here from other proposals in the literature that bifurcate context. Predelli (1998) argues that in making a recorded utterance, e.g. "I am not here now" on an answering machine, the speaker may *have in mind* a "context of interpretation" relative to which some of the context-sensitive expressions ("here," "now"), but not others ("I") are to be evaluated. Schlenker (2004) proposes distinguishing "context of utterance" (controlling the interpretation of tense and person) and "context of thought" (controlling the interpretation of other indexicals) in order to make sense of free indirect discourse ("Tomorrow was Monday, Monday, the beginning of another school week!") and the historical present ("Fifty eight years ago to this day, on January 22, 1944, just as the Americans are about to invade Europe, the Germans attack Vercors"). My notion of context of assessment, by contrast, has nothing to do with the speaker's (or author's) intentions, and is not fixed in any way (even "intentionally") by the context of use.

(83) A sentence S is true at a context c iff for every assignment a, S is true at $c, \langle w_c, t_c, a\rangle$. (547, simplified with notational changes)

Some features of context are only locally relevant (for example, the agent of the context, which figures in the semantic clause for "I"). And in principle, a feature might be only globally relevant (as the world of the context would be in a language not containing an actuality operator).[24]

Features of contexts of assessment can also be semantically relevant either locally or globally. For a toy example of the former case, imagine adding to English a word "noy" that works like "now," except that where "now" denotes the time of the context of use, "noy" denotes the time of the context of assessment.[25] To do semantics for a language containing "noy," we would need to recursively define truth, and more generally extension, relative to a *pair* of contexts (of use and assessment) and an index. Compare the recursive clauses for "now" and "noy":

(84) The extension of "now" at $c_1, c_2, \langle w, t, a\rangle$ is the time of c_1.

(85) The extension of "noy" at $c_1, c_2, \langle w, t, a\rangle$ is the time of c_2.

Admittedly, "noy" is a pretty silly word—one for which it is hard to find a use. But one can see how it differs from "now." A particular use of "Jim is hungry now" will have the same truth value as assessed from every context, while a use of "Jim is hungry noy" will have different truth values relative to different contexts of assessment. On the other

[24] As Kaplan (1989, p. 595) observes: "... it may appear that for a modal language *without* indexicals, without expressions that require a parameter, the notion of a context of use has no bearing. This is not correct. Truth in every model means truth in the 'designated' world of every model. This 'designated' world, the world at which truth is assessed, plays the role of actual-world."

[25] Kaplan (1989, 491 n. 12) reports Donnellan as having suggested something superficially similar: "if there were typically a significant lag between our production of speech and its audition (for example, if sound traveled very very slowly), our language might contain two forms of 'now': one for the time of production, another for the time of audition." Donnellan's second form of "now" is not the same as "noy," because audition is not the same as assessment; if one reassesses an assertion some time after first hearing it, the time of assessment is different, but the time of audition is the same. Still, Donnellan's proposal would require relativity of truth to contexts of assessment—not to the *time* of assessment, but to the *agent* of assessment, since the time of audition may vary from one assessor to another.

hand, as assessed from any given context, all uses of "noy" will denote the same time, whereas uses of "now" will typically denote different times.

For an example of a globally relevant feature of contexts of assessment, suppose we replace (81) with

(86) A sentence S is true as used at context c_1 and assessed from context c_2 iff for all indices $\langle m_{c_1}, h, a\rangle$, where $m_{c_2} \in h$ and a is any assignment, S is true at $c_1, \langle m_{c_1}, h, a\rangle$.[26]

This definition draws on the context of use to tell us which moment to look at, and on the context of assessment to tell us which *histories* to look at (the ones that run through the moment of the context of assessment).[27] The language may not contain any expression whose semantic clause makes reference to the moment of the context of assessment, but this feature of the context of assessment is still semantically relevant through its role in the definition of truth-at-contexts in terms of truth-at-contexts-and-an-index.

In a framework in which both context of use and context of assessment may be semantically relevant, context sensitivity comes in two flavors. We will say that

(87) An expression is *use-sensitive* if its extension (relative to a context of use and context of assessment) depends on features of the context of use, and

(88) An expression is *assessment-sensitive* if its extension (relative to a context of use and context of assessment) depends on features of the context of assessment.

Note that every contingent sentence counts as use-sensitive on this definition, since its truth value depends on the world of the context of

[26] Here we are defining truth at a context of use and context of assessment in terms of truth at a context of use and an index. One might ask why we do not define it in terms of truth at a context of use *and context of assessment* and an index. The answer is that this is not necessary unless the language contains expressions, like "noy," that are *locally* sensitive to features of contexts of assessment. The present definition highlights the fact that the moment of the context of assessment is only *globally* relevant; there is no way it *could* be locally relevant, because the recursive clauses for individual expressions see only the context of use and index.

[27] A cleaned-up version of (86) will be discussed in chapter **??**, below.

use.[28] It is useful, then, to parameterize the notions of use sensitivity and assessment sensitivity to indicate the *feature* of context on which an expression's extension depends. We will say that

(89) An expression is *F-use-sensitive* if its extension (relative to a context of use and context of assessment) depends on the *F* of the context of use, and

(90) An expression is *F-assessment-sensitive* if its extension (relative to a context of use and context of assessment) depends on the *F* of the context of assessment.[29]

So, for example, "The US stock market plunged on October 22, 2008" is world-use-sensitive but not time-use-sensitive, and "Jim is sitting noy" is time-assessment-sensitive.

3.3 Truth relativism as assessment sensitivity

Using these concepts, we can say what it is to be a relativist about truth, in the serious and philosophically nontrivial sense we have been seeking. To be a *relativist about truth* is to hold that languages with assessment-sensitive expressions are at least conceptually possible. This is a position one might endorse or reject on nonempirical, philosophical grounds; what it requires is that one come to understand *what it would be* for an expression's extension to depend on features of the context of assessment. To be a *relativist about truth in English* (or some other natural language) is to hold that some expressions of English are assessment-sensitive. Relativism about truth in English is at least partly an empirical thesis. It is coherent to hold that, although we can understand what it would be for an expression to be assessment-sensitive, assessment-sensitive expressions are not found in natural languages.

[28] This is what David Lewis was getting at when he said that "[c]ontingency is a kind of indexicality," Lewis (1998, p. 25); for further discussion, see MacFarlane (2009, §3).

[29] Note that "depends" in these definitions has causal/explanatory force. To show that the truth value of *S* depends on feature *F*, it is not enough just to find two contexts that differ with respect to *F* and relative to which *S* has different truth values. For the difference in truth values may be due to other differences between these contexts. As we have noted, it is generally not possible to find pairs of contexts that differ in respect *F* without differing in many other ways as well.

This characterization of truth relativism connects very naturally with the idea (explored in §3.1.2, above) that a truth relativist is a relativist about the truth of *utterances* or assertions (conceived as acts). An utterance of a sentence fixes a unique relevant context of use—the context in which the utterance occurs—but not a unique relevant context of assessment. So utterances of assessment-sensitive sentences can be assigned truth values (if at all) only relative to a context of assessment. However, our characterization of truth relativism in terms of assessment sensitivity avoids all the liabilities of talk of utterance truth. It avoids a linguistically odd application of a truth predicate to *acts*. And it is a purely semantic characterization, not one that requires notions from pragmatics or the theory of speech acts. So it is not subject to difficulties like the one raised by Zimmerman's "plastered" example.

Most importantly, it distinguishes clearly between three ways in which the truth of a sentence might be relative to some feature F:

1. The sentence's truth might vary with the F coordinate of the index.
2. The sentence might be F-use-sensitive.
3. The sentence might be F-assessment-sensitive.

Some examples may help to make these distinctions concrete.

(91) I swam in Lake Lucerne in 1984

is agent-use-sensitive, but not agent-assessment-sensitive, and there is no agent coordinate of the index (because there is no operator that would shift it).

(92) Socrates is sitting

is time-use-sensitive, but not time-assessment-sensitive; whether there is a time coordinate of the index will depend on whether tenses are analyzed as operators or referentially.[30]

(93) Socrates is sitting now

is time-use-sensitive, but not time-assessment-sensitive, and its truth does not vary with the time coordinate of the index (even if there is one).

[30] See King (2003) for a discussion.

(94) Socrates is sitting noy

is time-assessment-sensitive but not time-use-sensitive, and its truth does not vary with the time coordinate of the index. In a branching time framework, the truth of

(95) There will be a sea battle tomorrow

at a context and index depends on the history coordinate of the index, but the sentence is neither history-use-sensitive nor history-assessment-sensitive. (Recall that histories are not features of contexts.)

As these examples reveal, it is not the *kind* of thing to which truth is relativized that makes a position "seriously relativist," but the *way* in which truth is relativized to it.

To see this even more clearly, imagine a language with aesthetic terms ("beautiful," "ugly") and a sentential operator "by any aesthetic standard."[31] An operator needs a coordinate of the index to shift, so we will need an "aesthetic standard" coordinate in our indices. Whether a sentence like "That painting is beautiful" is true at a context and index will depend in part on the aesthetic standard coordinate of the index. The operator "by any aesthetic standard" can be treated as follows:

(96) ⌜By any aesthetic standard:ϕ⌝ is true at $c, \langle w, s, a\rangle$ iff for all aesthetic standards s', ϕ is true at $c, \langle w, s', a\rangle$.

In doing this much, have we committed ourselves to any serious kind of relativism about truth? Plainly not. For we might define truth at a context in the following way:

(97) A sentence S is true at a context c iff for all assignments a, S is true at $\langle w_c, s_G, a\rangle$, where w_c is the world of c and s_G is God's aesthetic standard.

On this semantics, the truth of aesthetic sentences would be completely insensitive to the aesthetic standards of the speaker or the assessor. The relativization to aesthetic standards in the index would have

[31] Put aside worries about whether "by any aesthetic standard" is best treated as a sentential operator in English; just stipulate that the language we are dealing with has a sentential operator with this meaning.

a merely technical role, for systematizing the truth conditions of sentences containing the operator "by any aesthetic standard."

Alternatively, we might define truth at a context as follows:

(98) A sentence S is true at a context c iff for all assignments a, S is true at c, $\langle w_c, s_c, a\rangle$, where w_c is the world of c and s_c is the aesthetic standard of the agent of c.

On this semantics, the truth of aesthetic sentences would depend on the speaker's aesthetic standards, but would be completely independent of the assessor's standards. Utterances of aesthetic sentences could be assigned absolute truth values.

The threshold of relative truth is only crossed when we give a semantically significant role to the context of assessment:

(99) A sentence S is true as used at a context c_1 and assessed from a context c_2 iff for all assignments a, S is true at c_1, $\langle w_{c_1}, s_{c_2}, a\rangle$, where w_{c_1} is the world of c_1 and s_{c_2} is the aesthetic standard of the agent of c_2.

This semantics does not allow us to assign absolute truth values to utterances of "That painting is beautiful." It holds that aesthetic sentences can be assigned truth values only relative to the aesthetic standard of the assessor. It is only at this point—at which assessment sensitivity is countenanced—that we run into *philosophical* issues concerning truth. The relativization of truth to aesthetic standards in (97) is just a technical device, like relativization of truth to an assignment. It is justified, if at all, by the technical requirements of the project of defining truth at a context, and if it is justified in that way, it requires no further defense. The relativity of truth to the aesthetic standards of the speaker in (98) is of the same character as the relativity of truth of sentences containing "here" to the location of the speaker. It is justified, if at all, by the sorts of considerations that normally support positing context sensitivity, and if it is justified in that way, it raises no philosophical problems not already raised by "here" and the like. But with the relativization of truth to the aesthetic standards of the assessor in (99), we find something genuinely new—something that needs philosophical clarification and justification. We will turn to that task in chapter 5, but first, it is worth asking how propositions are to be conceived in a framework that allows for assessment sensitivity.

Draft: December 13, 2009

4 Propositions

In §3.2, we followed Lewis in taking the semanticist's project to be that of articulating truth conditions for *sentences*. We have said nothing about propositions—the contents of assertions, beliefs, and other attitudes—and one might wonder how they fit into the picture. So in this chapter we will see how to fit assessment sensitivity into a Kaplan-style *two-stage* semantic framework. In a two-stage semantics, we do not directly assign truth-conditions to sentences; instead, we describe how a sentence's propositional *content* depends on its construction and features of context, and then give an account of the truth conditions of these contents.

Semanticists disagree about the usefulness—and even the feasibility—of a two-stage framework.[1] But what we have to say about propositions and assessment sensitivity should be of interest even if one thinks that one need not talk of propositions in one's *semantics*. Most theories of mental attitudes and speech acts take their contents to be proposition-like rather than sentence-like, at least in being language-independent. I can assert the same thing in English as Pierre asserts in French, and in doing so I can express the very same belief—say, that snow is white. So, if we are to understand the practical significance of (say) the proposal that "tasty" is assessment-sensitive, we had better understand what that proposal implies about propositions, like the proposition that Hen-of-the-Woods is tasty.

4.1 Content and character

Kaplan takes the linguistic meaning of an expression or phrase to be its *character*—a rule for determining its content in any given context. The

[1] For discussion, see Lewis (1980, §§9–12) and King (2003).

content of a sentence at a context is a proposition—"what is said" by the sentence at that context. The content of a subsentential expression is its contribution to the propositions expressed by sentences containing it: for a proper name, this will be an object; for a predicate, a property. Thus, the character of the indexical "I" will be the rule:

(100) Its content at a context c is the agent of c.

And the content of "I" at a context where Sam Smith is the agent will be Sam Smith. Similarly, the character of "I am swimming" will be a rule that maps contexts on to propositions, which vary systematically with the agent and time, while the content of "I am swimming" as spoken by Sam Smith at 3 PM Tuesday will be the proposition that Sam Smith is swimming at 3 PM Tuesday.

Knowledge of the character of an expression should suffice for competence with that expression. If I know that the French word "je" has the character given by (100), then I know enough to understand what French speakers are saying when they use it. Similarly, sameness of character seems to correspond to intuitive synonymy. So character can be identified with the notion of *linguistic meaning*—though of course one can know the meaning of an expression without being in a position to *articulate* its character (Kaplan, 1989, pp. 521, 577).

4.2 Circumstances of evaluation

If we settle the character of "I am swimming," then we understand how its content at a context will depend on features of that context. But we do not yet know how its *truth* will depend on features of the context. On a two-stage picture, the truth conditions of a sentence derive from the truth conditions of the proposition it expresses. The second stage of the semantics, then, describes how we get from a content to its extension (in the case of a sentence, its truth value). (See Fig. 4.2.)

In Kaplan's framework, contents have extensions only relative to a *circumstance of evaluation*:

> By ["circumstances"] I mean both actual and counterfactual situations with respect to which it is appropriate to ask for the extensions of a given well-formed expression. A circumstance will usually include a possible state or history of the world, a

Figure 4.1: Two-stage semantics.

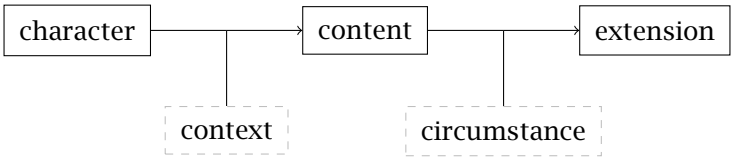

> time, and perhaps other features as well. The amount of information we require from a circumstance is linked to the degree of specificity of contents, and thus to the kinds of operators in the language. (Kaplan, 1989, p. 502)

The relativization of extensions to circumstances is orthodox and relatively uncontroversial. It makes sense to ask not only what the extension of "human being" *actually* is, but what it would have been in different circumstances. And we can model this by relativizing the "extension of" relation to circumstances. In the case of propositions, whose extensions are truth values, this amounts to a relativization of truth to circumstances. But as we have noted, it is a metaphysically innocuous relativization; we can use circumstance-relative truth as a modeling tool whatever our views about truth, possible worlds, and actuality. What is more controversial is Kaplan's decision to take *times* as a coordinate of circumstances. Because this decision has been widely discussed, and because we will be taking seriously Kaplan's invitation to countenance relativization to "others features as well," it is worth reflecting on the kinds of consideration that would tell for or against the relativization of propositional truth to times.

Kaplan's main motivation for relativizing the extensions of his contents to times is a semantic one. Following the tense logic tradition, he treats tenses as intensional operators. Just as "Joe must bake a cake" is standardly analyzed as "Must: Joe bakes a cake," so "Joe will bake a cake" is analyzed as "Will: Joe bakes a cake." Semantically, the temporal operators shift the time of evaluation, so that "Will: Joe bakes a cake" is true at t just in case the content expressed by "Joe bakes a cake" is true at some t' later than t. But this treatment of tense makes sense only if contents have truth values relative to times. As Kaplan explains:

> . . . if *what is said* is thought of as incorporating reference to a specific time, or state of the world, or whatever, it is otiose to ask whether what is said would have been true at another time, in another state of the world, or whatever. Temporal operators applied to eternal sentences (those whose contents incorporate a specific time of evaluation) are redundant. (Kaplan, 1989, p. 503)

This motivation is exactly the same as the motivation for coordinates of Lewisian indices: operators need something to shift. And it can be blocked by rejecting the treatment of tenses as intensional operators. King (2003) musters some powerful considerations against treating tenses as operators, and these undercut Kaplan's motivations for including times in circumstances of evaluation.

But it would be a mistake to suppose that these considerations settle the issue by themselves. The absence of time-shifting operators *would* demonstrate that Lewisian indices do not need a time coordinate, since the *sole* role of indices is to systematize the contributions to truth conditions of operators. But it does not show conclusively that circumstances of evaluation do not need a time coordinate. For there are constraints on circumstances of evaluation that are independent of considerations about intensional operators.[2] For the most part, these constraints concern the individuation of contents: when do we have one proposition, and when do we have many? If propositions have their truth values relative to times, then it makes sense to talk of the proposition *that Kenneth Starr is alive*, and not just of explicitly time-indexed propositions of the form *that Kenneth Starr is alive during time interval t*.

Why might it matter which ways we can talk? Propositions play a role not just in (two-stage) semantic theories, but in theories of speech acts and of mental attitudes, and issues about their individuation will have ramifications for these theories too. Consider, for example, what

[2]Stanley (2005, p. 150) argues, citing Lewis (1980): "... the difference between elements of the circumstance of evaluation and elements of the context of use is precisely that it is elements of the former that are shiftable by sentence operators." But Lewis is talking about indices, not circumstances of evaluation; indeed, a major point of his paper is that one need not have "contents" as intermediate semantic values in compositional semantics. What I am suggesting is that considerations that would tell decisively against inclusion of a coordinate in the Lewisian index need not settle the analogous issue about circumstances of evaluation.

we should say about a simple case of belief retention. Joe believed in 2004 that Kenneth Starr was alive, and he has continued to believe that Starr is alive up until this day. If we relativize the truth of contents to times, then we can see Joe as having retained through time a single propositional attitude: a belief with the tensed content *that Kenneth Starr is alive*. If we do not relativize the truth of contents to times, then we need a more complicated description of the situation. We will have to say (implausibly) that Joe has retained a belief with the (eternal) content *that Kenneth Starr is alive throughout the period P*, or (somewhat more plausibly) that Joe has had a sequence of systematically related beliefs—*that Kenneth Starr is alive at* t_1, *that Kenneth Starr is alive at* t_2, and so on—perhaps by maintaining the same representational vehicle in his "belief box." It is not my intention here to resolve this issue, but just to point out that the decisions we make about propositional truth have ramifications in propositional attitude psychology, and are therefore constrained in part by considerations from propositional attitude psychology. There might, then, be a motivation for relativizing propositional truth to times even if tenses are not intensional operators.

Similar considerations tell against the idea that we should relativize propositional truth to worlds only if the language we are considering contains world-shifting operators. Suppose we are studying a language that originally lacked modal operators, but evolved to contain them. Shall we say that speakers of the language came to express different propositions, even using sentences not containing modal operators, when the modal operators were added to the language? Or that they expressed the very same propositions, but that truth for these propositions came to be relative to possible worlds, when before it was not? Neither option seems attractive, but the idea that only world-shifting operators can justify relativizing propositional truth to worlds forces us to choose between them.

What these considerations show is that even in the absence of good candidates for intensional operators that shift a certain feature, a case might still be made on other grounds for including that feature as a coordinate of circumstances of evaluation—as something to which the extensions of contents might be relative.

It is sometimes argued that the relativity to worlds is special,

because there is always a privileged world—the actual world—that alone matters in assessing the accuracy of a claim. The thought is that the relativity of propositional truth to worlds can always be eliminated, by plugging in the actual world; whereas with times and locations, there is no unique "default value," so the relativity is more thoroughgoing.

In fact, however, the parallel between worlds and times is a good one, provided one accepts the indexical view of actuality (Lewis, 1970*a*). According to the indexical view, "the actual world" denotes the world of the context of use. So, when a speaker in world w talks about what "actually" happened, she is talking about what happens in w, not what happens in the world *we* call "actual." "Actually" is thus the modal analogue of "now," and the actual world is privileged in no deeper sense than the current time—it is the world we are in. It is sometimes thought that this conception of actuality makes sense only for modal realists, and not for modal "ersatzists," for whom "possible worlds" are abstract representations of alternative states of affairs, not concrete worlds. But this is not the case. Ersatzists can (and should) take "actually" to be indexical.[3] For even they will want to consider possible utterances in counterfactual states of affairs. And once contexts that take place at different worlds are on the table, we will need the indexical view of "actually" in order to make sense of the linguistic facts. I said "It is actually raining," and it actually was raining at the time, so I spoke truly. But if I had said this when it hadn't been raining, I would have spoken falsely—even though in fact, it actually was raining.

4.3 INCOMPLETENESS

Let us now return to the question that mainly concerns us: could the circumstances of evaluation to which propositional truth is relative include coordinates besides worlds and perhaps times—for example, tastes, standards of precision, information states, moral codes, or epistemic standards? Kaplan himself takes a fairly permissive stance on what might go into the circumstance:

> What sorts of intensional operators to admit seems to me largely a matter of language engineering. It is a question of which

[3] Stalnaker (1987, pp. 47–9) agrees, though for somewhat different reasons.

> features of what we intuitively think of as possible circumstances can be sufficiently well defined and isolated. If we wish to isolate location and regard it as a feature of possible circumstances we can introduce locational operators: 'Two miles north it is the case that', etc. ... However, to make such operators interesting we must have contents which are locationally neutral. That is, it must be appropriate to ask if *what is said* would be true in Pakistan. (For example, 'It is raining' seems to be locationally as well as temporally and modally neutral.) (Kaplan, 1989, p. 504)

If we can contenance contents that are world-neutral, time-neutral, or even location-neutral, then why should there be any conceptual obstacle to contents that are, for instance, *taste-neutral*?[4] What line is crossed when one moves from worlds to these more outlandish features?

Perhaps the most common objection is that time-neutral, location-neutral, and taste-neutral contents are "incomplete" and so not suited to play the role of propositions. An influential historical text here is this passage from Frege's unpublished article "Logic" (Frege, 1979, p. 135):

> If someone wished to cite, say, 'The total number of inhabitants of the German Empire is 52 000 000', as a counter-example to the timelessness of thoughts, I should reply: This sentence is not a complete expression of a thought at all, since it lacks a time-determination. If we add such a determination, for example, 'at noon on 1 January 1897 by central European time', then the thought is either true, in which case it is always, or better, timelessly, true, or it is false and in that case it is false without qualification.

[4] Kaplan tells me (p.c.) that although he is willing to consider times and locations as aspects of circumstances, he would himself draw the line at tastes and standards, on the grounds that these are too subjective and "perspectival" to be "features of what we intuitively think of as possible circumstances."

Interestingly, King (2003, p. 57), who argues against Kaplan's relativity to times, is open-minded about whether propositions might have truth values relative to *standards of precision*: "To say that the object of one of my beliefs is the claim that France is hexagonal, and that whether what I believe is true or false depends not just on what the world is like, but also on how much precision we require seems completely unobjectionable."

By "thought" here, Frege means essentially what we mean by "proposition": the sense of a declarative sentence, which he also takes to be the content of a propositional attitude. So one might expect this passage to give aid and comfort to the modern opponent of time-neutral propositions. But this is doubtful.

First, his discussion seems to confuse two very different issues. The paragraph from which our passage is taken begins with the issue of whether the propositions themselves are abstract and unchanging or concrete and mutable:

> Whereas ideas (in the psychological sense of the word) have no fixed boundaries, but are constantly changing and, Proteus-like, assume different forms, thoughts always remain the same. It is of the essence of a thought to be non-temporal and non-spatial.

And, after our passage, Frege returns to the theme of whether the propositions (thoughts) change. But this issue is orthogonal to the issue that concerns us—whether the truth of a proposition is relative to times. A temporalist can very well agree with Frege that propositions are abstract, timeless entities; that is perfectly compatible with the view that their truth values vary with the time of evaluation.

Second, Frege doesn't really give a reason for thinking that a complete expression of a thought must include a time determination. An intelligible motivation might be extracted from his general view that sense determines reference, and hence that thoughts or propositions determine truth values. This can be read in such a way as to exclude the temporalist view, on which a proposition determines only a function from times to truth values. But, read in this way, it also seems to exclude the modern eternalist view, on which a proposition determines a function from worlds to truth values. We can make the Fregean slogan compatible with contemporary eternalism by reading it as follows: a sense, together with the state of the world, determines a truth value. But we can also make it compatible with temporalism: a sense, together with the state of the world at a time, determines a truth value.

One might try to cash out Frege's "incompleteness" worry in the following way. Propositions are supposed to be the contents of beliefs and other propositional attitudes. But if we specify the content of someone's belief in a way that leaves it open whether the belief is accurate, we have not given its complete content. Thus, for example,

if we don't know whether the accuracy of Sam's belief that it is 0° C depends on the temperature in London on Tuesday or the temperature in Paris on Wednesday, we don't yet have the full story about what it is that Sam believes. Similarly, if we don't know whether the accuracy of Sam's belief that Hen-of-the-Woods is tasty depends on how Hen-of-the-Woods affects Sam or on how it affects Sal, then we don't yet know what it is that Sam believes. A location-neutral, time-neutral, or taste-neutral content would only incompletely determine the conditions for an attitude to be accurate, and so could not be the complete content of the attitude.

But this line of thought proves too much. For surely the accuracy of *any* contingent belief depends on features of the world in which the believer is situated (the world of the context of use). Even if we specify the content of Sam's belief in a way that builds in time and place—*that it is* 0° *C at the base of the Eiffel Tower at noon local time on February 22, 2005*—it is still not determined whether the accuracy of his belief depends on the temperature in Paris in world w_1 or on the temperature in Paris in world w_2. To know that, we would have to know not just what Sam believes—the content of his belief—but in what context, and in particular in what world, the belief occurs.

One might respond to these considerations by bringing the world of the context of use into the *content* of Sam's thought.[5] Intuitively, though, Sam could have had a thought with the very same content even if the world had been very different. Our ordinary ways of individuating thought contents do not support making the world of the context of use part of the content, except in exceptional circumstances. Moreover, bringing the world of the context into the content of Sam's thought would make this content a necessary truth about this possible world, rather than a contingent truth about the weather in Paris. We should not say, then, that Sam's thought is *about* the world of use. It is not *about* any particular world. Acknowledging the fact that it depends for its truth on the world of use, we may adopt John Perry's terminology and say that it *concerns* the world of use (Perry, 1986).

The objection from "incompleteness" may be motivated, in part, by an appreciation of the fact that the truth predicate we use in ordinary speech is monadic. We don't characterize claims as "true-in-w," or as

[5]As urged by Jonathan Schaffer in recent, unpublished work.

"true-in-w-at-t-on-s," but as "true" (simpliciter). But this no more shows that propositional truth is not relative to parameters than the fact that we normally say it's "3 PM," and not "3 PM Pacific Daylight Time," shows that the time of day is not relative to a time zone. The monadic predicate "true" is just another piece of vocabulary whose intension we can characterize using the relation of truth at a circumstance of evaluation (see §4.7, below).

Once we accept the relativity of propositional truth to worlds, we have accepted a kind of "incompleteness." We have accepted the idea that both the content of an assertion or belief and its context must be taken into account in assessing it for accuracy. The question is just *which* features of which contexts must be taken into account, and how. (This topic will be taken up in the next three sections.)

4.4 CIRCUMSTANCES AND CONTEXTS

We have now looked separately at the two stages of a two-stage semantic theory. The first stage assigns contents to expressions relative to a context of use, and the second assigns truth to these contents relative to a circumstance of evaluation. But in order to understand how the two-stage semantics relates to our Lewis-style one-stage semantics, we need to see how these two stages can be fit together to yield a definition of truth for sentences at contexts.

Here is how Kaplan fits them together:

> If c is a context, then an occurrence of [a sentence] ϕ in c is true iff the content expressed by ϕ in this context is true when evaluated with respect to the circumstance of the context. (Kaplan (1989, p. 522); cf. the formal version on 547)

Kaplan is entitled to talk of "*the* circumstance of the context," because his circumstances of evaluation are composed of a world and a time, and a context of use determines a unique world and time. But in the interest of full generality, we should not assume that context will always pick out a unique circumstance of evaluation. For example, in a framework with *overlapping* worlds or histories, a possible utterance will be contained in multiple overlapping worlds, so there will be no unique "world of the context of use." For this reason, I prefer to talk

of "all circumstances of evaluation compatible with the context" rather than of "the circumstance of the context."[6] Thus:

(101) A sentence *S* is true at context *c* iff the proposition expressed by *S* in *c* is true at all circumstances of evaluation compatible with *c*.[7]

(What "compatibility" amounts to must be worked out in detail for each semantic theory.) It will also be useful to define propositional truth relative to a context:

(102) A proposition *p* is true at a context of use *c* iff *p* is true at all circumstances of evaluation compatible with *c*.

We can now rephrase our definition of sentence truth at a context as follows:

(103) A sentence *S* is true at context *c* iff the proposition expressed by *S* in *c* is true at *c*.

It may seem strange to talk of a proposition being true at a context of use, because a proposition is not "used" in the way that a sentence is. But the notion makes perfect sense: as we have just seen, there is good reason to keep track of the truth of propositions relative to contexts at which *sentences* might be used.[8] And, in an extended sense, we can think of assertions or beliefs as "uses" of the propositions asserted or believed.

4.5 Nonindexical contextualism

It should be clear from (101) and (103) that the context of use plays two distinct roles in the definition of sentence truth at a context. It plays a *content-determining* role, since a sentence will express different propositions at different contexts. And it plays a *circumstance-determining*

[6] Compare our earlier discussion of Lewis on "the index of the context," p. 76, above.

[7] Kaplan suggests in a footnote that it "seems necessary for the definition of truth" that "a circumstance is an aspect of the context" (Kaplan, 1989, 511 n. 35). The definition given here shows that this is not so.

[8] Even in the case of sentences, talk of *S* being true at context *c* carries no commitment to there being an actual *use* or *utterance* of *S* at *c*. See Kaplan (1989, p. 522).

role, selecting the circumstances of evaluation that are relevant to the truth of an occurrence of a sentence at the context.[9]

What this means is that there are two distinct ways in which a sentence can be context-sensitive. It can be sensitive to a feature of context because that feature plays a content-determining role, or because it plays a circumstance-determining role. To see this point is to see that use sensitivity and use indexicality come apart:

(87) An expression is *use-sensitive* if its extension (relative to a context of use and context of assessment) depends on features of the context of use.

(104) An expression is *use-indexical* iff it expresses different contents at different contexts of use.[10]

(89) An expression is *F-use-sensitive* if its extension (relative to a context of use and context of assessment) depends on the *F* of the context of use.

(105) An expression is *F-use-indexical* iff the content it expresses at a context depends on the *F* of that context.

"I am over five feet tall" is both use-sensitive and use-indexical; it expresses different propositions (with different truth values) at different contexts of use. (To be precise, it is *agent*-use-indexical, because the content it expresses depends on the agent of the context; and on some views, it is also *time*-use-indexical.) But a sentence can be use-indexical without being use-sensitive, and even *F*-use-indexical without being *F*-use-sensitive. The sentence

(106) If it is raining now, it is raining.

is true at every context of use, and so not use-sensitive, but because it contains "now," it is (time-)use-indexical. This basic point is well known from Kaplan (1989), who argues that certain sentences containing

[9] For the point, see Belnap, Perloff, and Xu (2001, pp. 148–9), MacFarlane (2005*a*, pp. 326–7), Lasersohn (2005, p. 663).

[10] The sense of "indexical" defined by (104) is quite broad. It does not distinguish between different mechanisms by which an expression might express different contents at different contexts. Sometimes "indexicality" is used in a narrower sense to cover just some of these mechanisms (Stanley, 2000, p. 411). If you like, call the sense defined by (104) "broad indexicality."

indexicals, like "I am here now," can be logically true, or true at every context of use. Less well known is that the converse point also holds. A sentence can be use-sensitive without being use-indexical, and a sentence can be *F*-use-sensitive without being *F*-use-indexical: its truth value can depend on a feature of the context of use even though its content does not depend on this feature.

Here's an example. Suppose we hold (with Kaplan (1989) and other *temporalists*) that the contents of sentences have truth values relative to worlds and times. Then we will naturally take

(107) Socrates is sitting

to express, at every context of use, a time-neutral proposition—one that is true relative to some times of evaluation and false relative to others. Because we take (107) to express the same proposition at every context, we will not take it to be indexical. But we will still take it to be context-sensitive, in a way, since we take the truth of a use of this sentence to depend on the time of the context of use.[11] (101) shows how this is possible: the temporalist need only say that a circumstance of evaluation $\langle w, t \rangle$ is compatible with a context c just in case w is the world of c and t is the time of c, and the truth value of tensed sentences will depend on the time of the context of use, even if the content does not.[12]

[11]Some philosophers may take issue with my claim that, even on the temporalist analysis, (107) is context-sensitive. Indeed, many *define* "context-sensitive" the way I have defined "use-indexical" (Cappelen and Lepore (2005, p. 146), Stanley (2005, p. 16), Soames (2002, p. 245)). I would certainly concede that context sensitivity in the ordinary sense is not the same as use sensitivity: "I am here now" is not use-sensitive, but it is context-sensitive in the ordinary sense; conversely, "The population of China is 2002 is greater than one billion" is use-sensitive (since its truth varies with the world of the context of use) but not context-sensitive in the ordinary sense. But I would also maintain that context sensitivity in the ordinary sense is not the same as use indexicality. For example, it seems crazy for a temporalist to deny that tensed sentences are context-sensitive, even though, on the temporalist view, such sentences are not taken to be use-indexical. So it seems to me that the ordinary use of "context-sensitive" does not precisely track either "use-sensitive" or "use-indexical."

[12]Compare Percival (1989, pp. 193–5), defending the temporalist theory against Mellor (1981)'s objection that if tensed sentence expressed the same proposition at every time, all actual occurrences of this sentence should have the same truth value, regardless of the context in which they occur.

Thus, for the temporalist, (107) will be use-sensitive (specifically *time*-use-sensitive), but not use-indexical. The temporalist and the eternalist can agree that tensed sentences are use-sensitive—indeed, they can agree about what truth values such sentences get relative to every context of use—while disagreeing about whether this use-sensitivity derives from the context's content-determining role or its circumstance-determining role.

The template we see in temporalism—use-sensitivity sustained not by use-indexicality but by the circumstance-determining role of the context of use—can be reproduced for other parameters of propositional truth. For example, suppose we hold that propositions have truth values relative to worlds and aesthetic standards. We need not say that sentences like "The Mona Lisa is beautiful" are assessment-sensitive. For we could take the context of use to determine values for both parameters. We could say that a sentence is true at a context of use c just in case the proposition it expresses at c is true relative to the world of c and the aesthetic standard relevant at c. (This would amount to saying that a circumstance $\langle w, s \rangle$ is compatible with a context of use c just in case w is the world of c and s is the aesthetic standard relevant at c.)

The resulting position would resemble contextualist approaches in taking the truth of sentences about what is "beautiful" to depend on the taste of the speaker. But, unlike standard forms of contextualism, it would not take the *content* of such sentences to depend on the taste of the speaker. Because this is a view on which such sentences are aesthetic-standard-use-sensitive but not aesthetic-standard-use-indexical, it is aptly characterized not as relativism but as a kind of *nonindexical contextualism*.[13] However, I do not want to get too caught up in disputes about labeling. The important thing to see is that the position just described would have much in common with more standard forms of contextualism. To be sure, it would disagree with ordinary contextualism about the contents of aesthetic claims. But it would agree with ordinary contextualism on every question about the truth of sentences, and like standard contextualism it would give every use of a proposition an absolute truth value. It would remain on the safe side of the really interesting line—the line between use sensitivity

[13] See MacFarlane (2009) for a fuller discussion, with examples.

and assessment sensitivity. It is only when that line is crossed that we get a philosophically interesting and controversial notion of "relative truth." Even when we are talking about propositional truth, then, it is not *what* propositional truth is relativized to—worlds, times, tastes, standards—that matters most, but *how* it is relativized.

It is therefore unfortunate that many recent critiques of "relativist" doctrines in semantics characterize truth relativism as the relativization of propositional truth to something besides possible worlds.[14] If I am right, these critics are missing the interesting target.

4.6 Relativism

To make room for asssessment sensitivity in our two-stage semantic framework, we need truth to be relative to both a context of use and a context of assessment:

(108) A proposition p is true at as used at c_1 and assessed from c_2 iff p is true all circumstances of evaluation compatible with $\langle c_1, c_2 \rangle$.

(109) A sentence S is true as used at c_1 and assessed from c_2 iff the proposition expressed by S in c_1 as assessed from c_2 is true as used at c_1 and assessed from c_2.

The relation of "compatibility" now holds between circumstances and a *pair* of contexts—a context of use and context of assessment. As before, it must be defined on a per-application basis. Thus, for example, a relativist about aesthetic vocabulary who holds that circumstances of evaluation are world/aesthetic standard pairs might say that

(110) A circumstance $\langle w, s \rangle$ is compatible with $\langle c_1, c_2 \rangle$ iff w is the world of c_1 and s is the aesthetic standard relevant at c_2.

The upshot of this definition is that uses of propositions in context cannot be assigned truth values absolutely, but only relative to contexts of assessment. When we asssess an assertion, made yesterday by Ted, that the Mona Lisa is beautiful, what matters for its truth is not Ted's

[14] See for example Zimmerman (2007, p. 316), Stanley (2005, p. 137), Glanzberg (2007, p. 2), Cappelen and Hawthorne (2009).

aesthetic standards but our own. So, we say that Ted has spoken truly if the Mona Lisa is beautiful by *our* standards.

In a two-stage framework, there are two different ways in which a sentence might be assessment-sensitive. It might be assessment-sensitive because it is *assessment-indexical*.

(111) An expression is *assessment-indexical* iff it expresses different contents relative to different contexts of assessment.

(112) An expression is *F-assessment-indexical* iff the content it expresses as assessed from c depends on the F of c.

A word like "noy" (§3.2.3, above) would be (time-)assessment-indexical if we took it as a referring expression whose denotation, relative to a context of assessment, is the time of that context. A particular use of the sentence "It is 2 PM noy" would express different propositions as assessed from different contexts; there would be no "absolute" characterization of its content. We can call views that take some expressions to be assessment-indexical forms of *content relativism*.[15]

But, just as a sentence can be F-use-sensitive without being F-use-indexical, so a sentence can be F-assessment-sensitive without being F-assessment-indexical. It can do so by expressing a proposition that is itself F-assessment-sensitive:

(113) A proposition is assessment-sensitive if its truth value as used at c_1 and assessed from c_2 depends on features of c_2.

(114) A proposition is F-assessment-sensitive if its truth value as used at c_1 and assessed from c_2 depends on the F of c_2.

A plausible form of relativism about the tasty would take this form. According to such a view, which we can call *truth-value relativism*, the sentence "Hen-of-the-Woods is tasty at time t" expressess the same proposition relative to every context of use and context of assessment, but this proposition—the proposition that Hen-of-the-Woods is tasty at time t—is itself (taste-)assessment-sensitive, since its truth value (relative to a context of assessment) depends on the assessor's tastes.

[15]This term comes from Egan, Hawthorne, and Weatherson (2005). Note that MacFarlane (2005*a*) uses "expressive relativism" for this, and "propositional relativism" for what Egan, Hawthorne, and Weatherson call "truth-value relativism." I now prefer, and use, their terminology. For a similar distinction, see Percival (1994, pp. 192–3).

The views that will be developed in what follows are all forms of truth-value relativism, which I think is better suited to satisfy the motivations pushing us towards relativism. We will briefly consider content relativism in §??.

4.7 Monadic "true" and the Equivalence Schema

Relativism, as developed here, is the view that truth-conditional semantics should have as its output a definition of truth relative to a context of use and context of assessment. To resist relativism would be to defend the usual view that we need only truth relative to a context of use. But whichever view we take, the context-relativized truth predicate used in semantics is a technical term, which gets its meaning in part from an account of its pragmatic relevance (for example, in Lewis's theory, the view that speakers at c try to assert what is true at c, and trust others to be doing so). It is not the ordinary truth predicate used in everyday talk—a *monadic* predicate that applies to propositions, and is governed by the

> Equivalence Schema The proposition that Φ is true iff Φ.[16]

As we saw in §2.3, some philosophers have thought that the Equivalence Schema is somehow incompatible with relativism about truth. We noted there that there is no obvious inconsistency in taking the everyday monadic truth predicate to be governed by the Equivalence Schema while employing various kinds of relative truth predicates in one's semantic theorizing. A bit of unfinished business, though, was to say in more detail how these truth predicates relate. In particular, what account should the relativist give of monadic "true"?

[16]Note that the truth predicate needs to be monadic in order for "disquotation" to make sense. Once *any* relativization is added, it no longer makes sense to disquote: how would we continue "the proposition that Φ is true at w iff ..."? Perhaps an analogue of disquotation could be preserved here, by filling in the "..." with "in w, p," and understanding "in w" along the lines of "in Australia" (cf. Lewis, 1986, pp. 5–7). But then it is not clear why someone who relativized propositional truth to tastes could not simply use "by taste t" to the same purpose.

For the relativist, the monadic predicate "true" is just another predicate of the object language—the language for which she is giving a semantics. The natural semantics for it is this:[17]

(115) The extension of "true" at a circumstance of evaluation e is the set of propositions that are true at e.

Given this clause for "true", every instance of the Equivalence Schema will be true at every circumstance of evaluation, and hence also at every context of use and context of assessment.[18] And, if the language can express any assessment-sensitive propositions, "true" will also be assessment sensitive, since if p is assessment-sensitive, the proposition that p is true must be assessment-sensitive too.[19]

So the relativist can fully vindicate the Equivalence Schema, and the argument that relativism is incompatible with it falls flat. But there *is* a legitimate concern in the vicinity. Granted that our doubly relativized truth predicate is not the ordinary (monadic) truth predicate we use in ordinary speech, but a piece of *technical* vocabulary, we need to say something about how it is connected up with other parts of our theories of language and communication, so we can see the practical significance of going for a relativist semantic theory as opposed to a nonrelativist one. I want to emphasize, though, that this is a burden faced by nonrelativists, too—by anyone who uses "true at a context" in a truth-conditional semantic theory. (The point goes back at least to Dummett (1959).) We will return to this issue in chapter 5.

4.8 TRUTH BEARERS REVISITED

4.8.1 Newton-Smith

We are now in a position to revisit Newton-Smith's oft-cited objection to relative truth (§2.1). Here is the argument, as well as I can reconstruct

[17] It is, of course, a naive semantics, in the sense that it provides no solution to the semantic paradoxes. I am assuming, perhaps rashly, that the issues raised by the paradoxes are orthogonal to those we are worried about here, and can be dealt with separately.

[18] To see this, note that whatever circumstance e we choose, the right and left hand sides of the biconditional will have the same truth value at e. I assume here that "the proposition that Φ" rigidly denotes a proposition.

[19] Incidentally, this shows what is wrong with the thought that relativism about truth amounts to nothing more than an ordinary sort of contextualist (use-sensitive) semantics for "true."

it from the passage quoted on page 44, above:

1. Suppose, as the relativist holds, that there are sentences S_1 and S_2, a proposition p, and contexts Ψ and Θ[20] such that:
   a) S_1 is true in Ψ,
   b) S_2 is not true in Θ.
   c) p is expressed by S_1 in Ψ and by S_2 in Θ.
2. If S_1 and S_2 have different truth-conditions, they do not express the same proposition. (premise)
3. Hence S_1 and S_2 have the same truth-conditions. (by 2, 1c)
4. If S_1 and S_2 have different truth values, then they have different truth-conditions. (premise)
5. S_1 and S_2 have different truth values (by 1a and 1b).
6. So S_1 and S_2 have different truth-conditions. (by 4, 5)
7. This contradicts (3). So, by reductio, the clauses of (1) cannot all be true.

There are several major problems with this argument. First, if (2) is to support the move from (1c) to (3), it must be construed as follows:

2′. If S_1 and S_2 have different truth-conditions, then for all contexts c_1, c_2, the proposition S_1 expresses at c_1 is not the same as the proposition S_2 expresses at c_2.

But it is hard to see what notion of "truth-condition" for a sentence would make this claim plausible, even for a non-relativist. In general, sentences have truth values only relative to contexts of use. So the only reasonable notion of truth-condition for a sentence is the condition a *context* must satisfy in order for the sentence to be true (see above, §3.2.1). In this sense of "truth-condition," the sentences "I am here now" and "He was there then" have different truth-conditions, but for all that it may be the case that the proposition expressed by the former in one context is the same as the proposition expressed by the latter in another context (later, with a different speaker demonstrating the

[20] Newton-Smith says "social groups or theories," but I will present his argument in a more general way, since nothing in it depends on what the relevant features of contexts are—the agent's social group, theory, perspective, or whatever.

first speaker). Even without presupposing the coherence of a relativist framework, then, we can reject this step of Newton-Smith's argument.

Step (5) is also problematic. (1a) and (1b) use a context-relative truth predicate, "true in c." (It is unclear whether Newton-Smith is thinking here of relativity to a context of use or to a context of assessment; for now, it will not matter.) But (5), which is supposed to be a trivial consequence of these two premises, just speaks of "having different truth values." If (5) is to follow from (1a) and (1b), it must be construed as follows:

5′. The truth value of S_1 in Ψ is different from the truth value of S_2 in Θ.

But then, in order to get (6), we will need to construe premise (4) as

4′. If S_1 and S_2 have different truth values relative to different contexts, then they have different truth-conditions.

How plausible is (4′)? Not very plausible, if the truth-condition of a sentence is a condition a context must satisfy in order for the sentence to be true at that context. The English sentence "I am in Madrid" and the Spanish sentence "Estoy en Madrid" have different truth values relative to different contexts of use (imagine one speaker in Madrid, the other in Manhattan). But that does not mean that they have different truth-conditions; they are in fact synonymous. However, because Newton-Smith is considering a case where the two sentences express the same proposition, we can strengthen our premise as follows:

4″. If S_1 and S_2 have different truth values relative to different contexts, but express the same proposition at those different contexts, then they have different truth-conditions.

Strengthening the premise still leaves us with a valid argument, since we can use (1c) as an additional premise for the move to (6).

How plausible in (4″)? Here we need to decide whether "relative to" in (4″) means "as used at" or "as assessed from:"

$4''_u$. If for some contexts c_0, c_1, c_2, S_1 and S_2 have different truth values as used at c_1 and c_2 respectively and assessed from c_0, but express the same proposition at c_1 and c_2, then S_1 and S_2 have different truth-conditions.

$4''_a$. If for some contexts c_0, c_1, c_2, S_1 and S_2 have different truth values as used at c_0 and assessed from c_1 and c_2, respectively,

> but express the same proposition at c_0, then S_1 and S_2 have different truth-conditions.

Counterexamples to both versions are easy to produce. Suppose we are temporalists. Let S_1 = S_2 = "Socrates is sitting," which we will take to express the same proposition at noon and at midnight. S_1 may yet be true as used at noon, but false as used at midnight. So according to ($4''_u$) S_1 should have different truth-conditions than S_2—that is, different truth-conditions than itself. Clearly this is incoherent, and ($4''_u$) should be rejected. (If we are eternalists, we can construct a similar counterexample using a contingent sentence and two contexts located in different possible worlds.)

To get a counterexample to ($4''_a$), suppose we take S_1 = S_2 = "Hen-of-the-Woods is tasty," and suppose we take this sentence to express the same taste-variable proposition at every context of use (and context of assessment). Then S_1 will be true as assessed by Sal but S_2 will not be true as assessed by Sam. For all that, S_1 and S_2 do not have different truth-conditions, since they are the same sentence.[21] Newton-Smith may reply that this is a counterexample not to ($4''_a$) but to the very idea of taste-variable propositions, or of assessment sensitivity. But no non-question-begging reason has been given for accepting ($4''_a$), and as we have seen, similar principles are demonstrably false even in non-relativist frameworks.

I have discussed Newton-Smith's argument at some length because I think it exemplifies a general tendency in much of the literature on relative truth. Terms like "true in," "true for," and "truth-conditions" are deployed without any sensitivity to the various *kinds* of relativization of truth that are used in semantics. As we have seen, proper statement of a relativist position requires some care. A general argument against relative truth needs to take the same care.

4.8.2 The analytic argument

Another form of argument against the idea that the truth of a proposition can depend on anything other than the state of the world is an analytic one: "By 'proposition' we just *mean* something that partitions

[21] A relativist can think of a sentence's truth condition as a condition that a *pair* of contexts $\langle c_1, c_2 \rangle$ must satisfy if the sentence is to be true as used at c_1 and assessed from c_2.

the possible world-states into those at which it is true and those at which it is not true. So if your 'propositions' do not do that, they are not propositions properly so-called."

Clearly this verbal stipulation does nothing to settle the real question at issue, which is whether the contents of assertions, beliefs, and other so-called "propositional attitudes" have truth values relative to worlds only, or relative to worlds and something else. Someone who takes it to be analytic of "proposition" that propostions have truth values relative to worlds only will take this to be a question about whether the contents of these attitudes are propositions or something else. That is how Lewis (1979*a*) sees it; he concludes that the content of attitudes should be thought of as self-ascribed *properties*—which have truth values relative to a world, time, and agent—and not as propositions. I prefer to think of propositions as whatever the contents of these attitudes turns out to be; our views about propositional truth will then depend on more general theoretical considerations about the attitudes and speech acts. But nothing much hangs on this verbal decision.

In any case, as I have argued, the real crux is assessment sensitivity, and the relativity of propositional truth to things besides worlds is neither necessary nor sufficient for assessment sensitivity. It is not sufficient, because it is compatible with nonindexical contextualism, which does not countenance assessment sensitivity. Temporalists, who relativize propositional truth to times, and Lewis, who relativizes it to worlds, times, and agents,[22] can all agree that particular "uses"—particular assertions and beliefs with these propositions as their contents—have their truth values absolutely. It is also not necessary, because, as we will see in chapter **??**, one can describe a view on which even standard possible-worlds propositions are assessment-sensitive. The real issue is not the nature of the circumstances to which propositional truth is relative, but whether propositional truth is assessment-sensitive.

[22]Not, of course, under the label "propositional truth," but see above.

5 *Making Sense of Relative Truth*

In chapters 3 and 4, we have seen how the relativist's thesis ought to be stated. But do we really understand it? In order to understand it, we must grasp what is meant by "true as used at c_1 and assessed from c_2." And it is not clear that we do. For it is not clear that the concept of truth *admits* of relativization to assessors. Meiland (1977) states the problem very clearly as a dilemma. If "true" as it occurs in "true for X" is just the ordinary, nonrelative truth predicate, then it is unclear what "for X" adds.[1] On the other hand, if the occurrence of "true" in "true for X" is like the "cat" in "cattle"—an orthographic, not a semantic, part—then the relativist needs to explain what "true-for-X" means and what it has to do with truth, as ordinarily conceived. Meiland's own solution—explicating "true for X" as "corresponds to reality for X"—just pushes the problem back a level. The absolutist can say: my understanding of "correspondence to reality" leaves no room for an added "for X," so the proposed explicans is just as mysterious as the explanandum.

This, I think, is the hardest question for the relativist. Is assessment-sensitivity really intelligible? Do we have enough grip on the notion of assessment-sensitive truth to understand what relativist proposals in specific areas—say, predicates of personal taste or future contingents—amount to? Do we understand the practical difference between relativist and nonrelativist proposals sufficiently to tell what evidence would count in favor of each?

[1] As noted in §2.4, "true for X" can be used to specify the domain of a generalization or to say how things are "by X's lights," but neither of these uses captures what the truth relativist is aiming at.

5.1 A strategy

Relativists commonly try to meet this challenge by giving a *definition* of truth that makes its assessment-relativity plain. If truth is idealized justification, then, as we observed in §2.4, it might reasonably be thought to be assessor-relative, since ideal reasoners with different beliefs, propensities, or prior probabilities might take the same ideal body of evidence to support different conclusions. Similarly, if truth is defined pragmatically, as what is good to believe, then it might also be assessor-relative, insofar as different things are good for different assessors to believe. But although these coherentist and pragmatic definitions of truth capture the "relative" part of "relative truth," I do not believe they capture the "truth" part. Like Davidson (1997), I doubt that the concept of truth can be usefully illuminated by a definition in terms of more primitive concepts.

Of course, the relativist semanticist can give a formal definition of "true as used at c_1 and assessed from c_2" that fixes its extension over a particular class of sentences and contexts. But such a definition would not answer the challenge, for reasons Michael Dummett made clear in his classic paper "Truth" (Dummett, 1959). If our aim in giving a Tarskian truth definition is to explain the meanings of expressions by showing how they contribute to the truth conditions of sentences containing them, then we must have a grasp of the concept of truth that goes beyond what the Tarskian truth definition tells us. A recursive definition of "true in L" cannot simultaneously explain both the meanings of the expressions of L *and* the meaning of "true in L." It is only if we have some antecedent grasp of the significance of "true in L" that an assignment of truth-conditions can tell us something about the meanings of sentences and subsentential expressions.

Dummett illustrates his point by considering the concept of winning in a game—say, chess. Here is one kind of definition of "winning in chess:"

(116) White wins at chess just in case the current disposition of pieces on the board has been reached by a series of legal chess moves, with White and Black alternating, and Black's king is in checkmate.

a. Black's king is in checkmate iff Black's king is in check and Black has no legal move available that would result in Black's king not being in check.
b. Black's king is in check iff one of White's pieces could capture Black's king if it were White's move.
c. A chess move is legal iff ...

Someone who knew this definition would be in a position to tell when White had won a game of chess. But if she had *only* this knowledge, she would be missing a crucial aspect of the concept of winning: that winning is what one *aims* at in playing a game.[2] One can imagine a Martian who knows which chess positions are "winning" ones but believes, perhaps, that one's aim in playing chess one is *not* to reach a "winning" position. The Martian would have a formally correct definition of winning (in the mathematician's sense), but would not grasp the concept.[3]

In the same way, Dummett suggests, someone who had an extensionally correct Tarskian truth definition for a language but did not understand the *significance* of characterizing sentences as true would not grasp the concept of truth. Imagine, again, a Martian who has a correct definition of truth at a context of use for a language but thinks that speakers conventionally try to *avoid* uttering true sentences, and take others to be doing the same. The Martian's knowledge of the truth-conditions of sentences would not enable it to use these sentences to say anything, or to understand others' uses.

Dummett summarizes the general point as follows:

> If it was to be possible to explain the notion of meaning in terms of that of truth, if the meaning of an expression was to be regarded as a principle governing the contribution that it made to determining the truth-conditions of sentences containing it,

[2] What this means is presumably that one must *represent oneself* as trying to win in order to count as playing chess. One need not *actually* be trying to win—one might be trying to make one's opponent feel better by throwing the game, for example.

[3] One might object: isn't it at least conceivable that one day we should all begin to play games to lose? Dummett would say that we are really conceiving of a scenario in which (a) we have changed what counts as winning in all these games, so that what formerly counted as losing now counts as winning, and (b) we have started to use the word "lose" to mean what "win" used to mean. See Dummett (1981, p. 320).

> then it must be possible to say more about the concept of truth than under which conditions it applied to given sentences. Since meaning depends, ultimately and exhaustively, on use, what was required was a uniform means of characterising the use of a sentence, given its truth-conditions. (Dummett, 1978, p. xxi)[4]

This "uniform means of characterising the use of a sentence, given its truth-conditions" would be an account of the various illocutionary forces with which we can put forth sentences—for example, assertoric force: "corresponding to each different kind of force will be a different uniform pattern of derivation of the use of a sentence from its sense, considered as determined by its truth-conditions" (Dummett, 1981, p. 361). Hence, "[w]hat has to be added to a truth-definition for the sentences of a language, if the notion of truth is to be explained, is a description of the linguistic activity of making assertions. . . " (Dummett, 1978, p. 20). Although Dummett acknowledges that this task is one of "enormous complexity," he does propose, as one example of the shape such an account might take, that assertoric utterances are governed by the convention that one should utter only true sentences.[5] This is certainly a reasonable candidate for the knowledge that the Martian observer would need to acquire in order to use its correct specification of the truth conditions of English sentences to understand and speak to English speakers.

We will discuss this specific proposal in more detail shortly, but two general points are worth noting now. First, Dummett has given an example of an explication of "true (at c)" that does not take the form of a definition. Instead of defining "true (at c)," Dummett proposes to illuminate it by describing its role in a broader theory of language use—in particular, its connection to the speech act of assertion. As Davidson (1997) points out, most philosophically interesting concepts are not definable in simpler terms, but they can still be illuminated by articulating their theoretical connections to other concepts.

Second, if Dummett is right, then it is not just the relativist who owes an explication of the significance of her truth predicate. The absolutist owes one as well—at least if she is to use this predicate in

[4]For similar points, see Wiggins (1980) and Davidson (1990, p. 300).

[5]Dummett (1981, p. 302); compare Lewis (1983, §III) and Lewis (1980, §2), discussed in §3.2.1, above.

semantics.[6] So although Percival (1994, p. 208) is quite correct to say, of truth relativism, that "[i]n the absence of the clear statement of this doctrine's consequences for the evaluation of utterances, it is empty and worthless. . . ," the same could be said of *any* use of truth in giving a theory of meaning. It may be that the task is easier to discharge for the non-relativist, but the task is the same for both sides.

These two points suggest a strategy for the truth relativist. Start with an account of assertoric force (and the other illocutionary forces) that is acceptable to the nonrelativist. Such an account ought to explicate "true at c" by relating it to proprieties for the use of sentences or propositions. Then extend this to an explication of "true as used at c_1 and assessed from c_2" by finding a natural role for contexts of assessment to play in the account of assertoric force. Take care not to impose any explanatory demand on the truth relativist that it would not be fair to impose on the non-relativist engaged in the same general project of giving truth-conditional semantics for a class of expressions. If this strategy is successful, the relativist should be able to say to the absolutist: "If you can make sense of your absolute truth predicate, you should be able to make sense of my relative one, too, and see why it deserves to be called a *truth* predicate."

Since there is no universally accepted account of assertoric force, we will adopt "big tent" approach in this chapter, and consider the four most prominent kinds of account: the idea that assertion constitutively aims at truth, the idea that it constitutively aims at knowledge, the idea that assertion is the expression of belief, and the idea that assertion is a kind of commitment or responsibility-taking. We will see that assessment-relative truth can be illuminated in the context of the first, second, and fourth of these accounts, and that the third account is independently problematic. The upshot is that our ability to "make sense of relative truth" is not hostage to a particular story about

[6] *Semantic deflationists* hold that there is nothing more to the concept of truth than its role as a device for semantic ascent—a role that is captured (for a given language) by a Tarskian truth definition. The Dummett argument, if it is correct, shows that semantic deflationists should not use truth definitions to give the meanings of expressions. Most deflationists have accepted this argument, and consequently favor inferentialist explications of meaning over truth-conditional ones (Brandom, 1994; Field, 1994; Horwich, 1998). For a dissenting view, see M. Williams (1999); for a recent defense of the argument, see Patterson (2005).

assertoric force, but is compatible with the most plausible kinds of account.

5.2 The truth rule

Dummett's analogy with games suggests that the connection between truth and assertion is teleological: in making assertions, one aims to put forward truths and not falsehoods. No doubt there is something right about this, but as it stands it is not very helpful. The problem is not just that it seems perfectly possible to aim to assert a falsehood. For it may be a standing convention or expectation that one aims to assert only truths, so that in asserting one *represents oneself* as aiming to assert something true, even when one isn't (Dummett, 1981, pp. 299–301). A more fundamental problem is that truth is not the *only* thing it is conventional to aim at in making assertions. One is also expected to assert only that for which one has good evidence, and that which is relevant for the purposes of the conversation. Dummett himself notes that it is absurd to think that one could get a grip on the notion of truth simply by being told that it is the aim of assertion (Dummett, 1978, p. 20).

A more plausible way of getting at the root idea is by giving a normative account of assertion, in terms of what Williamson (1996) calls the "truth rule:"

(117) Assertion is the unique speech act whose sole constitutive rule is: assert (at a context c) only propositions that are true (at c).[7]

To say that the truth rule is *constitutive* of assertion is to say that nothing that is not subject to this rule can count as an assertion. It is crucial to this approach that there is a distinction between the "constitutive rules" that define the move of assertion and other kinds of norms. We can make such a distinction in the case of other game moves. For example, the rule of chess that says you can't castle if the king is in check is partially constitutive of the move of castling. A move

[7] Or, perhaps better: "whose sole constitutive rule, other than rules constitutive of action in general" (of the sort envisioned by Kant and Korsgaard). I am grateful to Eugene Chislenko for the point.

that was not subject to this rule would not be castling.[8] Since castling is nothing more than a move in chess, one can say what castling is by articulating all of the constitutive rules for castling: castling is the move that is subject to these rules. Similarly, the thought goes, to give an account of assertion, it is sufficient to articulate its constitutive rules.

(117) does not deny that there are other norms governing assertion—for example, norms of politeness, evidence, prudence, and relevance—or even that these norms can sometimes override the truth rule in one's deliberations about what to assert. But it denies that these norms are constitutive of assertion. They are, rather, derivative: given that assertion is governed by the truth rule, and given other facts about our interests and purposes in engaging in conversation, it will turn out that assertion is governed by these other norms as well. Similarly, other norms involving castling—for example, strategic norms about when you ought to castle if you want to defend against a certain kind of attack—are not constitutive rules, since one could still count as castling without being subject to them.

The truth rule is a semantic-pragmatic bridge principle. It connects a semantic theory—a theory whose output is a definition of truth at a context for arbitrary sentences of a language, and for the propositions they express—with norms for the use of these sentences and propositions. We need not think of either truth or assertion as more fundamental than the other; the bridge principle helps illuminate both. Thus, if we have a definition of truth at a context, the truth rule will help us predict what people will assert in what conditions, and determine what they are trying to get across in using the sentences they use. On the other hand, if we know what they assert, and what sentences they use to do so, the truth rule will allow us to test various definitions of truth at a context.[9]

[8] This is different from saying that a move that does not *obey* this rule would not be castling. A move may be subject to a rule either by obeying it or by being in violation of it. One can castle incorrectly. (If you are tempted to deny this, consider instead the move of *serving* in tennis. Clearly you can serve and violate the rules governing serving, even though being subject to these rules is what makes your movement a serve, and not just a racket-swing.)

[9] Compare our discussion of Lewis in §3.2.1, above.

5.2.1 Defending the truth rule

Recently Timothy Williamson has argued against the truth rule, in favor of the view that assertion is governed by the knowledge rule: "Assert only what you know" (Williamson (1996), Williamson (2000, ch. 11)). But his reasons for preferring the knowledge rule are not compelling.[10]

Williamson's main argument is that the truth rule can't explain why we shouldn't assert of a lottery ticket that it won't win. The assumption here is that, although we don't *know* that the ticket won't win, it is overwhelming likely to be *true* that it won't win. So we have excellent evidence that in asserting that the ticket won't win, we will satisfy the truth rule. Why, then, should it be wrong to assert this? The answer is clear if assertion is governed by the knowledge rule, since the merely statistical evidence we have that the ticket won't win is not sufficient for knowledge.

But the proponent of the truth rule can also answer the question, by invoking two plausible principles, both of which Williamson ought to accept. The first principle is that

(118) One acts reasonably in ϕing just in case one reasonably believes that ϕing is permissible.

Williamson explicitly accepts something that is a bit weaker than this: "One may reasonably do something impermissible because one reasonably but falsely believes it to be permissible" (Williamson, 2000, p. 256). That leaves it open whether one is unreasonable in ϕing when one believes unreasonably, or fails to believe, that ϕing is permissible. But this other direction seems plausible as well. The second principle is the knowledge rule for belief:

(119) One must: believe p only if one knows that p.

Williamson endorses this norm and inclines toward the idea that "to believe p is to treat p as if one knew p" (Williamson, 2000, pp. 46, 255–6). This commits him to the view that one should not even *believe* that a lottery ticket will lose on the basis of merely statistical evidence.

Consider now an agent who accepts the knowledge norm for belief and the truth rule for assertion (or something strictly stronger, like the

[10]For another defense of the truth rule, see Weiner (2005).

knowledge rule). Assume also that there are no other norms that might override the constitutive rule in this case (as considerations of urgency may override the knowledge norm when one yells "That's your train!", Williamson (2000, p. 256)). Then by (118), the agent can reasonably assert that a lottery ticket will lose only if she reasonably believes that the assertion would be permissible. Since there are no other overriding norms, and she takes the truth norm to be the standard for assertions, she reasonably believes that the assertion is permissible only if she reasonably believes that it would satisfy the truth norm. And she reasonably believes that it would satisfy the truth norm only if she reasonably believes that the ticket will lose. But she cannot reasonably believe that the ticket will lose, since she knows that if she did believe this, she would violate the knowledge norm on belief. It follows that she cannot reasonably assert that the ticket will lose.

Thus the proponent of the truth rule *can* explain what is wrong with asserting that a randomly selected lottery ticket will lose. Such an assertion might be permissible, if the ticket turns out to be a losing one, but it would be unreasonable.[11]

Indeed, in combination with the knowledge norm for belief, the truth rule for assertion can explain many of the other phenomena that Williamson takes to support the knowledge norm, such as the fact that assertions can be challenged with "How do you know?", and the fact that it never seems appropriate to assert a proposition of the form

(120) *P*, but I don't know that *P*.

For in asserting the second conjunct, the speaker is conceding that either she does not believe the first conjunct, or she believes it impermissibly. In either case, the assertion is infelicitous.

5.2.2 Relativism and the truth rule

Suppose we start with the truth rule, then, as our basic way of explicating "true (as used) at context *c*." Pursuing our strategy (§5.1), then, let us ask whether this explication of "true as used at *c*" can be

[11] Williamson might urge that the knowledge rule still has some advantages, insofar as it can explain what is wrong with lottery assertions even when the asserter does not accept a knowledge norm for belief. But is it so clear, in such cases, that the problem lies with the assertion, and not with the belief it expresses?

expanded in a natural way to an explication of "true as used at c_1 and assessed from c_2." How can we restate the truth rule in (117) using a truth predicate that is relativized to both contexts of use and contexts of assessment? It seems that there are three basic options for dealing with the extra context of assessment parameter.

First, we could *relativize the norm itself to contexts of assessment:*

(121) Relative to context c_2, an agent is permitted to assert that p at c_1 only if p is true as used at c_1 and assessed from c_2.

(Here we have stated the rule as a deontic principle rather than an imperative; this will make things easier later.) On this view, there is no "absolute" answer to the question "what is the norm governing assertion?", but only a perspective-relative answer.

Second, we could *quantify over contexts of assessment:*

(122) An agent is permitted to assert that p at context c_1 only if p is true as used at c_1 and assessed from

a. *some* contexts.

b. *all* contexts.

c. *most* contexts.

All three variants give us "absolute" norms for assessing assertions.

Finally, we could *privilege one context of assessment.* The only natural choice is the context occupied by the asserter in making the assertion:[12]

(123) An agent is permitted to assert that p at context c_1 only if p is true as used at c_1 and assessed from c_1.

Since it is a metaphysical truth that an agent must be in a context in order to assert anything at it, (123) has the same normative consequences as

(124) An agent at c_2 is permitted to assert that p at context c_1 only if p is true as used at c_1 and assessed from c_2.

[12] In some cases, it may also make sense to privilege another context that the speaker has in mind, but quite often there will not be a unique such context, since the speaker will intend her assertion to be assessable from many different contexts. In any case, the objections in what follows to (123) will apply equally to a proposal to fix the relevant context of assessment as the one the speaker has in mind.

Since (123) and (124) are normatively equivalent, we will switch between them freely. In this section, we will use the first, simpler, formulation, though the second formulation will later prove useful in deriving retraction norms.

It turns out that none of these options gives us what we are looking for: a practical grip on the doubly relativized predicate "true as used at c_1 and assessed from c_2."

Option (121) just explains one mysterious relativization in terms of another. What is it for assertion to be governed by one constitutive rule from one context of assessment, and by another from another? We can readily make sense of game rules whose contents make reference to context—for example, "if you're on a corner square, do this; if not, do that"—but what is envisioned here is that it is a context-relative matter what the rule *is*. It is not helpful to be reminded that the rules for football are different in the US and in Australia; clearly, there are two different games here, "American football" and "Australian rules football." Any particular game is going to subject to the rules of one or the other. To get an analogy with (121), we'd need to imagine a pass in a single televised game that was legal as assessed from America but not as assessed from Australia. We could make sense of such a thing if we could understand what it was for the claim *that that particular pass was legal* to be true as assessed from one context, but not as assessed from another. But that is just what we are hoping that (121) would help illuminate. (121) presupposes, rather than providing, an understanding of assessment-relative truth.

Options (122a–122c) are at least intelligible, but they will not serve the relativist's purposes. It is too easy to assert something that is true at *some* context of assessment, and if we require truth at *every* context of assessment, the resulting norm will forbid asserting anything assessment-sensitive. *Most* seems the best choice of quantifier for the relativist, but there is something arbitrary about it; majority rule looks misplaced here. Nor is it clear what "most" means in this context, if, as seems likely, there are infinitely many possible contexts of assessment.[13]

[13]This proposal would also face a version of the problem for (123) discussed in the rest of this section.

What about option (123)? It makes sense to privilege the context the asserter occupies when she makes the assertion as the one relative to which she should assert only truths.[14] But this option will not help us make sense of relative truth, for it leaves contexts of assessment without any *essential* role to play. Given (123) and a definition of truth at a context of use and context of assessment, we can always construct a nonrelativist theory that does not posit any assessment sensitivity but has exactly the same consequences for the correctness of assertions. Unless we can say something more about the significance of contexts of assessment, then, they will be an idle wheel.

The point is most easily seen with a concrete example. Suppose we are working with contents that have truth values relative to worlds and tastes. And suppose we accept (123) as our basic account of how truth at a context of use and context of assessment relates to proprieties for language use. Let us now compare two theories. The first theory, T_1, defines truth at a context of use and context of assessment as follows (cf. §4.6):

(125) A proposition p is true at as used at c_1 and assessed from c_2 iff p is true at $\langle w_{c_1}, g_{c_2} \rangle$, where w_{c_1} is the world of c_1 and g_{c_2} is the taste of the agent of c_2.[15]

According to this theory, "tasty" is assessment-sensitive. The second theory, T_2, defines truth at a context of use and context of assessment as follows (cf. §4.5):

(126) A proposition p is true at as used at c_1 and assessed from c_2 iff p is true at $\langle w_{c_1}, g_{c_1} \rangle$, where w_{c_1} is the world of c_1 and g_{c_1} is the taste of the agent of c_1.

According to this theory, "tasty" is use-sensitive, but not assessment-sensitive.

We would like to see some difference in practice between the relativist theory T_1 and the nonindexical contextualist theory T_2. But, given (123), the two theories have exactly the same consequences for normative assessment of assertions. They both predict that agents should assert that a food is tasty only when that food tastes good to

[14] See Kölbel (2002, p. 125), Egan, Hawthorne, and Weatherson (2005, p. 153).

[15] This amalgamates (110) and (108).

them. This is a problem for the relativist. The problem is not that the prediction is implausible, or one the relativist should reject. The problem is that, if (123) is our sole point of connection between the semantic theory and facts (in this case, normative facts) about the use of language, then the relativist has not explained the practical difference between a relativist theory and a nonrelativist one.

Here is another way of seeing the point. Suppose there are three possible contexts: c_1, c_2, and c_3. The contexts all have the same agent but take place at different times (t_1, t_2, t_3). The agent likes licorice at t_1 and t_2, but not at t_3. Let p be the proposition that licorice is tasty.[16] We can compare T_1 and T_2 by looking at the truth values they assign to p at each possible combination of a context of use and context of assessment (see Figures 5.1 and 5.2).

Figure 5.1: T_1 (relativist)

		assessed		
		c_1	c_2	c_3
used	c_1	*T*	*T*	*F*
	c_2	*T*	*T*	*F*
	c_3	*T*	*T*	*F*

Figure 5.2: T_2 (contextualist)

		assessed		
		c_1	c_2	c_3
used	c_1	*T*	*T*	*T*
	c_2	*T*	*T*	*T*
	c_3	*F*	*F*	*F*

Note that the only cells of the table that matter to the propriety of assertions, according to (123), are the shaded cells on the diagonal (where the context of assessment is the same as the context of use). Since T_1 and T_2 agree on these cells, (123) does not help us to distinguish between them. They are "normatively equivalent" theories.

Indeed, the figure shows how, given *any* relativist theory T_r, we can construct a "normatively equivalent" contextualist theory T_c: we simply need to ensure that the truth value T_c assigns to a sentence or proposition at context of use c_1 and any context of assessment c_2 is the truth value that T_r assigns it at context of use c_1 and context of assessment c_1. The upshot is that (123) does not give us a basis for "making sense of relative truth." The antirelativist can say to the relativist:

[16] If you like, you can add "throughout the period $t_1 \ldots t_3$"; we will assume that the taste of licorice does not change during this period.

> What you call "truth as used at and assessed from c," and identify with the norm of assertion, is what I call "truth as used at c." At any rate, they are identical in their normative and empirical import. But you have done nothing to explain what "truth as used at c_1 and assessed from c_2" means, when $c_1 \neq c_2$. If you had, we would be able to see a difference in the consequences for language use between a relativist theory and a nonindexical contextualist theory that coincides with it "on the diagonal," as T_2 coincides with T_1.

It might be protested that even if the difference between T_1 and T_2 does not manifest itself as a difference in the norms for asserting p, it manifests itself as a difference in the norms for asserting that particular assertions of p are "true." One might expect that the relativist and the nonindexical contextualist theories would disagree at least about this. It turns out, though, that they do not. Recall the natural semantics for "true" given in §4.7, above:

(115) The extension of "true" at a circumstance of evaluation e is the set of propositions that are true at e.

Suppose that Jake asserts p (the proposition that licorice is tasty) at c_1, and we are assessing his assertion from c_3. As we have already seen, T_1 and T_2 disagree about whether p is true as used at c_1 and assessed from c_3. But they do not disagree about whether the proposition expressed at c_3 by

(127) What Jake said at t_1 is true.[17]

—call it $T(p)$—is true as used at and assessed from c_3. For, on both accounts, (127) will be true as used at and assessed from c_3 just in case $T(p)$ is true at $\langle w_{c_3}, g_{c_3}\rangle$, where w_{c_3} is the world of c_3 and g_{c_3} the taste of the agent at c_3. And, given (115), $T(p)$ will be true at $\langle w_{c_3}, g_{c_3}\rangle$ just in case p is true at $\langle w_{c_3}, g_{c_3}\rangle$. Since T_1 and T_2 agree that p is false at $\langle w_{c_3}, g_{c_3}\rangle$, they will agree that $T(p)$ is false as used at and assessed from c_3. And, given (123), they will agree that what Jake said cannot be correctly said at c_3 to be "true."

[17] Or "was true." Since we are operating with eternalist propositions, whose truth values do not vary with time, there is no significant difference.

Granted, the two theories will make different predictions about whether an assessor at c_3 could correctly call Jake's *utterance* (in the "act" sense) "true." But, as noted in §3.1.2, the monadic predicate "true" in ordinary use is a predicate of propositions, not utterances. Perhaps we can understand utterance truth as a technical notion, by saying that an utterance at c_1 is true (as assessed from c) just in case the sentence uttered is true as used at c_1 and assessed from c_2. But precisely because utterance truth is a technical semantic notion, we should not expect to be able to adjudicate between two theories (T_1 and T_2) by looking at their predictions about the use of sentences that characterize utterances as true. Theorists who accept T_1 will apply "true" to utterances in one way; those who accept T_2 will apply it in another way; and ordinary speakers will not apply "true" to utterances at all. And what happens if the language we are studying does not contain "true" as a predicate of utterances? Do we then lose our grip on the significance of assessment-relative truth assignments?

We must conclude, then, that if we follow Dummett and Lewis in taking the truth rule to be the fundamental semantic-pragmatic bridge principle connecting a truth-conditional semantic theory with proprieties for the use of language, then we cannot make sense of assessment-relative truth. (Or, what comes to the same thing: we cannot see any practical difference between semantic theories that posit assessment sensitivity and those that do not.) Parallel considerations will rule out explaining the significance of relative truth by talking of truth as the norm of belief, rather than assertion. Some philosophers have concluded on this basis that relative truth talk is incoherent.[18]

5.3 Retraction

I want to suggest a less bleak diagnosis. The basic thought is that the pragmatic difference between T_1 and T_2 manifests itself in norms for the *retraction* of assertions rather than norms for the *making* of assertions. T_1 predicts that an assertion of p at c_1 ought to be retracted by the asserter in c_3, while T_2 predicts that it need not be retracted. Thus, (123) is not so much wrong as incomplete. It needs to be supplemented by a constitutive norm for retraction:

[18] In addition to Evans (1985), see the nuanced discussions in Percival (1994) and Campbell (1997, pp. 165–6).

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

Together, (123) and (128) allow us to see the practical difference between semantic theories that posit assessment sensitivity (like T_1) and theories that do not (like T_2).[19]

By "retraction," I mean the speech act one performs in saying "I take that back" or "I retract that." The target of a retraction is another speech act, which may be an assertion, a question, a command, an offer, or a speech act of another kind.[20] For this reason, there are always *two* contexts relevant to retractions: the context in which the retraction itself takes place and the context in which the original speech act took place. It is by exploiting this fact that (128) gives a normative role to contexts of assessment. It requires retraction when the proposition asserted is untrue as used at the context of the original assertion and assessed from the context in which the retraction is being considered.

The effect of retracting a speech act is to "undo" the normative changes effected by the original speech act. So, for example, in retracting a question, one releases the audience from an obligation to answer it, and in retracting an offer, one withdraws a permission that one has extended. Similarly, in retracting an assertion, one disavows the assertoric commitment undertaken in the original assertion. This means, among other things, that one is no longer obliged to respond to challenges to the assertion (since one has already conceded, in effect), and that others are no longer entitled to rely on one's authority for the accuracy of this assertion. (One can, of course, still be held morally accountable if others relied on one's assertion before they knew that it was retracted.)

Note that (128) uses an "if" rather than an "only if," as in (123). In other words, (128) *obliges* an act under certain conditions, while (123) *forbids* an act under certain conditions. This is as it should be. One

[19] A different, but compatible approach would be to describe a norm for challenging others' assertions, as in MacFarlane (2007, §5.2).

[20] Explicit retractions of assertions are relatively rare, because it is usually assumed that in acknowledging the inaccuracy of the original assertion, one implicitly retracts it. But this presumption can be defeated: I might say, for example, "I know that what I said was almost certainly false, but I'm standing by it and not retracting."

is not forbidden to retract an assertion that is true. In many cases one might have good reason to retract an assertion one still thinks is true—for example, one may not want others to rely on one's word in this matter, or one may not want to take on the obligation of defending the assertion—and doing so is not "insincere" in the way that asserting something one does not believe to be true is.[21]

Since retraction is supposed to be a kind of "undoing" of the speech act it retracts, the norm for retracting a speech act ought to be derivable from the norm for performing it. Taking account of the point made in the previous paragraph, the obvious relationship is this:

(129) Retraction of a speech act A is *required* if the original act A is *forbidden* by its constitutive norm.

Disambiguating this by specifying *for whom* each act is required or forbidden, we get:

(130) Retraction of a speech act A is *required* of an agent in context c if the original act A is *forbidden* for an agent in context c by its constitutive norm.

Applying (130) to (124), we get (128). So our modified truth norm for making assertions and our norm for retracting assertions fit each other like hand and glove.

This combination of norms allows that someone who asserts that p in c_1 might be compelled to retract this assertion in a later context c_2, even though the assertion was permissible for her to make at c_1. (This can happen if p is true as used at and assessed from c_1, but not true as used at c_1 and assessed from c_2.) This may seem odd. Percival (1994, p. 209) asks, "How can I believe both that the aims given A, for him, by the language he employs were successfully pursued, and that I have every right to force him to withdraw his utterance?"

Here it is important to keep in mind that withdrawing an assertion (or other speech act) is not tantamount to conceding that one was at

[21] This point is missed by one of the few explicit accounts of retractions in the literature on speech act theory. Bach and Harnish (1979, p. 43) say that "In uttering e, S retracts the claim that P if S expresses: (i.) that he no longer believes that P, contrary to what he previously indicated he believed, and (ii.) the intention that H [the hearer] not believe that P."

fault in making it. Suppose one's evidence all strongly suggests that Uncle Jack is coming to lunch, and on the strength of that evidence you assert that Uncle Jack is coming. A bit later, Aunt Sally calls to say that Uncle Jack has broken his leg. This makes it quite unlikely that he is coming, so you retract your assertion. Nonetheless, you were perfectly reasonable in making it, and cannot be criticized for having done so. Retracting it is not admitting fault.

The case the relativist allows is similar, only the difference between the later context and the earlier one is not (just) a difference in the *evidence* one has for a claim, but a difference in the very *truth* of the claim. Perhaps there is something odd about this, but its oddity cannot consist in the fact that you can be compelled to withdraw an assertion that you had every right to make, since we see that in the epistemic case as well.

5.4 Other accounts of assertion

So far we have been trying to "make sense of relative truth" within the framework of a Dummett/Lewis style account of assertion as normatively aiming at truth. What we have found is that if such an account is complemented with a normative account of retraction, a role can be found for assessment-relative truth. The combined norms allow us to derive normative predictions from semantic theories, and give us a practical handle on the difference between assessment-sensitive and non-assessment-sensitive accounts. An assertion of an assessment-sensitive proposition must be retracted when its content is untrue as assessed from the context in which the retraction is being considered, while an assertion of a non-assessment-sensitive proposition need not be.

But what if we reject the idea that the truth rule is a constitutive norm for assertion? What if we accept a different constitutive norm, such as the knowledge norm (Williamson, 1996, 2000)? What if we reject the idea that assertion is to be understood as a move defined by constitutive norms, and hold instead that assertion is to be understood as the expression of belief (Bach and Harnish (1979, p. 42), B. Williams (2002, pp. 73–5), Owens (2006)), or as a commitment to the truth of the asserted content (Peirce (1934, p. 384), Searle (1979, p. 12), Brandom

(1983), Brandom (1994, ch. 3), Wright (1992), Watson (2004))? Will we still be able to make sense of relative truth?

5.4.1 The knowledge norm

We might state the knowledge norm for assertion as follows:

(131) An agent is permitted to assert that p only if she knows that p.

Is this norm compatible with relativism about truth? It is sometimes argued that the factivity of knowledge precludes relativism. The thought is something like this: if it can be known that p, then p must be a fact, and hence true absolutely. Thus nothing that can be known can be assessment-sensitive. The knowledge norm would then imply that nothing assessment-sensitive can permissibly be asserted. Such an account would certainly not help to illuminate relative truth.

But the line of thought just scouted is faulty. The *factivity of knowledge* is the claim that instances of the following schema are analytically true:

(132) If α knows that ϕ, then ϕ

What we can conclude from this is that if ϕ is assessment-sensitive, so is $\ulcorner \alpha$ knows that $\phi \urcorner$, and so is the predicate "knows" (since its extension varies as we shift the context of assessment). Thus the argument above relies on the tacit premise that "knows" is not assessment-sensitive. But it is question-begging to assume this in an argument that purports to rule out the possibility of assessment sensitivity.[22]

Let us then allow for the possibility that "knows" is assessment-sensitive. If we do that, then (131) needs to be restated in non-assessment-sensitive language. A reasonable candidate, following (124), is

(133) An agent at c_2 is permitted to assert that p (at a context c_1) only if the proposition *that she knows (at* t_{c_1}*) that* p is true as used at c_1 and assessed from c_2.

Because an agent at c can only assert things at c, this norm gives us the same predictions as

[22] Perhaps there are independent arguments that "knows" cannot be assessment-sensitive. I do not know of any. For independent arguments *for* the assessment sensitivity of "knows," see MacFarlane (2005*b*) and chapter **??**, below.

(134) An agent at c is permitted to assert that p (at a context c) only if the proposition *that she knows (at t_c) that p* is true as used at and assessed from c.

By itself, then, it won't help us make sense of relative truth. (The situation is the same as with (124).) But the outlook improves if we use (130) to derive a corresponding retraction norm:

(135) An agent at c_2 is required to retract an assertion that p (made in c_1) if the proposition *that she knows (at t_{c_1}) that p* is not true as used at c_1 and assessed from c_2.

From the agent's standpoint, this amounts to saying that you should retract an earlier assertion that p if, from your present perspective, you did not know that p.

One might quibble with (135). Suppose that at the beginning of a basketball tournament one asserted, without very good evidence, that team X would win. Now it is near the end of the tournament, and it is clear to everyone that X will win. You concede that you did not know this at the time you made the assertion, but are you required to retract it? That seems odd. I suggest, however, that this is a reason to worry about (133) rather than the claim that (135) is the right complementary retraction norm.

Setting that substantive concern aside, however, it should be clear that the pair of norms (133) and (135) does allow us to make sense of relative truth, by giving us a clear practical grip on the distinction between relativist and nonindexical contextualist theories. Suppose that propositions about taste are assessment-sensitive. Let c_1 be a context centered on Joey, with the tastes of a ten-year-old. The proposition that fish sticks are tasty is true as used at and assessed from c_1, and Joey's evidence is sufficient to establish this. So, plausibly, the proposition that Joey knows (at t_{c_1}) that fish sticks are tasty is true as used at and assessed from c_1, and Joey is permitted to assert that fish sticks are tasty. Let us suppose that he does. Let c_2 be a context centered on Joey, ten years later, with the tastes of an educated twenty-year-old. The proposition that fish sticks are tasty is false as used at c_1 and assessed from c_2, so (by factivity) the proposition that Joey knew (at t_{c_1}) that fish sticks are tasty must also be false as used at c_1 and assessed from c_2. Thus (135) says that Joey is now required

to retract his earlier assertion. A nonindexical contextualist account of taste predicates, by contrast, would not predict this, since it would take the proposition that fish sticks are tasty to be true as used at c_1 and assessed from c_2 iff it is true as used at c_1 and assessed from c_1.

5.4.2 Assertion as expression of belief

We have been trying to illuminate relative truth in a framework that understands assertion through its normative relation to truth (or to a truth-involving concept like knowledge). But one prominent strand of thinking about assertion understands it instead through its relation to *belief*. According to this strand, to assert something is to express an attitude—usually belief, but sometimes also an intention that the hearer have a belief.[23] If assertion aims at truth, on such accounts, it is only because truth expresses belief and *belief* aims at truth. Do expressive accounts of assertion give us the means to understand assessment sensitivity? Do they rule it out in principle?

Here the verdict is mixed. Expressive accounts do not rule out relativism about truth; they just kick the problem upstairs. If we can make sense of what it is to have a *belief* with an assessment-sensitive content, then an expressive account will tell us what it is to assert such a content. The problem, as we will see in §5.5, below, is that we cannot distinguish practically between relativist and nonindexical contextualist theories just at the level of belief. So expressive accounts of assertion are bad news for the project of understanding relative truth by seeing what it is to assert assessment-sensitive propositions.

That need only trouble the relativist, though, if expressive accounts are more plausible than their rivals. And they face at least two major problems. The first is that they make it difficult to distinguish between what one has *asserted* and what one has only implied (or, to use Paul Grice's term, "implicated"). If I effusively praise a candidate's handwriting in a letter of recommendation, saying nothing about his suitability for the job, then surely I have expressed my belief that the

[23] B. Williams (2002, p. 74) gives a disjunctive account: "A asserts that p where A utters a sentence S which means that p, in doing which either he expresses his belief that p, or he intends the person addressed to take it that he believes that p." Bach and Harnish (1979, p. 42) say that to assert that P in uttering e to a hearer H is to express "(i) the belief that P, and (ii) the intention that H believe that P."

candidate is not suitable, and my intention to convey this belief to the reader. Still, one might want to say that I have not *asserted* that the candidate is unsuitable, but only implied this. The expressive account of assertion threatens to erase this intuitive distinction.

One might try to reinstate the distinction by saying that assertion is *literal* expression of belief—expression of belief by means of a sentence whose literal content (in context) matches that of the belief expressed (so B. Williams (2002, p. 74)). But that is too restrictive. When Geoffrey Nunberg's waitress says, "The ham sandwich left without paying," she has not asserted that the ham sandwich left without paying (Nunberg, 1979). Nonetheless, she has made an assertion. (How else would you characterize the illocutionary force of her utterance?) Perhaps it is even possible to make assertions using improvised gestures that lack any conventional meaning. Stephen Schiffer has an example in which a husband communicates to his wife that he is bored at a party by wiggling his ears (Schiffer, 1972, p. 126). Perhaps this is not an assertion, but if it isn't one, that's not merely because it lacks a linguistic vehicle. If it turns out that all assertions are linguistic, this ought to be the result of argument, not stipulation about the meaning of "assertion."

The second problem for expressive accounts of assertion is that assertions can be insincere. Indeed, they can be *openly* insincere (as often happens in politics); in such cases the speaker lacks both the belief that is being expressed and the intention that the audience come to have this belief. Thus an expressive account needs to understand expression as something other than the the outward manifestation of an inner state; it must be possible to express attitudes one does not have.[24] Such a conception of expression is not entirely unintuitive: we might naturally say of a con man who duped us by pretending to be lost that he "expressed great consternation," and not just that he pretended to do so. Plausibly, he has expressed consternation because he has acted with the intention of giving us a reason to think him in a state of consternation.

[24]B. Williams (2002, pp. 73–5) and Owens (2006) hold that expressions of beliefs must be caused by the beliefs they express, and concede that only sincere assertions are expressions of belief. They say that insincere assertions are acts of representing oneself as if one is expressing a belief, and therefore count as assertions in a parasitic sense.

Bach and Harnish turn this thought into an account of what it is to express an attitude:

(136) For *S* to *express* an attitude is for *S* to R-intend[25] the hearer to take *S*'s utterance as reason to think that *S* has that attitude. (Bach and Harnish, 1979, p. 15)

But it is not easy to see how, in the case of an *openly* insincere assertion, where it is common ground that the speaker lacks the belief being expressed, the speaker can be intending to give the hearers a reason to attribute the belief. Bach and Harnish point out that a reason can be *pro tanto*, and need not be conclusive:

> *S*'s utterance is, and can be R-intended to be taken to be, *a* reason, despite the fact that it can be overridden by mutual contextual beliefs to the contrary. Even when defeated, a reason is a reason. (Bach and Harnish, 1979, p. 58)

And they offer a reformulation that does not assume that a defeated reason is still a reason:

> Instead of saying that expressing an attitude is R-intending *H* to take one's utterance as reason to believe that one has that attitude, we can say that it is R-intending *H* to take one's utterance as sufficient reason, *unless there is mutually believed reason to the contrary*, to believe that one has that attitude. (Bach and Harnish, 1979, 291, emphasis added)

But this borders on unintelligibility. We can make sense of *intending that Jane take out the trash today, unless it is a holiday*, in a case where it might be a holiday. But when Jane knows that today is a holiday, and the speaker knows that she knows this, what is it for the speaker to intend that Jane take out the trash today, unless it is a holiday? Similarly, if it is mutually known that the speaker's utterance is not a sufficient reason to attribute an attitude to her, what is it for the speaker to intend the hearers to take it as a sufficient reason, unless there is mutually believed reason to the contrary?

5.4.3 Assertion as commitment

Both problems for expressive accounts are handily dealt with by a third kind of account of assertion, according to which "to assert a

[25] To "R-intend" an effect is to intend to bring it about by means of the recognition of this very intention.

proposition is to make oneself responsible for its truth" (Peirce, 1934, p. 384), or to "commit oneself (overtly) to its truth."[26] In Grice's handwriting case, one has implied that the applicant is a poor candidate for the job by committing oneself to the truth of another proposition—that he has nice handwriting. A sign that one has not asserted that the applicant is a poor candidate is that it would be unreasonable to demand that one *retract* this assertion. Retraction isn't appropriate here, because one has not taken responsibility for the claim in the first place. In Nunberg's waitress case, by contrast, one has committed oneself to the truth of the proposition that the person who ordered the ham sandwich left without paying. Here, it would be appropriate to demand retraction if it turned out that the person in question had simply gone to the restroom. Finally, in the case of an openly insincere assertion, one has committed oneself to the truth of the asserted proposition, and can subsequently be held responsible for it, even if one has not expressed the corresponding belief. If someone objects to the assertion, offers evidence, and demands that it be retracted, it is not an adequate defense to say, "But I didn't actually *believe* that."

At first glance such accounts of assertion may seem to leave little room for relativism about truth. If the truth of p is relative to contexts of assessment, how can I commit myself to the truth of p? What, exactly, am I committing myself to? It might seem that the target of my commitment should be something absolute. Recall Burnyeat's charge (discussed in §2.2.1, above) that "[n]o amount of maneuvering with his relativizing qualifiers will extricate Protagoras from the commitment to truth absolute which is bound up with the very act of assertion" (M. F. Burnyeat, 1976*b*, p. 195).

Before we can clarify and address this objection, however, we need to get clearer about what "commitment to the truth of a proposition" amounts to. Commitments are, in the first instance, commitments to *do* something (or to refrain from doing something), and we ought to be able to understand other kinds of commitment talk in terms of commitments-to-do. If a political leader says that she is committed to the fairness of elections, for example, this means that she is committed

[26] Variations on this theme are defended in Searle (1969, p. 29), Searle (1979, p. 12), Brandom (1983), Brandom (1994, ch. 3), Wright (1992), Watson (2004), MacFarlane (2003), and MacFarlane (2005*a*); for criticism, see Pagin (2004).

to doing what it takes to ensure that the elections are fair, and that she can be held responsible if the election turns out to be unfair through her failure to take feasible action. What, then, is a commitment to the truth of a proposition a commitment to do? What would count as honoring such a commitment, and what would count as violating it? Until we can answer these questions, we do not really understand what is meant by "commitment to truth."

Here are three things that might be involved in a commitment (undertaken in a context c_1) to the truth of a proposition p:

(W) Commitment to withdraw the assertion if and when p is shown to be untrue as used at c_1.

(J) Commitment to vindicate the assertion (by providing grounds for the truth of p as used at c_1, or perhaps by deferring to someone else who can) when it is appropriately challenged.

(R) Commitment to be be held responsible if someone else acts or reasons on the basis of the assertion and p proves to be untrue as used at c_1.

I take it that everyone should be able to agree that the commitment one undertakes in making an assertion includes at least (W). Imagine someone saying: "I concede that what I asserted wasn't true, but I stand by what I said anyway." We would have a very difficult time taking such a person seriously as an asserter. If she continued to manifest this kind of indifference to established truth, we would stop regarding the noises coming out of her mouth as assertions. We might continue to regard them as expressions of beliefs and other attitudes (just as we might regard a dog's whining as an expression of a desire for food). We might even find them useful sources of information. But we would not regard them as commitments to truth, and hence not as assertions.

There will be less agreement about (J). Brandom has argued that assertoric commitment includes (J) as well as (W) (Brandom, 1983, 1994), but this may be overgeneralizing from seminar-room assertions to assertions in general. Suppose someone were to say: "You've given some very good reasons to doubt the truth of what I asserted. I have nothing to say in answer to your objections, yet I continue to stand by my claim." She would not be playing the game of assertion the way

philosophers play it, but perhaps philosophers do not get to set the rules here. We would surely take her assertions less seriously than we would if she were responsive to reasons. But would we cease treating her as an asserter at all? That is not so clear.

What about (R)? Asserting is a bit like *giving one's word* that something is so, and our reactions to assertions that turn out to be untrue are a bit like our reactions to broken promises. When someone tells us something that turns out to have been untrue, we feel a legitimate sense of grievance, especially if we have acted on what we were told. Suppose someone tells you that there will be a talk by an interesting celebrity at a nearby university. You cancel some appointments and spend considerable time and energy getting there—but there is no talk. When you confront your informant, you imagine that she will apologize profusely. Even if she has an excuse (perhaps she heard about the talk from someone else she regarded as reliable, or perhaps there was a typo in the schedule), she will accept some measure of responsibility. You don't expect her to say: "You actually *acted* on my assertion? Well, that's not my problem. I can say whatever I like: it's up to you to sort out what's worth taking seriously." But why not? What is so wrong with this kind of response? After all, it *is* up to us whether to believe what we are told, and we criticize people who believe whatever anyone says for their gullibility. A plausible answer (though not the only possible one) is that part of making an assertion is accepting partial responsibility for the accuracy of what one says.

These three putative components of assertoric commitment are largely independent: one could reasonably hold that assertion is constitutively governed by just (W), or by (W) and (J), or by (W) and (R), or by all three. Yet they are arguably connected teleologically. It is *because* proven falsity requires withdrawing one's assertion that one has an obligation to justify it in the face of an appropriate challenge. Providing a justification is a way of guaranteeing (to oneself and others) that the assertion won't have to be withdrawn. And it is *because* one can be held responsible for one's assertions that they must be withdrawn when shown to be untrue: by withdrawing them, one signals that one is no longer to be held responsible. Thus, arguably, (R) provides a rationale for (W), which in turn provides a rationale for (J).

Suppose we understand the "commitment to truth" involved in

an assertion in terms of some combination of (W), (J), and (R). Can we understand what it would be to commit oneself to the truth of an assessment-sensitive proposition? That is, can we construe (W), (J), and (R) in a way that allows that truth might be relative to contexts of assessment?

(W) talks of the asserted proposition being "shown to be untrue as used at c_1." In a framework that allows assessment sensitivity, we must also specify a context of assessment. There are three natural options:

1. The relevant context of assessment is the context in which the proposition was asserted (that is, c_1).
2. Quantify over contexts of assessment: the proposition must be shown to be untrue relative to some/all/most contexts of assessment.
3. The relevant context of assessment is the context to which the asserter *currently* belongs.

It should be plain from our parallel discussion of the aim of assertion that the first two options will not help make sense of assessment sensitivity. They imply that for any given assessment-sensitive proposition, there will be a systematically related assessment-invariant proposition whose assertion results in exactly the same commitments. But unless we can see some difference in practice between asserting an assessment-sensitive proposition and asserting a related assessment-invariant one, we lack a real understanding of assessment sensitivity. Only the third option gives an essential and ineliminable role to contexts of assessment.

I conclude that the relativist should construe (W) along the lines of the third option, which privileges contexts the asserter occupies, while still allowing the relevant context of assessment to diverge from the context of use:

(W*) Commitment to withdrawing the assertion (in any future context c_2) if p is shown to be untrue as used at c_1 and assessed from c_2.

There should be no worries about the intelligibility of (W*). Logically, it is no more complex than a commitment to refill the pitcher (at any future time t) if and when it is shown to be empty (at t). And it agrees with (W) in the special case where p is not assessment-sensitive.

If (W) rationalizes (J), we must generalize them both in the same way. (W*) requires that an assertion be withdrawn when its content is proven untrue relative to the asserter's current context of assessment, so the justification demanded by (J*) must consist in grounds for the truth of the asserted proposition relative to the asserter's current context of assessment.

(J*) Commitment to justifying the assertion (that is, providing grounds for the truth of p as used at c_1 and assessed from c_2) if and when the assertion is appropriately challenged at c_2.

Finally, if (R) rationalizes (W), we must generalize (R) as follows:

(R*) Commitment to accepting responsibility (at any future context c_2) if on the basis of this assertion someone else takes p to be true (as used at c_1 and assessed from c_2) and it proves not to be.

An account of assertoric commitment based on some combination of (W*), (J*), and (R*) will help us make sense of relative truth by showing us what, exactly, one is committing oneself to do in asserting an assessment-sensitive content, and how that commitment differs from commitments one could undertake by asserting assessment-invariant contents. For example, suppose we are trying to choose between a nonindexical contextualist account of "tasty" and a relativist one. According to the nonindexical contextualist account, the truth of the proposition that fish sticks is tasty, as used at c_1 and assessed from c_2, depends on the tastes of the agent at c_1; according to the relativist account, by contrast, it depends on the tastes of the agent at c_2. So according to the nonindexical contextualist account, someone who asserts that fish sticks are tasty is undertaking a commitment to withdrawing this assertion if fish sticks are shown to be tasty by the lights of her tastes at the time she made the assertion, while according to the relativist account, she is undertaking a commitment to withdrawing this assertion if fish sticks are shown to be tasty by the lights of her *current* tastes (at the time she is considering whether she should retract the assertion). Since these are clearly different commitments, (W*) gives us a grip on the practical difference between a nonindexical contextualist and a relativist account, and tells us what to look for in choosing between them.

Pace Burnyeat and Passmore, then, we can make good sense of the idea of commitment to truth even if truth is relative.

5.5 Believing relative truths

We have sought to understand what it is for truth to be assessment-relative by understanding the practical difference between *asserting* assessment-sensitive and assessment-invariant contents. One might wonder: why this focus on assertion? Why not make sense of relative truth by trying to understand what it is to *believe* an assessment-sensitive content? More generally, why should we focus on the significance of assessment sensitivity for speech acts rather than for mental attitudes?

The problem is not that we cannot make sense of beliefs with assessment-sensitive contents. Once we understand what it is to assert an assessment-sensitive proposition, there is no obstacle to countenancing beliefs with these propositions as their contents. The problem is, rather, that we cannot make sense of assessment sensitivity *by* understanding what it is to believe assessment-sensitive contents. To put it starkly: for creatures that were only believers, and did not also make assertions, we could not discern any practical difference between an assessment-sensitive and a non-assessment-sensitive semantic theory.

The reason is simple: there is nothing corresponding to the *retraction* of a belief. Recall that we could make sense of the distinction between an assessment-sensitive and an assessment-invariant theory that agreed on the intensions of propositions[27] only by considering norms for retraction (or commitments to retract). Retraction was the key to making sense of assessment sensitivity because in retraction there are always two significant contexts: the context in which the retraction is being considered and the context in which the assertion whose retraction is being contemplated was made. This gives both the context of assessment and the context of use a job to do in a norm for retraction.

Why is there nothing like retraction in the case of belief? Can't one give up a belief, just as one can retract an assertion? Here it is important to keep in mind a metaphysical difference between assertions and beliefs. An assertion is an *action*, and hence also an *event* that takes place at a given place and time; a belief is not an action but a *state* that an agent can be in over a period of time. The

[27] That is, their truth values relative to circumstances of evaluation.

inception of a belief may be an event (though usually not an action), but the belief itself is not. When one gives up a belief that p, one transitions from being in the state of believing that p to being in the state of not believing that p, but this transition is not directed towards any particular past event. Retraction, by contrast, is always retraction *of* some particular dated act.

Here is another way to see the point. Suppose that at t_0 we query Jim about p, and he asserts that p.[28] Then, at some later time t_1, we query him again, and—perhaps because he has come to doubt the grounds he had before—he refuses to assert that p. This refusal to assert that p at t_1 is not itself a retraction of his earlier assertion that p at t_0. Jim might *also* retract the earlier assertion, signaling that he is no longer committed to the truth of p, but if he does, that is a distinct act. It is an act he *ought* to perform if he wants to be coherent, but he might fail to perform it. We can distinguish, then, between no longer being willing to assert that p and retracting an earlier assertion that p. In the case of belief, though, we can't make a comparable distinction. We can certainly imagine Jim believing that p at t_0 and then, in response to new evidence, ceasing believing that p at t_1. But there is no further backwards-directed act he needs to perform in order to be coherent. Ceasing believing is all he needs to do; he need not somehow "undo" his earlier belief.

Consider this analogy. Two men, A and B, are walking around an Italian garden. A carries with him a supply of stakes with flags on them. He stops periodically, drives a stake into the ground, and writes on the flag, "Viola is the loveliest woman on earth." B is also an admirer of Viola, but instead of planting flags in the ground, he simply *carries* his flag reading "Viola is the loveliest woman on earth." Halfway through the garden, A and B both spot Cynthia and are immediately smitten. A begins to plant flags reading "Cynthia is the loveliest woman on earth," and B repaints his flag to read the same. At this point, A must go back and pick up all the flags declaring Viola to be the loveliest woman on earth. B faces no comparable task; it is enough for him to simply change his flag.

[28] Assume, to avoid complications, that p is a normal eternalist proposition, with truth values relative to possible worlds.

It is in the norms for retraction that we find an independent role for the notion of a context of assessment, so if there is nothing like retraction for beliefs, then it is not clear how we could distinguish between relativist and nonindexical contextualist theories just by looking at their predictions about when people should believe the propositions in question. The problem is not, as is sometimes supposed, that the relativist can't make sense of belief as "aiming at truth," or of norms for belief. To say that belief aims at truth is to say that a belief in context c succeeds in its aim if its propositional content is true as used at and assessed from c.[29] To say that truth is a norm for belief is to say that at a context c one ought to believe only propositions that are true as used at and assessed from c (cf. Kölbel (2002, pp. 32, 91)).[30] The problem, rather, is exactly the same as the problem about assertion we discussed in §5.2.2: saying these things isn't enough to distinguish relativist views from nonindexical contextualist variants of them that generate exactly the same normative predictions. The solution there was to bring in norms for retraction, but that solution is not available for belief.

I think this is an interesting and surprising result. What makes relative truth intelligible is the potential difference between the context at which an assertion is made and the contexts at which challenges to it will have to be met and retractions considered. Thus, even though assessment-sensitive propositions can be believed, judged, doubted, supposed, and so on, there would be no theoretical need for relative truth if we did not also make assertions.

In §2.2, we concluded that the solid core of the self-refutation objection was a challenge for the relativist. The relativist cannot understand

[29] Of course, beliefs don't literally "aim at" anything. For attempts to unpack the metaphor, see Velleman (2000) and Wedgwood (2002).

[30] Zimmerman (2007, p. 337) argues that a rational agent seeking to believe the truth will not clearheadedly "believe a proposition that is both relatively true and relatively false." But surely it is perfectly rational for an agent to believe, for example, that Dodos are extinct, even though this proposition is true at some circumstances of evaluation (worlds) and false at others. Similarly, temporalists will hold that it is rational to believe that it is raining, even though this (tensed) proposition is true at some times and false at others. What matters is whether the proposition at issue is true relative to the context the believer currently occupies.

asserting that p as putting p forward as absolutely true. But what is it to put p forward true, but only relatively so?

In this chapter we have tried to meet this challenge head on, by showing how the assertion of assessment-sensitive contents can be rendered intelligible in the context of several different kinds of accounts of assertion. Given one of these accounts, we can say precisely what the difference in practice is between asserting an assessment-sensitive content and asserting an assessment-invariant one. The accounts are conservative, in the sense that they give orthodox predictions about assertions of assessment-invariant contents. So they provide a framework within which we can consider relativist theories and compare them with non-relativist ones. The framework itself is neutral, but it tells us what to look for in arguing for or against a relativist theory.

It is time to put to rest the common but unsupported view that relativism about truth is self-refuting or incoherent, and ask instead whether it is supported by the (broadly linguistic) evidence. First, however, we need to get clearer about the thorny issue of disagreement, in light of our new understanding of assessment-sensitive truth. That is the task of the next chapter.

Draft: December 13, 2009

6 Disagreement

The achilles heel of contextualism is the problem of *lost disagreement*. If in saying "That's tasty" Sal is asserting that the food tastes good to him, and in saying "That's not tasty" Sam is asserting that it doesn't taste good to *him*, then their claims are compatible and it is mysterious why they should regard themselves as disagreeing. Sophisticated contextualists attempt to regain the lost disagreement by taking "tasty" to express the property of tasting good to a contextually relevant *group*, or to a suitably idealized version of the speaker. As we saw in § 1.2, such moves face a dilemma. If the group is kept small and surveyable, and the idealization mild, then it is always possible to find cases of apparent disagreement it will not explain. But if we expand the group (or idealization) far enough to capture all the apparent disagreement, we can no longer understand how speakers could regard themselves as suitably placed to make the relevant assertions in the first place. A primary selling point of relativist views against contextualist ones is that they purport to capture the subjectivity of claims of taste without losing the disagreement.

But this claim needs more scrutiny. Some critics have charged that relativists about truth are unable to account for disagreement at all. Thus Frege, who favors a contextualist account of taste predicates,[1] 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

[1] "As regards a sentence containing a judgement of taste like, 'This rose is beautiful', the identity of the speaker is essential to the sense, even though the word 'I' does not occur in it" (Frege, 1979, p. 235).

> 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, p. 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.

However, it is easy to ask the wrong question. 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 may be 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, our strategy will be 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.

6.1 CLARIFYING THE TARGET

Cappelen and Hawthorne (2009, pp. 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. 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.

Here we will be primarily concerned to illuminate the state sense. It seems plausible that any account of the activity sense will make reference to a state sense: having a disagreement requires taking oneself to disagree. If this is right, then the state sense is more fundamental.

What is the logical form of the relation we seek to explicate? We could take as our target the relation

(137) x disagrees with y.

But this concept 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 some respects, while disagreeing in others. So we might take our target to be the relation

(138) 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

(139) x disagrees with y's ϕing 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. Given this relation between a person and an attitude or speech act in context, we can presumably define 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. So we do not lose anything by focusing on this admittedly somewhat artificial relation.

6.2 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.[2] In other words, I disagree with attitudes that are not *cotenable* with my current attitudes.[3]

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

[2]

[3]This notion can be extended from attitudes to claims: I disagree with someone's claim if I could not coherently make the same claim—a claim with the same content and force—without changing my mind or retracting one or more of my own claims. This extension might be useful in cases where the parties are playing "devil's advocate" or in some other way speaking against their own beliefs. If Lawyer *A* says "my client is innocent" and Lawyer *B* says, "no, he is guilty," then they have made noncotenable claims; no one person could make both claims without incoherence. But they may not have noncotenable beliefs, since both may believe that the client is guilty. Here there is a disagreement in claims, but not in beliefs.

> THE SIMPLE VIEW OF DISAGREEMENT To disagree with someone's belief that p is to have beliefs whose contents are jointly incompatible with p.[4]

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.

Or consider the following (Huvenes, 2008):

(140) Bob: The hypothesis is false.
Carol: I disagree, we need to do further testing.

Here Carol seems to be disagreeing with Bob, even though what Bob has said is not incompatible with anything she believes. (She may think Bob's claim is more likely than not to be true.) We can understand the disagreement in terms of non-cotenability. In asserting that the hypothesis is false, Bob has expressed a high degree of confidence that it is false. This confidence is not cotenable with Carol's attitudes, which warrant a lower degree of confidence pending further tests.

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

[4]Of course, if contents are individuated coarsely—for example, as sets of possible worlds or Russellian propositions—more must be said. We might not want to say that Hammurabi disagrees with Sammurabi's belief that Hesperus is visible in virtue of believing that Phosphorus is not visible. One solution is to adopt a conception of contents, and of compatibility, on which *Hesperus is visible* is compatible with *Phosphorus is not visible*. Another is to require not just that the beliefs be incompatible, but that it be possible to come to know that they are so without further empirical investigation.

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 is enough if they just have different attitudes towards licorice.

So, non-cotenability is a 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.[5]

6.3 PRECLUSION OF JOINT SATISFACTION

Does Jane and Sarah's difference in attitude Bob really amount to a disagreement? A difference, certainly, but a disagreement? Well, it does seem natural to say that they disagree in their attitude towards Bob. But perhaps that is rather thin. A disagreement, we might think, is a kind of conflict or dispute. To disagree with someone is not just to have a different attitude, but to be in a state of tension that can only be resolved by one or both parties *changing* their minds. Mere practical noncotenability does not always give us that. If Jane loves Bob more than anyone in the world, and Bob loves Jane more than anyone in the world, then their attitudes are not practically cotenable, but far from disagreeing, they seem to be in a happy state of concord.

We might, then, want to think about disagreement in attitude in a somewhat different way, following C. L. Stevenson:

> This occurs when Mr. A has a favorable attitude to something, when Mr. B has an unfavorable or less favorable attitude to it,

[5]Kölbel (2004*b*, p. 305), defending a kind of truth relativism, 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 doxastic non-cotenability. We will see 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.

> and *when neither is content to let the other's attitude remain unchanged*. (C. L. Stevenson, 1963, 1, emphasis added)

This won't quite do if we're trying to explicate the "state" sense of disagreement, which is not supposed to depend on the parties' attitudes towards each other. But Stevenson later recharacterizes "disagreement in attitude" in terms that are more suitable for our purposes:

> 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)

I disagree with someone's attitude, on this account, if its satisfaction precludes satisfaction of my own. Call this sense of disagreement *preclusion of joint satisfaction*.

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-cotenability and preclusion of joint satisfaction would become purely notional, at least in the case of desirings-to-do. (It is far from clear how the strategy could be extended to attitudes like *loving Jane more than anyone in the world*.)

This is not the place to settle a dispute about the contents of desires. But the controversy here provides no reason to resist distinguishing between practical noncotenability and preclusion of joint satisfaction. Even if it turns out that the two notions are necessarily equivalent, so that the distinction between them is merely notional, that wouldn't show that the distinction is pointless.

6.4 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 $\langle w, t, i \rangle$ 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. (Lewis (1979*a*), who originated this approach, talks 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 its context) is eating a sandwich (at the time of its context), 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*. 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.[6]

We have, then, another 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. Though it is initially tempting to give a modal analysis of "the accuracy of A precludes the accuracy of B"—perhaps as "it is impossible for B to be accurate if A is"—this will not work. For if it is impossible for B to be accurate if A is because it is impossible for B to be accurate, quite independently of any relation to A, then it would be wrong to say that A *precludes* the accuracy of B. Although it is difficult to say what preclusion amounts to in other terms, 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).

[6]This may seem to contradict the view, propounded by some advocates of truth relativism, that disputes of taste involve "faultless disagreement." Whether it does depends on how we disambiguate "faultless" and "disagreement." For further discussion, see §6.8, below.

I have used "accuracy" in an informal way, but one can say precisely how it is related to the various relativized notions of truth we have looked at in chapters 3–5, above. 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). More precisely: to say that an attitude or speech act is accurate is to say that is true relative to its context. Or, if we are working in a framework with assessment-relative truth: to say that an attitude or speech act is accurate, as assessed from a context c, is to say that it is true as used at the context in which it occurs and assessed from c.

The distinction between truth and accuracy doesn't matter much when we're considering whether to assert or believe something ourselves. 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 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, non-tensed 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—as a premise in an argument against centered propositions and other nonstandard contents. This is essentially what Cappelen and Hawthorne do in their recent book (Cappelen and Hawthorne, 2009, pp. 96–8). They argue, more or less, as follows:

(141) Two parties agree if they both accept the same proposition.

(142) Suppose that tensed propositions can be the contents of beliefs.

(143) Then it should follow that if Bill believed, two days ago, the tensed proposition *It is raining in Boston*, and Janet believed the same tensed proposition two weeks ago, they agreed.

(144) But this pattern of attitudes does not constitute agreement.

(145) So, by reductio, tensed propositions cannot be the contents of beliefs.

Once we have distinguished between disagreement as doxastic noncotenability and disagreement as preclusion of joint accuracy—and between the corresponding senses of agreement—we can see that this argument has no force. If tensed propositions *can* be the contents of beliefs, then premise (141) is only true if "agree" is taken in the first sense, while (144) is only true if "agree" is taken in the second sense. So Cappelen and Hawthorne can only resist the charge of equivocation by assuming at the beginning of the argument what they seek to prove—that belief contents cannot be tensed. And that is begging the question.

Arguably, a need for the distinction between doxastic noncotenability and preclusion of joint accuracy 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. Borrowing

some terminology from Perry (1986), we might say that although neither belief is *about* any particular world, Jane's belief *concerns* our world, while June's concerns hers, and both beliefs are accurate.[7] In at least one important sense of "disagree," 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 MacFarlane (2007), and Cappelen and Hawthorne (2009, ch. 2, § 17) rightly called me on it.) 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.

6.5 PRECLUSION OF JOINT REFLEXIVE ACCURACY

There is still a further distinction to be made. In order to motivate it, though, we will need to review the difference between objectivist, non-indexical 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 Hen-of-the-Woods is tasty*—not *tasty to Sal*, or *tasty to most people*, but just *tasty*. But they differ in what they say about the intension of this

[7] For this technical use of "concerns," see §4.2, above.

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-in-context is accurate if its content is true relative to the world of the context and the taste relevant at that context (normally, the believer's own taste). 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 of the context of use and the taste relevant at the context of assessment (normally, the assessor's taste). Its accuracy (at a context of assessment) does not depend at all on the believer's tastes (unless the believer is also the assessor).

Note that relative to any one context of asesssment, at most one of two incompatible taste propositions will be accurate. So, for the relativist, as for the objectivist, doxastic non-cotenability and preclusion of joint accuracy go together.

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 asesssed 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.

6.6 DISAGREEMENT IN DISPUTES OF TASTE

Enough distinguishing! Recall our 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 Hen-of-the-Woods that Sal cannot coherently take on board himself without changing his own attitudes towards Hen-of-the-Woods. 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 Hen-of-the-Woods, making it congruent with his own, and to do this he may try to call Sam's attention to various salient facts about the Hen-of-the-Woods. 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 Hen-of-the-Woods. 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 mushroom, respectively. Indeed, even if Sal had said "I like this" and Sam had said "Well, I hate it," they could be said to disagree.

However, some of the ways in which Sam might naturally express his disagreement with Sal 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:

(146) Sal: I like this.
Sam: No, I don't like it.

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

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

This is hard to explain unless Sam takes himself to disagree with what Sal has asserted, or with a belief Sal 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, explored by de Sa (2008), 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. If Sam is presupposing that Sal's tastes are like his, then the belief expressed by Sal's claim, on the contextualist account—Sal's

belief that Hen-of-the-Woods tastes good to him—*is not* doxastically co-tenable with Sam's attitudes: Sam could take it on board only by rejecting his belief that Sal'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 foods tend to be very different. The dialogue with which we 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. On such a view, an assertion that p made at context c counts as accurate just in case p is true at the world of c and the taste of the agent at c.[8] This approach would retain the key contextualist idea that the accuracy of Sam's belief about the tastiness of Hen-of-the-Woods 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. Even though Sam could acknowledge that Sal's belief is accurate, he could not regard its content as true,[9] and he could not come to believe what Sal believes without giving up a belief of his own. This would be enough to vindicate responses like those in (147), above.

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 things tip in favor of relativism if the parties to such disagreements think of themselves not just as trying to change the other party's attitudes, but as trying to *refute* them—where the sign of successful refutation is not just that the other party now

[8] Compare the account of "beautiful" discussed in §4.5, above.

[9] Cf. MacFarlane (2009, § 7).

holds the content of her original claim to be false, but that she retracts her original assertion as inaccurate.

Disputes of taste do seem to have this flavor. If Sam eventually gets Sal to dislike the taste of Hen-of-the-Woods, Sal will feel pressure to withdraw his earlier assertion that the it is 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 liking for Hen-of-the-Woods, 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 his 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.

6.7 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 (see Figure 6.1[10]).

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

Figure 6.1: Semantic views and types of disagreement.

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.

[10]The chart assumes 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.

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

6.8 Appendix: On "faultless disagreement"

Some recent advocates of relativism about truth—most prominently Max Kölbel—have argued that disputes of taste are characterized by "faultless disagreement," and that only the relativist can explain how faultless disagreement is possible (Kölbel, 2002, 2004*a*, 2008*a*). I have avoided using this phrase here, because it is dangerously ambiguous. Both "faultless" and "disagreement" can be understood in several ways, and how we understand them matters greatly for the plausibility of "faultless disagreement" and its significance for the debate about relative truth.

We have already discussed some possible senses of "disagreement." The ones that will matter most for us here will be

- $\textit{disagreement}_n$ = doxastic non-cotenability
- $\textit{disagreement}_p$ = preclusion of joint accuracy

What about "faultless"? What is it for a belief or assertion to be faultless? Here are four possibilities:

- $\textit{faultless}_w$ = epistemically warranted
- $\textit{faultless}_t$ = true
- $\textit{faultless}_a$ = accurate
- $\textit{faultless}_n$ = not in violation of constitutive norms governing belief/assertion

Now, what of the claim that faultless disagreement is possible?

- Clearly faultless$_w$ disagreement is possible, no matter what we mean by "disagreement." Two people can hold contradictory but equally warranted beliefs about a perfectly objective subject matter—say, the age of the earth—if one of them has misleading evidence. It does not seem, though, that one needs to invoke relative truth to explain faultless$_w$ disagreement.
- Faultless$_t$ disagreement is not possible on either construal of "disagreement." If you can coherently characterize another's belief as "true" (using the monadic propositional truth predicate), then you could come to have a belief with the same content without giving up any of your current beliefs, so the other's attitude is doxastically cotenable with your own. It is not coherent to say, "I disagree with you about that, but what you believe is true." (I suspect that many opponents of truth relativism take its goal to be vindicating faultless disagreement in this sense. Clearly *that* is not a viable goal.)
- Faultless$_a$ disagreement$_p$ is not possible. To say that joint accuracy is precluded is to say that at least one of the disagreeing attitudes must be at fault in the sense of being inaccurate.
- However, faultless$_a$ disagreement$_n$ *is* possible. As we have seen, on a nonindexical contextualist treatment of "tasty," Sam's and Sal's beliefs can both be accurate, even though they are not doxastically cotenable.
- Assuming, as we have been, that the norms governing assertion and belief are keyed to accuracy relative to the asserter's or believer's context of assessment, faultless$_n$ disagreement$_p$ is possible. Preclusion of joint accuracy means that there is no single context of assessment relative to which both beliefs or assertions are accurate. But it may be that the beliefs or assertions are both reflexively accurate—that is, both accurate as assessed by those who hold them. So it may be that both satisfy the relevant norms.

There are, then, at least three coherent, but very different, ways to construe relativists' claim that there is "faultless disagreement" in disputes of taste. (For a summary, see Table 6.2.) First, the claim might be that both parties in such a dispute can be *warranted* in holding

the views they do. Of course, this general phenomenon—disagreeing views, both of which are justified or warranted—is hardly distinctive of relativism. But it could be that truth relativism explains better than rival views how the things the disagreeing parties take to warrant their claims—their own subjective reactions—could possibly do so.

Second, the claim might be that two parties who hold contradictory beliefs might both be “getting it right,” in the sense that their beliefs are accurate. If this is the point, then the view supported is nonindexical contextualism, and the kind of disagreement at stake is doxastic non-cotenability.

Third, the claim might be that two parties whose beliefs or assertions preclude each others’ accuracy are both succeeding in living up to the norms governing the formation and retention of beliefs and the making and retracting of assertions. This is predicted by the kind of truth relativism we have been articulating in these pages—a truth relativism that countenances genuine assessment sensitivity and makes sense of it by relating truth at a context of assessment to norms governing belief and assertion. From Sam’s point of view, Sal’s assertion that the Hen-of-the-Woods is tasty is inaccurate. But Sam can agree that Sal has succeeding in conforming his assertions to the truth rule (as revised in §5.2.2), which only forbids Sal from asserting things that are inaccurate as assessed from his own perspective, and only requires him to retract things that are inaccurate as assessed from his own perspective.

Sense of “faultless”	w	w	t	t	a	a	n	n
Sense of “disagreement”	n	p	n	p	n	p	n	p
Standard contextualism	Y	Y	N	N	N	N	N	N
Nonindexical contextualism	Y	Y	N	N	Y	N	Y	N
Relativism	Y	Y	N	N	Y	N	Y	Y
Objectivism	Y	Y	N	N	N	N	N	N

Figure 6.2: Can there be “faultless disagreement” about matters of taste?

Kölbel's official definition of "faultless disagreement" does not by itself discriminate between these three construals:

> A faultless disagreement is a situation where there is a thinker *A*, a thinker *B*, and a proposition (content of judgement) *p*, such that:
>
> (a) *A* believes (judges) that *p* and *B* believes (judges) that not-*p*
>
> (a) Neither *A* nor *B* has made a mistake (is at fault). (Kölbel, 2004*a*, pp. 53–4)

As we have seen, condition (a) is too weak to capture preclusion of joint accuracy, so the notion of disagreement at stake here seems to be doxastic non-cotenability. (This may simply be an oversight, since Kölbel (2004*a*) does not have in view the distinction between relativism and nonindexical contextualism.) And the talk of "mistake" in condition (b) is too generic to select between the various interpretations of "faultlessness" considered above.

However, Kölbel's subsequent commentary points strongly towards construing "faultless" as *faultless*$_n$, and hence towards the third construal above. He endorses the principle

> TR It is a mistake to believe a proposition that is not true in one's own perspective,

which closely resembles our relativized truth norm (121). (He does not state the additional norm governing retraction that would be required to distinguish relativism from nonindexical contextualism.) So it seems that the "mistakes" he has in mind are violations of constitutive norms, and not (say) violations of epistemic norms.

If by "faultless disagreement" Kölbel means faultless$_n$ disagreement$_p$, then we can agree that the notion is coherent and distinctive of truth relativism. But it is that it is all too easy to give "faultless disagreement" other construals—including construals on which it is incoherent. Moreover, it is not clear we should expect people to have intuitive judgements about whether disagreements of taste are *faultless*$_n$: the notion of being faultless$_n$ is too theoretical, as it depends on one's views about constitutive norms governing assertion and belief. For all these reasons, we have avoided talk of "faultless disagreement" in motivating and explaining truth relativism.

Draft: December 13, 2009

Part II

Applications

7 *Tasty*

If "tasty" is assessment-sensitive, then we can acknowledge that each of the three standard views about "tasty" gets something right. The objectivist is right that there can be genuine disagreements about whether something whose taste both parties have experienced is tasty—disagreements that do not reduce to "disagreements in attitude" or to disagreements about the taste's of some particular person or group. The contextualist is right that "tasty" is, somehow, contextually sensitive to tastes. Finally, the expressivist is right that in calling something "tasty" one generally *expresses*, but does not *assert*, one's liking for it. For in asserting that something is tasty, one performs a speech act that is permitted only if one likes the food in question. One thereby give others reason to take one to like the food, and in that sense one expresses one's liking for it. But one is not *asserting* that one likes the food.

Thus a relativist account of "tasty" saves what seems right about each of the rival accounts. But it also avoids what seems wrong about them. Unlike objectivism, it does not require us to take ordinary speakers to be excessively chauvinistic about matters of taste. Unlike contextualism, it allows us to make sense of disagreement about claims of taste, and to understand why speakers would retract such claims when their tastes have changed. And unlike classical expressivism, it does not require us to abandon the force/content distinction and the framework of truth-conditional semantics. We can plug our semantics for "tasty" into existing truth-conditional frameworks, and we will get, for free, an account of the meanings of arbitrary sentences in which "tasty" occurs, no matter how deeply embedded.[1]

[1]I say "classical expressivism" because, as noted in §1.3, Gibbard's modern form

This chapter explores some issues about the compositional behavior of "tasty" and other assessment-sensitive expressions. "Tasty" is unlike some adjectives in allowing explicit relativization to a taster. We can say not just that the cookies are tasty, but that they are "tasty for kids" or "tasty to John." A compositional semantics for "tasty" ought to explain how these constructions work. It should also yield sensible predictions about embeddings of "tasty" under propositional attitude verbs ("Jill believes it is tasty"), quantifiers ("everyone got a tasty cookie"), tense modifiers ("it was tasty"), and modals ("it would be tasty"). All of these constructions raise interesting questions, and some have been used as the basis of arguments for contextualist treatments of "tasty," so it is important to explore what a relativist can say about them.

7.1 ATOMIC FORMULAS

We'll build up our semantics bit by bit, starting with atomic formulas and adding constructions that form more complex sentences one by one.

GRAMMAR

- *Singular terms:* "Joe", "Sally", "Hen-of-the-Woods", "Two Buck Chuck", "I", "you"
- *Variables:* "x", "y", "z"
- *One-place predicates:* "is tasty", "is poisonous", "is a person", "is a cookie".
- *Two-place predicates:* "likes the taste of", "gets"
- *Atomic formulas:* If α and β are singular terms or variables, Φ is a one-place predicate, and Ψ is a two-place predicate, $\ulcorner \alpha\ \Phi \urcorner$ and $\ulcorner \alpha\ \Psi\ \beta \urcorner$ are atomic formulas.

of expressivism appropriates all the machinery of truth-conditional semantics to give a unified account of content (normative, factual, and mixed). It is a subtle question, raised briefly at the end of §1.3, and to be discussed further in chapter **??**, what makes his view deserve the title "expressivism," and how it differs from the relativist view being described here.

Anticipating needs that will arise when we add quantifiers, tense and modal operators, and other constructions to our language, we define an index as follows:

INDEX An *index* is a quadruple $\langle w, t, g, a \rangle$, where w is a possible world, t a time, g a gustatory standard or "taste," and a an assignment of values to the variables.

We use a time parameter here because we will be treating tense and temporal modifiers as sentential operators. There are strong considerations in favor of treating tense using quantifiers instead (see King, 2003). However, the operator approach allows a simpler account of predicates and predication, so we prefer it here. Our discussion below of interactions with tense can be transposed, with small modifications, to either framework.

The "gustatory standard" coordinate also deserves notice. Lasersohn (2005) uses a *judge* coordinate, rather than a taste or standard of taste, in his indices. This has a couple of drawbacks. First, a single judge may have different tastes at different times. So, strictly speaking, what we need is not just a judge, but a judge and a time. We already have a time coordinate in our indices, in order to deal with tense and temporal modifiers. But, as I will argue in §7.9, below, if we let *that* be the relevant time, we get incorrect results. We could add a *second* time coordinate, whose role is just to answer the question, "the judge's tastes *when*?" But it seems simpler just to let the taste, or "gustatory standard," serve itself as the coordinate.

Second, if we use a judge rather than a gustatory standard, we close off the possibility that some judge-dependent expressions are assessment-sensitive while others are use-sensitive. For if, in the definition of truth at a context of use and context of assessment, the judge is initialized by the context of assessment, *every* expression whose extension is sensitive to the judge will be assessment-sensitive. This excludes, for example, combining a relativist treatment of "tasty" with a nonindexical contextualist treatment of "beautiful." The point here is not that we should want such a combination, but that our basic framework should not rule it out from the start. If we have separate coordinates for gustatory and aesthetic standards, we have

the possibility of letting one be initialized by the context of assessment and the other by the context of use.[2]

One final note on the notion of a gustatory standard. The talk of "standards" can suggest something intellectual: a set of principles the agent uses in assessing whether something is tasty. Nothing like that is intended here. Think of a standard, rather, as something that determines a scale. The International Prototype Kilogram in Sèvres, France is, in this sense, a standard for weight—the "standard kilogram."[3] One's "tastes," too, determine a gustatory standard, quite independently of whether one can articulate this standard.

We can now define truth at a context of use and index for atomic formulas. We do this by defining the extension of arbitrary expressions at a context and index. The extension of a term at a context and index is an object, the extension of a one-place predicate is a set, the extension of a two-place predicate is a set of pairs, and the extension of a formula is a truth value.

SEMANTICS We use $|\alpha|_{c,\langle w,t,g,a\rangle}$ to denote the extension of α at $c, \langle w, t, g, a\rangle$.

Singular terms:

- $|$"Joe"$|_{c,\langle w,t,g,a\rangle}$ = Joe
- $|$"Sally"$|_{c,\langle w,t,g,a\rangle}$ = Sally
- $|$"Hen-of-the-Woods"$|_{c,\langle w,t,g,a\rangle}$ = Hen-of-the-Woods
- $|$"Two Buck Chuck"$|_{c,\langle w,t,g,a\rangle}$ = Two Buck Chuck[4]
- $|$"I"$|_{c,\langle w,t,g,a\rangle}$ = the agent of c
- $|$"you"$|_{c,\langle w,t,g,a\rangle}$ = the addressee at c

Variables:

- $|\alpha|_{c,\langle w,t,g,a\rangle} = a(\alpha)$, where α is a variable

One-place predicates:

- $|$"is tasty"$|_{c,\langle w,t,g,a\rangle}$ =
  $\{x \mid x$ is tasty at world w and time t by the gustatory standard $g\}$

[2] Of course, it would also be possible to have two separate judge coordinates.

[3] At least it was in 2009. It is possible that by the time this sees print, this standard will have been replaced by a definition in terms of natural constants.

[4] A notoriously cheap California wine.

- $|\text{“is poisonous”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{x \mid x \text{ is poisonous at world } w \text{ and time } t\}$
- $|\text{“is a person”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{x \mid x \text{ is a person at world } w \text{ and time } t\}$
- $|\text{“is a cookie”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{x \mid x \text{ is a cookie at world } w \text{ and time } t\}$

Two-place predicates:

- $|\text{“likes the taste of”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{\langle x, y\rangle \mid x \text{ likes the taste of } y \text{ at world } w \text{ and time } t\}$
- $|\text{“gets”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{\langle x, y\rangle \mid x \text{ gets } y \text{ at world } w \text{ and time } t\}$

Atomic formulas:

- $|\alpha\ \Phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\alpha|_{c,\langle w,t,g,a\rangle} \in |\Phi|_{c,\langle w,t,g,a\rangle} \\ \text{False} & \text{otherwise} \end{cases}$
- $|\alpha\ \Psi\ \beta|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } \langle |\alpha|_{c,\langle w,t,g,a\rangle}, |\beta|_{c,\langle w,t,g,a\rangle}\rangle \\ & \in |\Psi|_{c,\langle w,t,g,a\rangle} \\ \text{False} & \text{otherwise.} \end{cases}$

7.2 Postsemantics

Note that up to this point, we have not needed to mention contexts of assessment. That is because contexts of assessment are not *locally relevant* in the sense of §3.2.3. Contexts of assessment are needed only in the next phase, the definition of truth relative to a context of use and context of assessment in terms of truth at a context of use and index. To distinguish this phase from the definition of truth at a context of use and index, we call it “postsemantics.”

Postsemantics A sentence S is true as used at context c_1 and assessed from a context c_2 iff for all assignments a, $|S|_{c_1,\langle w_{c_1},t_{c_1},g_{c_2},a\rangle}$ = True, where w_{c_1} is the world of c_1, t_{c_1} is the time of c_1, and g_{c_2} is the gustatory standard (“taste”) of the agent of c_2 (that is, the assessor’s taste at the time of c_2).

At this stage, we can also define some LOGICAL NOTIONS Two sentences S, T are *strictly equivalent* iff for all c, w, t, g, a, $|S|_{c,\langle w,t,g,a\rangle} = |T|_{c,\langle w,t,g,a\rangle}$.

Two sentences S, T are *logically equivalent* iff for all contexts c_1, c_2, S is true as used at c_1 and assessed from c_2 iff T is true as used at c_1 and assessed from c_2.

Two sentences S, T are *diagonally equivalent* iff for all contexts c, S is true as used at c and assessed from c iff T is true as used at c and assessed from c.

A sentence T is a *logical consequence* of a set Γ of sentences iff for all contexts c_1, c_2, if every member of Γ is true as used at c_1 and assessed from c_2, then T is true as used at c_1 and assessed from c_2.

A sentence T is a *diagonal consequence* of a set Γ of sentences iff for all contexts c, if every member of Γ is true as used at c and assessed from c, then T is true as used at c and assessed from c.

7.3 CONTENTS AND CIRCUMSTANCES

The semantics is not a two-stage semantics; it does not associate sentences with propositions and then define truth for propositions, but defines truth directly for sentences. But, following Kaplan (1989, p. 546), we can superimpose a theory of contents (properties and propositions) onto this semantics.

CIRCUMSTANCE OF EVALUATION Let a *circumstance of evaluation* be a triple $\langle w, t, g\rangle$, where w is a world, t a time, and g a gustatory standard.

CONTENT Where α is a formula, predicate, or singular term, let $[\![\alpha]\!]^a_c$ denote its content at context of use c under the assignment a.

INTENSIONS OF CONTENTS The intension of $[\![\alpha]\!]^a_c$ is the function f from circumstances of evaluation to extensions such that $f(\langle w, t, g\rangle) = |\alpha|_{c,\langle w,t,g,a\rangle}$.

Contents are relativized to contexts of use, but not to contexts of assessment, because we are developing a form of truth-value relativism, not content relativism. Since the semantics we have given implies that $|\text{“is tasty”}|_{c_1,\langle w,t,g,a_1\rangle} = |\text{“is tasty”}|_{c_2,\langle w,t,g,a_2\rangle}$ for all $c_1, c_2, w, t, g, a_1, a_2$, the intension of “is tasty” is independent of the context of use and the assignment, and we can consistently stipulate that “tasty” invariantly expresses a single property, the property of being tasty.

7.4 Truth functional connectives

It is easy to add truth-functional sentential connectives to our language:

Grammar Where ϕ and ψ are formulas, $\ulcorner\neg\phi\urcorner$, $\ulcorner\phi \wedge \psi\urcorner$, and $\ulcorner\phi \vee \psi\urcorner$ are formulas.

Semantics

$$|\neg\phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} False & \text{if } |\phi|_{c,\langle w,t,g,a\rangle} = \text{True} \\ \text{True} & \text{otherwise.} \end{cases}$$

$$|\phi \wedge \psi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\phi|_{c,\langle w,t,g,a\rangle} = \text{True} \\ & \text{and } |\psi|_{c,\langle w,t,g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\phi \vee \psi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\phi|_{c,\langle w,t,g,a\rangle} = \text{True} \\ & \text{or } |\psi|_{c,\langle w,t,g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

These clauses are simple and straightforward, and they allow us to extend our account of truth at a context of use and context of assessment from atomic sentences to arbitrary truth-functional compounds of such sentences. Operating in a truth-conditional framework gives us a simple solution to the embedding problem that proves so difficult for classical expressivism (§1.3.2).

7.5 Explicit relativizations

As noted in §1.2, “tasty” can be explicitly relativized to a judge:

(148) That brand of peanut butter is very tasty to young kids.

(149) Yuk, that isn't tasty at all to me.

(150) Have you got anything that will be tasty to everybody?

There is reason to think this capacity for relativization is a semantic feature of "tasty," because similar relativizations seem out of place for other kinds of adjectives—even adjectives whose application requires a judgement call:

(151) #Sam is strong to/for young kids.[5]

(152) #Sam may be bald to/for you, but he isn't bald to/for me at all.

(153) #Can you send someone who will be intelligent to/for everybody?

It is difficult to see how an objectivist about "tasty" would account for these data. And it is easy to see how a certain kind of contextualist can account for them. If "tasty" expresses the relational property of being *tasty to* a judge, then we should expect to find sentences in which the judge is explicitly specified, as well as sentences in which it is not (in which case the judge argument place is to be filled by context). But what account of (148–150) can a relativist give?

The approach favored here, due to Lasersohn (2005), is to treat "tasty to Sal" as a complex predicate, as follows:

GRAMMAR Where α is a one-place predicate and β a singular term or variable, $\ulcorner \alpha \textit{ to } \beta \urcorner$ is a one-place predicate.[6]

SEMANTICS $|\alpha \textit{ to } \beta|_{c,\langle w,t,g,a\rangle} = |\alpha|_{c,\langle w,t,g',a\rangle}$, where g' = the gustatory standard of $|\beta|_{c,\langle w,t,g,a\rangle}$ at t.[7]

[5] "Sam is strong for a kid" is okay, but here "for a kid" gives a comparison class, not a judge. To say that Sam is strong for a kid is not to say that he is strong as judged by the standards of kids, but that he is stronger than the norm for kids. If it meant the former rather than the latter, then "Hulk Hogan is strong for a kid" should be equally acceptable, and it is not.

[6] Perhaps α should be required to be a predicate whose extension varies with the gustatory standard of the index. For simplicity, we do not require that here, so "is a cookie to Joe" counts as grammatically well-formed, though semantically it will be equivalent to "is a cookie."

[7] This clause differs slightly from the one at Lasersohn (2005, p. 666), because his indices contain judges while mine contain gustatory standards.

We can now see how, on a relativist account, the intension of "tasty" (as used at some context) differs from the intension of "tasty to me" (as used at the same context). As we look at different circumstances of evaluation that agree on the world and time but differ in the gustatory standard, the extension of "tasty" varies, while the extension of "tasty to me" stays the same. Logically speaking, "Two Buck Chuck is tasty" and "Two Buck Chuck is tasty to me" are neither strictly equivalent nor logically equivalent. But they are *diagonally equivalent*; for any context c, "Two Buck Chuck is tasty" is true as used at and assessed from c just in case "Two Buck Chuck is tasty to me" is true as used at and assessed from c. This means that the Truth Rule for assertion will license asserting one just when it licenses asserting the other, which explains why it is so tempting to think that they are equivalent in some stronger sense.[8]

Stephenson (2007) takes an interesting alternative approach. Like Lasersohn, she includes a "judge" coordinate in her indices, but instead of taking "tasty" to be a one-place predicate whose extension is sensitive to this coordinate, she takes it to be a two-place predicate. The extra argument place can be filled either (i) by a regular pronoun (as in "tasty to Sal"), (ii) by a semantically equivalent null pronoun pro_{Sal}, or (iii) by a special null pronoun PRO_J, which denotes (at any context and index) the judge of the index. When "tasty" occurs without explicit qualification, Stephenson holds, its underlying syntax can be either (ii)—which is semantically equivalent to (i)—or (iii), whose extension is sensitive to the judge coordinate of the index.

Stephenson's view differs from ours both in its syntax and in its semantics. The syntactic difference is that "tasty" is taken to be a two-place predicate. The main semantic difference is that a judge is used instead of a gustatory standard; this has implications for temporal embeddings, which will be discussed below. Our view could be revised to resemble Stephenson's syntactically without matching it semantically. We could think of "tasty" as a two-place predicate with the meaning: "ψ tastes good by the gustatory standards ξ." The semantic value "to" in "to Sal" could be understood as a function from judges to their gustatory standards at the time of the index. Instead of PRO_J, which denotes the judge of the index, we could use PRO_G,

[8]Compare cf. Lasersohn (2005, 688 Remarks 8 and 9).

which denotes the gustatory standard of the assessor at the time of assessment (or more flexibly, the gustatory standard that is relevant at the context of assessment). The resulting view would agree with ours in its predictions about the truth (at a context of use, context of assessment, and index) of every sentence. The syntactic issues, then, can be factored out from the semantic ones.

7.6 Implicit relativizations

Lasersohn (2005) points out that there are some cases in which predicates like "tasty" that are clearly intended to be evaluated with respect to a particular judge or standard. He calls these uses "exocentric", distinguishing them from the more usual "autocentric" uses, which we evaluate relative to ourselves as judge (whether as speaker or as a third-party assessor). For example, in the dialogue

(154) Mary: How did Bill like the rides?
John: Well, the merry-go-round was fun, but the water slide was a little too scary,

"we intuitively regard John's utterance as true if the merry-go-round was fun for Bill, not if it was fun for ourselves (or for John)" (Lasersohn, 2005, p. 672). Similarly, in buying dog food we might ask ourselves,

(155) I wonder which brand is most tasty?

In answering the question, we try to figure out which brand would be most tasty to the dog, not to ourselves.

Lasersohn thinks that these uses need to be distinguished from autocentric uses in the pragmatics. He takes John to be asserting with (154) the very same proposition he would be asserting if he used the same sentence autocentrically. A proper appreciation of the significance of John's speech act, then, requires not just a grasp of its force and content, but also an awareness of whether John was adopting an exocentric or an autocentric stance towards the asserted proposition. If we judge that he was adopting an autocentric stance, it is appropriate for us to take an autocentric stance in evaluating his claim. But if we judge that he was adopting an exocentric stance, then we should evaluate his claim relative to the judge he intended.

This approach is incompatible with the framework being developed here, on which it is an intrinsic property of a *content* that it is assessment-sensitive (or not). So I am committed to offering an alternative account of exocentric uses. It seems to me that in (154), John is not asserting the same proposition he would be asserting if he used "fun" autocentrically. He is, rather, asserting what would be literally expressed by the sentence "the merry-go-round was fun *for Bill*, but the water slide was a little too scary *for him*." Here "for Bill" works in much the same way as "to Bill" in "tasty to Bill"; it converts assessment-sensitive predicates ("fun," "scary") into non-assessment-sensitive ones. John doesn't use the words "for Bill," because he doesn't need to: it is obvious from context which proposition he is asserting, and he uses the minimum possible linguistic resources to get that across.

It is not difficult to find other examples of this kind of linguistic flexibility. John might say

(156) June's a friend, but Sam is an enemy,

and thereby assert that June is a friend *of Bill* and that Sam is an enemy *of Bill*. This kind of laziness is to be expected. We tend not to make things explicit unless our audience is likely to misunderstand us.[9]

One might object that if we allow this kind of flexibility, John should be able to utter (156) and thereby assert that June is a friend of Bill and that Sam is a enemy of Sarah. But it is easy to explain why this should be difficult. It would take a very special kind of contextual setup for hearers to be able to fill in the relativizations that way. Without such a setup in place, speakers cannot reasonably expect hearers to divine their intentions without explicit relativizations. So they cannot expect to assert this proposition using (156).

Aside from the fact that it is compatible with the framework of chapters 3–4, there are two additional considerations favoring the content-centered approach over Lasersohn's use-centered approach. First, as we will see in §7.8, we need to posit implicit relativizations

[9]Stephenson's approach, discussed in §7.5, above, is similar to this one in taking the difference between exocentric and autocentric uses to be a difference in the *contents* of the asserted propositions. But on her view, there is also a syntactic difference: in the exocentric uses, the "judge" argument place of "tasty" is filled by a null pronoun, while in the autocentric uses, it is filled by PRO_J.

anyway to account for binding phenomena. If we need these resources anyway, there is no additional cost to using them here.

Second, as we will see in §7.7, Lasersohn's account of attitude verbs is warped by his account of exocentric uses. By contrast, the implicit content approach advocated here is compatible with a simple and conservative semantics for attitude verbs.[10]

7.7 ATTITUDE VERBS

The simplest semantics for attitude verbs takes them to express relations between persons and contents.[11] For example:

GRAMMAR

- "believes" is a two-place predicate.
- If ϕ is a sentence (formula with no free variables), ⌜that ϕ⌝ is a singular term.

[10]Stephenson (2007) criticizes Lasersohn's account of exocentric uses on other grounds. She thinks that exocentric uses are available for "tasty," but not for epistemic modals, and that this must be explained by a semantic difference between them. (She takes "tasty," but not epistemic modals, to have an extra argument place for a judge.) Thus, for example, in

(1) The cat food might be tasty,

"tasty" can mean "tasty to the cat" or it can be interpreted in the relativist ("judge-dependent") way, but (she claims) "might" can only have the judge-dependent meaning (Stephenson, 2007, p. 499). As she notes, it is hard to see how Lasersohn could explain this (assuming he took "might" to be judge-dependent), because taking an exocentric perspective on which the relevant judge is the cat would affect "tasty" and "might" equally. The problem is less severe on the view being proposed here, because "tasty" and "might" will be affected by different coordinates of the index. But it is not clear what kind of principled explanation the present view could give for the unavailability of a reading that implicitly qualifies "might" as "might, as far as the cat knows" and "tasty" as "tasty to the cat."

[11]Another alternative, pioneered by Hintikka (1962) and favored by Stephenson (2007) and other linguists, construes attitude verbs as modal operators: roughly, "Joe believes that p" is true at w just in case p is true at all the worlds w' that are doxastic live possibilities for Joe at w. (Stephenson argues, plausibly, that these operators shift the judge as well.) We prefer the relational analysis here, because it more smoothly handles complements that are not that-clauses ("Joe believes Goldbach's conjecture") and avoids the Hintikka approach's commitment to the closure of belief over entailment.

SEMANTICS

- $|\text{“believes”}|_{c,\langle w,t,g,a\rangle}$ =
  $\{\langle x, y\rangle \mid x$ has a belief with content y at world w and time $t\}$
- $|\ulcorner\text{that } \phi\urcorner|_{c,\langle w,t,g,a\rangle} = [\![\phi]\!]^a_c$

On Lasersohn's view, however, this cannot suffice. For a believer can take either an autocentric or an exocentric perspective towards a proposition she believes. And one may believe a proposition exocentrically while not believing it autocentrically. It may be, for example, that Joe does not believe (autocentrically) that the dog food is tasty, but does believe (exocentrically, taking the dog as relevant judge) that the dog food is tasty. Thus, Lasersohn concludes, "we must now treat *believe* as a 3-place relation between an individual, a context, and a sentence content" (Lasersohn, 2005, p. 676). The context argument allows us to distinguish between autocentric and various exocentric uses.[12]

The suggestion is implausible, because there is no independent evidence that "believe" has a third argument place for a context, judge, or standard. If "believe" did have such an argument place, it should be possible to specify a value for it, and to bind it with quantifiers. There should be a natural English way to express explicitly the thought that Joe believes that the dog food is tasty taking the dog as judge, or that for every judge, Joe believes that snow is white taking that judge as judge. But these seem to be things we can express only in Lasersohn's own quasi-technical metalanguage; there are no natural

[12]In Lasersohn (forthcoming), Lasersohn distinguishes "believe" from "consider," which he takes to be two-place, always requiring an autocentric perspective. For comment, see Cappelen and Hawthorne (2009, 106 n. 11). Lasersohn also raises the possibility that genuinely autocentric uses of "believe" express a two-place relation, like "consider." For the three-place relation cannot distinguish between autocentric uses and exocentric uses that are targeted on the speaker, but not in a *de se* way:

> While riding the roller coaster past the mirror, John mistakes his own reflection for someone else, realizes that "that person" must enjoy roller coasters, and assesses *Roller coasters are fun* as true relative to "that person." Surely we should count this as taking an exocentric stance rather than an autocentric one. (Lasersohn, forthcoming, §4).

English equivalents. This strikes heavily against the proposal that "believe" in English has the third argument place.

If, instead of taking the difference between exocentric and autocentric perspectives to be pragmatic, we understand it as a difference in the *contents* of the relevant beliefs, there is no longer any need to take "believe" to have an extra argument place. We simply say that the sentence

(157) I believe that California Natural dog food is tasty.

can (depending on the context) be used to assert either that the speaker stands in the belief relation to the proposition that California Natural dog food is tasty, or that he stands in the belief relation to the proposition that California Natural dog food is tasty *to the dog*.

7.8 QUANTIFIERS AND BINDING

We can add binary quantifiers to the language in the usual way:

GRAMMAR Where ϕ and ψ are formulas and α a variable, $\ulcorner All_\alpha(\phi,\psi)\urcorner$, $\ulcorner Some_\alpha(\phi,\psi)\urcorner$, and $\ulcorner Most_\alpha(\phi,\psi)\urcorner$ are formulas.

SEMANTICS

Notation: $a[x/\alpha](y) = x$ if $y = \alpha$ and $a(y)$ otherwise.

$$|All_\alpha(\phi,\psi)|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for every object } x \text{ such that} \\ & |\phi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True}, \\ & |\psi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|Some_\alpha(\phi,\psi)|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some object } x \text{ such that} \\ & |\phi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True}, \\ & |\psi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|Most_\alpha(\phi,\psi)|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for most objects } x \text{ such that} \\ & |\phi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True}, \\ & |\psi|_{c,\langle w,t,g,a[x/\alpha]\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

The quantifiers shift the assignment function and leave the other coordinates of the index, including the gustatory standard, alone. This creates a prima facie problem with sentences like

(158) Every person gets some tasty cookies
All_x(x is a person, $Some_y$(y is a cookie$\wedge y$ is tasty, x gets y))[13]

(158) seems to have (at least) two different readings. On the first reading, it says that every person gets some cookies that are tasty (full stop). One should endorse the claim, thus construed, only if one finds every cookie given out to be tasty. On the second reading, it says that every person gets some cookies that are tasty *to that person*. One can endorse the claim, thus construed, even if one finds some of the cookies given out disgusting. Suppose that children and adults have different tastes in cookies. The green cookies are tasty to children, but not to adults, while the red cookies are tasty to adults, but not to children. All the children get two green cookies (and no red ones), and all the adults get two red cookies (and no green ones). On the first reading, (158) is not true as assessed by an adult; on the second reading, it is true.

The problem is that the relativist account seems to predict only the first reading. By contrast, a standard sort of contextualist account—one that takes "tasty" to work like "local"—easily predicts both readings. If "tasty" has an extra argument place for a judge or gustatory standard, then we should expect that it can be either filled by context (the first reading) or bound by a quantifier (the second reading).

There are two ways for a relativist to meet the objection. The first way is to join the contextualist in positing an extra argument place in "tasty," but argue, with Stephenson, that it can sometimes be filled by a special null pronoun whose denotation is determined by features of the context of assessment. A relativist who goes this way can give the same explanation of the ambiguity as the contextualist. Indeed, Stephenson's account seems to predict that there should be *three* interpretations of (158):

(159) a. All_x(x is a person, $Some_y$(y is a cookie $\wedge$
y is tasty to PRO_J, x gets y))

[13]Cf. Lasersohn (2005, p. 681).

b. All_x(x is a person, $Some_y$(y is a cookie ∧ y is tasty to pro_{Sally}, x gets y))

c. All_x(x is a person, $Some_y$(y is a cookie ∧ y is tasty to x, x gets y))

(159a) is assessment-sensitive; it is true (relative to a context of assessment) if every person gets a cookie whose taste is pleasing to the assessor. (159b) involves an exocentric use of "tasty;" it is true if every person gets a cookie whose tasty is pleasing to Sally. (159c) is the bound reading; it is true if every person gets a cookie whose taste is pleasing to him or her.

Another approach—one that is consistent with the semantics given above, which does not take "tasty" to have an extra argument place—is to hold that the sentence (158) can be used to assert the proposition that would be literally expressed by

(160) All_x(x is a person, $Some_y$(y is a cookie ∧ y is tasty *to* x, x gets y)),

using the predicate-modifying operator "ϕ *to* α" defined above. Here no syntactic mechanism is posited that gets us this interpretation. Instead, we simply suppose that the speaker expects the hearer to be able discern that the proposition expressed by (160) is the one she intends to assert. The binding is implicit, not explicit.

The second approach is in many ways more conservative, as it does not require positing a hitherto unknown syntactic element, PRO_J. It will not be acceptable to those who reject implicit content (most notably Stanley (2007)). This is not the place to enter into an extended discussion of that debate. But the availability of the first approach as another option shows that the debate over relativism is not going to be settled if the argument is settled in Stanley's favor.[14]

[14]For several different kinds of defenses of implicit content against Stanley's arguments, see, for example, Recanati (2003), Carston (2004), Neale (2004), Collins (2007). See also Lasersohn (2005, p. 681) for a syntactic argument against any proposal that explains the bound readings by positing an extra argument place in "tasty" that gets filled by a syntactically realized but unpronounced pronoun. (This would include Stephenson's proposal, discussed above, and the kinds of contextualism Stanley favors.)

7.9 Tense

For present purposes, we treat temporal modifiers as sentential operators, in the tradition of tense logic. As noted above, nothing essential hangs on this decision. We need a definite syntax and semantics in order to discuss interactions between temporal modifiers and "tasty," and the operator approach is simple. But the discussion below could be recast in a framework that treats temporal modifiers quantificationally.

Grammar Where ϕ is a formula, ⌜*Now:* ϕ⌝ ⌜*Will:* ϕ⌝, ⌜*Was:* ϕ⌝, ⌜*Yesterday:* ϕ⌝, ⌜*Tomorrow:* ϕ⌝, and ⌜*One year ago:* ϕ⌝ are formulas.

Semantics

$$|\textit{Now: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\phi|_{c,\langle w,t_c,g,a\rangle} = \text{True}, \\ & \text{where } t_c \text{ is the time of } c \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\textit{Will: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some } t' > t, \\ & |\phi|_{c,\langle w,t',g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\textit{Was: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some } t' < t, \\ & |\phi|_{c,\langle w,t',g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\textit{Yesterday: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some } t' \text{ belonging to} \\ & \text{the day before the day of } c, \\ & |\phi|_{c,\langle w,t',g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\textit{Tomorrow: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some } t' \text{ belonging to} \\ & \text{the day after the day of } c, \\ & |\phi|_{c,\langle w,t',g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\textit{One year ago: } \phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\phi|_{c,\langle w,t',g,a\rangle} = \text{True}, \\ & \text{where } t' = t - 1 \text{ year} \\ \text{False} & \text{otherwise.} \end{cases}$$

These temporal operators shift the time coordinate of the index, but they leave the gustatory standard coordinate untouched. This has a consequence that might seem surprising. Suppose that one's tastes change. At c_1, one likes the taste of Hen-of-the-Woods, while at c_2 (one year later), one dislikes the taste—not because the taste has changed, but because one's reactions to it have changed. Then at c_2 one can permissibly assert not only

(161) Hen-of-the-Woods is not tasty
¬(Hen-of-the-Woods is tasty)

but also

(162) Hen-of-the-Woods was not tasty a year ago
¬*One year ago:* Hen-of-the-Woods is tasty.

Some readers may disagree with this prediction, and think that it should be correct in such a situation to assert

(163) Hen-of-the-Woods was tasty a year ago
One year ago: Hen-of-the-Woods is tasty.

Lasersohn (2005) and Stephenson (2007) both offer semantics for "tasty" that give this result. On their views, "is tasty" (or, for Stephenson, "is tasty PRO_J") is true of an object at an index if the judge of the index likes the taste of the object at the time of the index. So, even if the flavor of Hen-of-the-Woods has not changed through an interval,

(164) Hen-of-the-Woods was tasty before, but is not tasty any longer

can be true, because the judge's tastes have changed during that interval.

Here is one way of seeing the difference between the two approaches. On Lasersohn's and Stephenson's approaches, the time of the index plays a double role. It tells us not only what time-slice of the object to look at (which is important because objects can change their flavors over time), but also what time-slice of the judge to look at (which is important because judges can change their tastes over time). On the approach being recommended here, by contrast, the time of the index plays only the first role; there is no need to determine which time-slice of the judge is relevant, because the index already contains a complete gustatory standard.

Although it may at first seem a good feature of Lasersohn's and Stephenson's views that they endorse (163), there are several good reasons for rejecting the judge-in-the-index approach. First, it predicts that

(165) Hen-of-the-Woods will still be tasty in fifty years

should entail

(166) Someone will be alive in fifty years.

(I assume here that it is not the case that any dead creature likes the taste of Hen-of-the-Woods, or any other food, while dead.) But this entailment seems dubious.[15]

Second, it would seem odd to say:

(167) Last year Hen-of-the-Woods was tasty, but this year it isn't. It has exactly the same flavor this year that it did last year, but after my mushroom tasting class I now find it unappealing.

By contrast, it would not be at all odd to say:

(168) Last year Hen-of-the-Woods was tasty to me, but this year it isn't. It has exactly the same flavor now this year that it did last year, but after my mushroom tasting class I now find it unappealing.

But the judge-in-index view would not predict a difference in acceptability here.

Third, endorsing (163) does not sit well with saying that an assertion in c_1 of

(169) Hen-of-the-Woods is tasty

must be retracted in c_2. It would be odd (at the very least) to say:

(170) Last year I asserted that Hen-of-the-Woods was tasty. And last year Hen-of-the-Woods *was* tasty. Still, my assertion was not accurate and must be retracted.

This is a decisive reason for an assessment-sensitive semantics for "tasty" to reject the judge-in-index approach. (It is not clear that it would be a decisive reason for Lasersohn and Stephenson, though, since they may not accept the retraction prediction.)

[15] See Lasersohn (2005, 663 n. 13) for a similar point (attributed to an anonymous referee), using a modal operator rather than a temporal one.

7.10 ALETHIC MODALS AND COUNTERFACTUALS

Alethic necessity and possibility operators and counterfactual conditionals can be added in the standard way.

GRAMMAR Where ϕ and ψ are formulas, $\ulcorner\Box\phi\urcorner$, $\ulcorner\Diamond\phi\urcorner$, and $\ulcorner\phi \boxright \psi\urcorner$ are formulas.

SEMANTICS

$$|\Box\phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for all worlds } w' \text{ accessible from} \\ & w, |\phi|_{c,\langle w',t,g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\Diamond\phi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if for some world } w' \text{ accessible from} \\ & w, |\phi|_{c,\langle w',t,g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

$$|\phi \boxright \psi|_{c,\langle w,t,g,a\rangle} = \begin{cases} \text{True} & \text{if } |\psi|_{c,\langle w',t,g,a\rangle} = \text{True, where } w' \\ & \text{is the closest world to } w \text{ (by the} \\ & \text{metric relevant at } c\text{) such that} \\ & |\phi|_{c,\langle w',t,g,a\rangle} = \text{True} \\ \text{False} & \text{otherwise.} \end{cases}$$

These connectives shift the world coordinate of indices, leaving the other coordinates—including the gustatory standard—alone. Thus,

(171) Hen-of-the-Woods could have been tasty
◊Hen-of-the-Woods is tasty

is accurate (as assessed by someone with gustatory standard g) just in case Hen-of-the-Woods could have had a flavor that is tasty according to g. The fact that the assessor (or anyone else) could have had a different gustatory standard than g is irrelevant to the truth of (171). This is as it should be. To wonder what things would be like if horse manure were tasty is to wonder what things would be like if horse manure had a different flavor than it in fact has, not to wonder what things would be like if one had (say) the gustatory standards of a dog.

For similar reasons, although the counterfactual conditional

(172) If I had not trained my pallet on many better wines, Two Buck Chuck would be tasty to me

is true,

(173) If I had not trained my pallet on many better wines, Two Buck Chuck would be tasty

is false (as used and assessed by me now). For, the closest possible world where I had not tried many better wines—call it w'—is presumably a world where Two Buck Chuck has the same flavor it has in the actual world. And it is false, in w', relative to my (the assessor's) actual standard of taste, that Two Buck Chuck is tasty.[16]

One might have supposed that on the relativist view, the property of tastiness is mind-dependent; after all, our tastiness judgements seem to be projections of our own reactions onto the things that cause them. But in at least one sense of "mind-dependent," this charge is not valid. For the counterfactual

(174) If no sentient beings had ever existed, nothing would be tasty

comes out false on the proposed semantics. What matters for its truth is not whether, in the imagined humanless world, anything would be pleasing to a creature, but whether the flavors things would have had in such a world are pleasing by the gustatory standard of the assessor (here, us). This means that quick defenses of the objectivity of values that appeal to our intuitions about such conditionals cannot rule out relativist views.[17]

[16] More formally, $|\text{Two Buck Chuck is tasty } \textit{to } I|_{c,\langle w',t_c,g_c,a\rangle}$ = True (for any assignment a), whereas $|\text{Two Buck Chuck is tasty}|_{c,\langle w',t_c,g_c,a\rangle}$ = False (for any a).

[17] Indeed, many contextualist views will also pass muster—those that take the proposition expressed to concern the speaker's standards, rather than the speaker herself. And some relativist views, including those of Lasersohn and Stephenson, will have trouble with (174); see Lasersohn (2005, 663 n. 13). This is another reason to put a standard of taste, rather than a judge, in the index.

Bibliography

Ayer, A. J. (1959). *Language, Truth, and Logic*. New York: Dover. See pp. 21–23.

Bach, Kent and Robert M. Harnish (1979). *Linguistic Communication and Speech Acts*. Cambridge: MIT Press. See pp. 125, 126, 129, 131.

Baghramian, Maria (2004). *Relativism*. New York: Routledge.

Bar-Hillel, Y. (1973). "Primary Truth Bearers". In: *Dialectica* 27, pp. 303–312. See p. 65.

Belnap, Nuel and Mitchell Green (1994). "Indeterminism and the Thin Red Line". In: *Philosophical Perspectives* 8, pp. 365–388. See p. 77.

Belnap, Nuel, Michael Perloff, and Ming Xu (2001). *Facing the Future: Agents and Choices in Our Indeterministic World*. Oxford: Oxford University Press. See pp. 77, 98.

Bennigson, Thomas (1999). "Is Relativism Really Self-Refuting?" In: *Philosophical Studies* 94, pp. 211–236. See pp. 43, 46–48, 57.

Blackburn, Simon (1984). *Spreading the Word*. Oxford: Oxford University Press. See pp. 28–30, 32.

— (1988). "Attitudes and Contents". In: *Ethics* 98, pp. 501–517. See pp. 31, 32.

— (1993). *Essays in Quasi-Realism*. New York: Oxford University Press. See p. 32.

Boghossian, Paul (2006). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press. See pp. 43, 47, 48.

Brandom, Robert (1983). "Asserting". In: *Noûs* 17, pp. 637–650. See pp. 126, 132, 133.

— (1994). *Making It Explicit*. Cambridge: Harvard University Press. See pp. 113, 127, 132, 133.

Burke, T. E. (1979). "The Limits of Relativism". In: *Philosophical Quarterly* 29, pp. 193–207. See p. 44.

Burnyeat, M. F. (1976*a*). "Protagoras and Self-Refutation in Later Greek Philosophy". In: *Philosophical Review* 85, pp. 44–69. See p. 45.

— (1976*b*). "Protagoras and Self-Refutation in Plato's *Theaetetus*". In: *Philosophical Review* 85, pp. 172–195. See pp. 45–48, 132.

Burnyeat, Myles (1990). *The Theaetetus of Plato*. Indianapolis: Hackett.

Campbell, John (1997). "The Realism of Memory". In: *Language, Thought, and Logic*. Oxford: Oxford University Press, pp. 157–181. See p. 123.

Cappelen, Herman and John Hawthorne (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press. See pp. 12, 15, 101, 142, 151, 152, 177.

Cappelen, Herman and Ernie Lepore (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell. See pp. 7, 99.

Carston, Robyn (2004). "Explicature in Semantics". In: *Semantics: A Reader*. Ed. by S. Davis and B. Gillon. Oxford: Oxford University Press, pp. 817–845. See p. 180.

Cartwright, Richard (1962). "Propositions". In: *Analytic Philosophy*. Ed. by R. Butler. Vol. 1. Blackwell. See p. 67.

Collins, John (2007). "Syntax, More or Less". In: *Mind* 116, pp. 806–850. See p. 180.

Davidson, Donald (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press. See p. 36.

— (1990). "The Structure and Content of Truth". In: *Journal of Philosophy* 87, pp. 279–328. See pp. 65, 112.

— (1997). "The Folly of Trying to Define Truth". In: *Journal of Philosophy* 94, pp. 263–278. See pp. 110, 112.

de Sa, Dan López (2008). "Presuppositions of Commonality: An Indexical Relativist Account of Disagreement". In: *Relative Truth*. Ed. by Manuel García-Carpintero and Max Köbel. Oxford: Oxford University Press, pp. 297–310. See pp. 18, 155.

DeRose, Keith (2004). "Single Scoreboard Semantics". In: *Philosophical Studies* 119, pp. 1–21. See pp. 18, 20.

Dreier, James (1996). "Expressivist Embeddings and Minimalist Truth". In: *Philosophical Studies* 83, pp. 29–51. See p. 34.

— (1999). "Transforming Expressivism". In: *Noûs* 33, pp. 558–572. See pp. 23, 34.

Dummett, Michael (1959). “Truth”. In: *Proceedings of the Aristotelian Society* n.s. 59, pp. 141–62. See pp. 36, 104, 110.

— (1978). *Truth and Other Enigmas*. Cambridge: Harvard University Press. See pp. 112, 114.

— (1981). *Frege: Philosophy of Language*. 2nd. Cambridge: Harvard University Press. See pp. 111, 112, 114.

Egan, Andy (forthcoming). “De Gustibus Non Disputandum Est (At Least, Not Always)”. In: *Disagreement*. Ed. by Feldman and Warfield. URL: `http://sitemaker.umich.edu/egana/files/de_gustibus.2007.10.18.pdf`. See p. 13.

Egan, Andy, John Hawthorne, and Brian Weatherson (2005). “Epistemic Modals in Context”. In: *Contextualism in Philosophy*. Ed. by G. Preyer and P. Peter. Oxford: Oxford University Press. Chap. 6, pp. 131–168. See pp. 102, 120.

Evans, Gareth (1985). “Does Tense Logic Rest Upon a Mistake?” In: *Collected Papers*. Oxford: Oxford University Press, pp. 343–363. See pp. 71, 123.

Feldman, Richard (2001). “Skeptical Problems, Contextualist Solutions”. In: *Philosophical Studies* 103, pp. 61–85. See p. 20.

Fellman, Anita Clair (2008). *Little House, Long Shadow: Laura Ingalls Wilder’s Impact on American Culture*. Columbia: University of Missouri Press. See p. 55.

Field, Hartry (1994). “Deflationist Views of Meaning and Content”. In: *Mind* 103, pp. 249–85. See p. 113.

Fine, Gail (19983). “Plato’s Refutation of Protagoras in the *Theaetetus*”. In: *Apeiron* 31, pp. 201–234. See p. 45.

Fitelson, Branden (2005). “Inductive Logic”. In: *Philosophy of Science: An Encyclopedia*. Ed. by J. Pfeifer and S. Sarkar. New York: Routledge. See p. 57.

Fox, John F. (1994). “How Must Relativism Be Construed to Be Coherent?” In: *Philosophy of the Social Sciences* 24, pp. 55–75. See pp. 53, 56.

Frege, Gottlob (1979). *Posthumous Writings*. Ed. by Friedrich Kaulbach Hans Hermes Friedrich Kambartel. Chicago: University of Chicago Press. See pp. 93, 141, 142.

Geach, Peter (1960). “Ascriptivism”. In: *Philosophical Review* 69, pp. 221–225. See pp. 27, 28.

Geach, Peter (1965). "Assertion". In: *Philosophical Review* 74, pp. 449–465. See p. 28.

Gibbard, Allan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press. See pp. 32, 33, 35.

— (2003). *Thinking How to Live*. Cambridge: Harvard University Press. See pp. 32–35.

Glanzberg, Michael (2007). "Context, Content, and Relativism". In: *Philosophical Studies* 136, pp. 1–29. See pp. 6, 101.

— (forthcoming). "Semantics and truth relative to a world". In: *Synthese*. URL: http://dx.doi.org/10.1007/s11229-007-9285-3.

Goodman, Nelson (1979). *Fact, Fiction, and Forecast*. Cambridge: Harvard University Press. See p. 57.

Grice, Paul (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press. See p. 9.

Groenendijk, Jeroen and Martin Stokhof (1997). "Questions". In: *Handbook of Logic and Language*. Ed. by J. van Benthem and A. Ter Meulen. Cambridge: MIT Press, pp. 1055–1124. See p. 26.

Hale, Bob (1986). "The Compleat Projectivist". In: *Philosophical Quarterly* 36, pp. 65–84. See pp. 29, 31.

Hales, Steven D. (1997*a*). "A Consistent Relativism". In: *Mind* 106, pp. 33–52. See pp. 44, 45.

— (1997*b*). "A Reply to Shogenji on Relativism". In: *Mind* 106, pp. 749–750. See p. 46.

Hamblin, C. L. (1973). "Questions in Montague English". In: *Foundations of Language* 10, pp. 41–53. See p. 26.

Harman, Gilbert (1975). "Moral Relativism Defended". In: *Philosophical Review* 84, pp. 3–22. See p. 49.

Hintikka, Jaakko (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press. See p. 176.

Horn, Laurence R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press. See p. 10.

Horwich, Paul (1998). *Meaning*. Oxford: Basil Blackwell. See p. 113.

Husserl, Edmund (2001). *Logical Investigations*. Vol. 1. Translated by J. N. Findlay. New York: Routledge. See p. 44.

Huvenes, Torfinn (2008). "Varieties of Disagreement and Predicates of Taste". manuscript, University of St. Andrews. See p. 145.

Jackson, Dennis E. (1917). "An Experimental Investigation of the Pharmacological Action of Nitrous Oxid". In: *American Journal of Surgery* 31, p. 70. See p. 55.

Jackson, Frank and Philip Pettit (1998). "A problem for expressivism". In: *Analysis* 58, pp. 239–251. See pp. 21, 23.

James, William (1909). *The Meaning of Truth*. Longmans, Green, and Co. See p. 53.

— (1978). *Pragmatism and The Meaning of Truth*. Cambridge: Harvard University Press. See pp. 53, 57.

Kamp, Hans (1971). "Formal Properties of 'Now'". In: *Theoria* 37, pp. 227–273. See p. 74.

Kaplan, David (1989). "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". In: *Themes from Kaplan*. Ed. by Joseph Almog, John Perry, and Howard Wettstein. Oxford: Oxford University Press, pp. 481–566. See pp. 62, 65, 68–72, 74, 75, 77, 80, 81, 88–90, 93, 96–99, 170.

Karttunen, Lauri (1977). "Syntax and Semantics of Questions". In: *Linguistics and Philosophy* 1, pp. 3–44. See p. 26.

Kennedy, Christopher (2007). "Vagueness and grammar: the semantics of relative and absolute gradable adjectives". In: *Linguistics and Philosophy* 30, pp. 1–45. See p. 6.

King, Jeffrey C. (2003). "Tense, Modality, and Semantic Values". In: *Philosophical Perspectives* 17, pp. 195–245. URL: http://dx.doi.org/10.1111/j.1520-8583.2003.00009.x. See pp. 69, 84, 87, 90, 93, 167.

Kölbel, Max (2002). *Truth Without Objectivity*. London: Routledge. See pp. 10, 28, 31, 34, 44, 47, 49, 67, 120, 139, 159.

— (2004*a*). "Faultless Disagreement". In: *Proceedings of the Aristotelian Society* 104. See pp. 159, 162.

— (2004*b*). "Indexical Relativism versus Genuine Relativism". In: *International Journal of Philosophical Studies* 12, pp. 297–313. See p. 146.

— (2008*a*). "Introduction: Motivations for Relativism". In: *Relative Truth*. Ed. by Manuel García-Carpintero and Max Kölbel. Oxford: Oxford University Press, pp. 1–38. See pp. 68, 159.

Kölbel, Max (2008*b*). ""True" as Ambiguous". In: *Philosophy and Phenomenological Research*. See p. 53.

Kripke, Saul (1972). "Naming and Necessity". In: *Semantics of Natural Languages*. Ed. by Donald Davidson and Gilbert Harman. Dordrecht: Reidel, pp. 254–355. See p. 68.

Kruger, Justin and David Dunning (1999). "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments". In: *Journal of Personality and Social Psychology* 77, pp. 1121–1135. URL: http://www.apa.org/journals/features/psp7761121.pdf. See pp. 4, 5.

Künne, Wolfgang (2003). *Conceptions of Truth*. Oxford: Oxford University Press. See p. 58.

Lasersohn, Peter (2005). "Context Dependence, Disagreement, and Predicates of Personal Taste". In: *Linguistics and Philosophy* 28, pp. 643–686. URL: http://dx.doi.org/10.1007/s10988-005-0596-x. See pp. 7, 14, 16, 18, 98, 167, 172–174, 177, 179, 180, 182, 183, 185.

— (forthcoming). "Relative Truth, Speaker Commitment, and Control of Implicit Arguments". In: *Proceedings of NELS* 37. URL: http://semanticsarchive.net/Archive/jIOMmNlO. See p. 177.

Lewis, David (1970*a*). "Anselm and Actuality". In: *Noûs* 4, pp. 175–188. See p. 92.

— (1970*b*). "General Semantics". In: *Synthese* 22, pp. 18–67. See p. 77.

— (1979*a*). "Attitudes De Dicto and De Se". In: *Philosophical Review* 87, pp. 513–545. See pp. 108, 148.

— (1979*b*). "Scorekeeping in a Language Game". In: *Journal of Philosophical Logic* 8, pp. 339–359. See p. 17.

— (1980). "Index, context, and content". In: *Philosophy and Grammar*. Ed. by Stig Kanger and Sven Öhman. Dordrecht: Reidel. See pp. 36, 71, 72, 77, 87, 90, 112.

— (1983). "Languages and Language". In: *Philosophical Papers*. Vol. 1. Oxford: Oxford University Press, pp. 163–188. URL: http://dx.doi.org/10.1093/0195032047.003.0011. See pp. 36, 73, 112.

— (1986). *On the Plurality of Worlds*. Oxford: Blackwell. See pp. 68, 77, 103.

— (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press. See pp. 72, 73, 75, 76, 83.

Lockie, Robert (2003). “Relativism and Reflexivity”. In: *International Journal of Philosophical Studies* 11, pp. 319–339. See p. 47.

MacFarlane, John (2003). “Future Contingents and Relative Truth”. In: *Philosophical Quarterly* 53.212, pp. 321–336. See pp. 64, 76, 132.

— (2005*a*). “Making Sense of Relative Truth”. In: *Proceedings of the Aristotelian Society* 105, pp. 321–339. See pp. 50, 98, 102, 132.

— (2005*b*). “The Assessment Sensitivity of Knowledge Attributions”. Ed. by Tamar Szabó Gendler and John Hawthorne. In: *Oxford Studies in Epistemology* 1, pp. 197–233. See p. 127.

— (2007). “Relativism and Disagreement”. In: *Philosophical Studies* 132, pp. 17–31. URL: http://www.springerlink.com/content/q9163h52l58v3108/. See pp. 124, 152.

— (2008). “Truth in the Garden of Forking Paths”. In: *Relative Truth*. Ed. by Manuel García-Carpintero and Max Köbel. Oxford: Oxford University Press. See p. 77.

— (2009). “Nonindexical Contextualism”. In: *Synthese* 166, pp. 231–50. URL: http://dx.doi.org/10.1007/s11229-007-9286-2. See pp. 83, 100, 156.

— (forthcoming). “What Is Assertion?” See p. 21.

Maclay, Kathleen (2001). *UC Berkeley's Davitt Moroney shares abiding love of music through teaching and performing*. UC Berkeley press release. Oct. 2001. URL: http://berkeley.edu/news/media/releases/2001/10/25_moroney.html. See p. 55.

Margolis, Joseph (1991). *The Truth About Relativism*. Oxford: Blackwell. See p. 43.

Maudlin, Tim (2007). “Review of Patrick Grennough and Michael Lynch, eds., *Truth and Realism*”. In: *Notre Dame Philosophical Reviews* 2007.06.18. URL: http://ndpr.nd.edu/review.cfm?id=10143. See p. 22.

Meiland, Jack (1977). “Concepts of Relative Truth”. In: *The Monist* 60, pp. 568–582. See pp. 50, 58, 61, 109.

Meiland, Jack W. and Michael Krausz, eds. (1982). *Relativism: Cognitive and Moral*. Notre Dame: University of Notre Dame Press. See pp. 51, 56.

Mellor, D. H. (1981). *Real Time*. Cambridge: Cambridge University Press. See p. 99.

Neale, Stephen (2004). "This, That, and the Other". In: *Descriptions and Beyond*. Ed. by Anne Bezuidenhout and Marga Reimer. Oxford: Oxford University Press, pp. 68–182. See p. 180.

Newton-Smith, W. (1981). *The Rationality of Science*. New York: Routledge. See p. 44.

— (1982). "Relativism and the Possibility of Interpretation". In: *Rationality and Relativism*. Ed. by Martin Hollis and Steven Lukes. Cambridge: MIT Press. See p. 44.

Nozick, Robert (2001). *Invariances: The Structure of the Objective World*. Cambridge: Harvard University Press. See pp. 50, 53, 68, 70.

Nunberg, Geoffrey (1979). "The non-uniqueness of semantic solutions: Polysemy". In: *Linguistics and Philosophy* 3, pp. 143–184. See p. 130.

Owens, David (2006). "Testimony and Assertion". In: *Philosophical Studies* 130, pp. 105–129. See pp. 126, 130.

Pagin, Peter (2004). "Is assertion social?" In: *Journal of Pragmatics* 36, pp. 833–859. See p. 132.

Passmore, John (1961). *Philosophical Reasoning*. London: Duckworth. See p. 46.

Patterson, Douglas (2005). "Deflationism and the Truth Conditional Theory of Meaning". In: *Philosophical Studies* 124, pp. 271–294. See p. 113.

Peirce, C. S. (1934). "Belief and Judgement". In: ed. by Charles Hartshorne and Paul Weiss. Cambridge: Harvard University Press. See pp. 126, 132.

Percival, Philip (1989). "Indices of Truth and Temporal Propositions". In: *Philosophical Quarterly* 39, pp. 190–199. See p. 99.

— (1994). "Absolute Truth". In: *Proceedings of the Aristotelian Society* 94, pp. 189–213. See pp. 64, 102, 113, 123, 125.

Perry, John (1986). "Thought without Representation". In: *Proceedings of the Aristotelian Society*. See pp. 95, 152.

— (2001). *Reference and Reflexivity*. Stanford: CSLI. See p. 64.

Predelli, Stefano (1998). "I am not here now". In: *Analysis* 58, pp. 107–115. See p. 80.

— (2005). *Contexts: Meaning, Truth, and the Use of Language*. Oxford: Oxford University Press. See p. 72.

Preyer, G. and G. Peter, eds. (2007). *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*. Oxford: Oxford University Press. See p. 7.

Prior, Arthur (1957). *Time and Modality*. Westport, Connecticut: Greenwood Press. See p. 69.

— (2003). *Papers on Time and Tense*. Ed. by Per Hasle et al. New. Oxford: Oxford University Press. See p. 69.

Putnam, Hilary (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press. See pp. 43, 47, 49.

Ramsey, F. P. (2001). "The Nature of Truth". In: *The Nature of Truth*. Ed. by Michael P. Lynch. Cambridge: MIT Press, pp. 433–446. See p. 53.

Recanati, François (2003). *Literal Meaning*. Cambridge: Cambridge University Press. See p. 180.

Richard, Mark (1980). "Temporalism and Eternalism". In: *Philosophical Studies* 39, pp. 1–13. See p. 69.

— (1982). "Tense, Propositions, and Meanings". In: *Philosophical Studies* 41, pp. 337–351. See p. 69.

— (2003). "Introduction to Part I". In: *Time, Tense, and Reference*. Ed. by Aleksandar Jokić and Quentin Smith. Cambridge: MIT Press, pp. 27–45. See p. 69.

— (2004). "Contextualism and Relativism". In: *Philosophical Studies* 119, pp. 215–242. See p. 15.

Rorty, Richard (1998). *Truth and Progress: Philosophical Papers, Vol. 3*. Cambridge: Cambridge University Press. See p. 43.

Salmon, Nathan (1986). *Frege's Puzzle*. Cambridge, MA: MIT Press. See p. 69.

— (2003). "Tense and Intension". In: *Time, Tense, and Reference*. Ed. by Aleksandar Jokić and Quentin Smith. Cambridge: MIT Press, pp. 109–154. See p. 69.

Schiffer, Stephen (1972). *Meaning*. Oxford: Oxford University Press. See p. 130.

Schlenker, Philippe (2003). "A Plea for Monsters". In: *Linguistics and Philosophy* 26, pp. 29–120. See p. 74.

— (2004). "Context of Thought and Context of Utterance (A Note on Free Indirect Discourse and the Historical Present)". In: *Mind and Language* 19, pp. 279–304. See p. 80.

Searle, John R. (1962). "Meaning and Speech Acts". In: *Philosophical Review* 71, pp. 423–432. See p. 28.

— (1969). *Speech Acts*. Cambridge: Cambridge University Press. See pp. 26, 28, 132.

— (1979). *Expression and Meaning*. Cambridge: Cambridge University Press. See pp. 26, 126, 132.

Sheffer, Gabriel (2003). *Diaspora Politics: At Home Abroad*. Cambridge: Cambridge University Press. See p. 55.

Shissler, A. Holly (2003). *Between Two Empires: Ahmet Agaoglu and the New Turkey*. New York: I. B. Tauris. See p. 54.

Shogenji, Tomoji (1997). "The Consistency of Global Relativism". In: *Mind* 106, pp. 745–747. See p. 46.

Soames, Scott (2002). "Replies". In: *Philosophy and Phenomenological Research* 65, pp. 429–452. See p. 99.

Stalnaker, Robert (1978). "Assertion". In: *Syntax and Semantics, Vol. 9: Pragmatics*. Ed. by P. Cole. New York: Academic Press. See p. 72.

— (1987). *Inquiry*. Cambridge: MIT Press. See pp. 68, 92.

Stanley, Jason (2000). "Context and Logical Form". In: *Linguistics and Philosophy* 23, pp. 391–434. See p. 98.

— (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press. See pp. 68, 90, 99, 101.

— (2007). *Language in Context: Selected Essays*. Oxford: Oxford University Press. See pp. 7, 180.

Stephenson, Tamina (2007). "Judge dependence, epistemic modals, and predicates of personal taste". In: *Linguistics and Philosophy* 30, pp. 487–525. See pp. 173, 176, 182.

Stevenson, Charles L. (1963). *Facts and Values*. New Haven: Yale University Press. See p. 147.

Stevenson, Leslie (1988). "Can Truth Be Relativized to Kinds of Mind?" In: *Mind* 97, pp. 281–284. See p. 44.

Strawson, P. F. (1950). "Truth". In: *Proceedings of the Aristotelian Society* s.v. 24, pp. 129–156. See p. 65.

Swoyer, Chris (1982). "True For". In: *Relativism: Cognitive and Moral*. Ed. by Jack W. Meiland and Michael Krausz. Notre Dame: University of Notre Dame Press, pp. 84–108. See pp. 43, 44, 50, 56.

Tarski, Alfred (1935). "Der Wahrheitsbegriff in den formalisierten Sprachen". In: *Studia Philosophica* 1, pp. 261–405. See p. 62.

— (1983). “The Concept of Truth in Formalized Languages”. In: *Logic, Semantics, Metamathematics*. Ed. by John Corcoran. Indianapolis: Hackett. Chap. 8, pp. 152–278. See p. 62.
Thomason, Richmond H. (1970). “Indeterminist Time and Truth-Value Gaps”. In: *Theoria* 36, pp. 264–281. See p. 77.
Unwin, Nicholas (1987). “Beyond Truth: Towards a New Conception of Knowledge and Communication”. In: *Mind* 96, pp. 299–317. See p. 53.
Vallicella, William F. (1984). “Relativism, Truth and the Symmetry Thesis”. In: *The Monist* 67, pp. 452–456. See pp. 47, 56.
Velleman, J. David (2000). *The Possibility of Practical reason*. Oxford: Oxford University Press. See p. 139.
Watson, Gary (2004). “Asserting and Promising”. In: *Philosophical Studies* 117, pp. 57–77. See pp. 127, 132.
Wedgwood, Ralph (2002). “The Aim of Belief”. In: *Philosophical Perspectives* 16, pp. 267–297. See p. 139.
Weiner, Matthew (2005). “Must We Know What We Say?” In: *Philosophical Review* 114, pp. 227–251. See p. 116.
White, F. C. (1986). “On a Proposed Refutation of Relativism”. In: *Australasian Journal of Philosophy* 64, pp. 331–334. See p. 44.
Whyte, Jamie T. (1993). “Relativism is absolutely false”. In: *Cogito* 7, pp. 112–118. See p. 43.
Wiggins, David (1980). “What Would Be a Substantial Theory of Truth?” In: *Philosophical Subjects: Essays Presented to P. F. Strawson*. Ed. by Zak van Straaten. Oxford: Clarendon Press, pp. 189–221. See p. 112.
Williams, Bernard (2002). *Truth and Truthfulness*. Princeton: Princeton University Press. See pp. 126, 129, 130.
Williams, Michael (1999). “Meaning and Deflationary Truth”. In: *Journal of Philosophy* 96, pp. 545–564. See p. 113.
Williamson, Timothy (1996). “Knowing and Asserting”. In: *Philosophical Review* 105, pp. 489–523. See pp. 114, 116, 126.
— (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press. See pp. 116, 117, 126.
Wright, Crispin (1992). *Truth and Objectivity*. Cambridge: Harvard University Press. See pp. 127, 132.
— (2001). “On Being in a Quandary: Relativism, Vagueness, Logical Revisionism”. In: *Mind* 110, pp. 45–98. See p. 38.

Wright, Crispin (forthcoming). “Fear of Relativism?” In: *Philosophy and Phenomenological Research*. See p. 49.

Zaret, Barry L. and Genell J. Shubak-Sharpe (2007). *Heart Care for Life: Developing the Program that Works Best for You*. New Haven: Yale University Press. See p. 55.

Zimmerman, Aaron (2007). “Against Relativism”. In: *Philosophical Studies* 133, pp. 313–348. URL: http://dx.doi.org/10.1007/s11098-005-4604-3. See pp. 66, 68, 101, 139.

Index