

Truth, Relativism, and the Objective Epistemic Ought

Jack Spencer

Working Draft of 18-Apr-2021

This essay is concerned with what truth relativists should think it takes for beliefs and assertions to be correct. After introducing some terminology, I put forward three conceptions of correctness that truth relativists could adopt: my preferred view, which I call ‘the targeted view’, and two initially plausible alternatives, which I call ‘the simple view’ and ‘the reflexive view’. I then defend the targeted view, by arguing that it avoids the problems that beset the simple and reflexive views.

§1/ Relative Truth

Let me begin with some terminology and assumptions.

I take a *circumstance* to be an n -tuple that contains a value for each of the n parameters across which truth-value can vary. A circumstance thus might contain a possible world, an instant of time, an information state, a gustatory propensity, and so on. A momentary agent *occupies* a circumstance just if each of the circumstance’s coordinates are the agent’s own. If w is the world that is actual for me, t is the time that is present for me, s is my information state, and so on, then $\langle w, t, s, \dots \rangle$ is the circumstance I occupy.¹

A circumstance is *actual* if its worldly coordinate is the actual world, *present* if its temporal coordinate is the present moment, and *occupied* if there is, in fact, some agent who occupies it. There are non-present agents, but there are no

¹ Circumstances are thus like what others call centered words; see Liao (2012).

non-actual agents—or so I will assume. So some non-present circumstances are occupied, but no non-actual circumstances are occupied.

I take *propositions* to be the objects of belief and assertion, and I assume that every circumstance decides the truth-value of every proposition. If c is a circumstance and p is a proposition, then p is either true-at- c or false-at- c , and not both.

I assume that beliefs and assertions inherit their alethic properties from their propositional contents:²

Inheritance. A belief or assertion is true/false just if its propositional content is true/false, and a belief or assertion is true-at- c /false-at- c just if its propositional content is true-at- c /false-at- c .

I also assume that truth is stable. According to a familiar claim in modal logic: relative to any world w , a proposition is true just if it's true-at- w . The claim I assume generalizes this familiar thought to circumstances:

Stability. Relative to any circumstance c , a proposition is true just if it is true-at- c .

I will say that a belief or assertion is *absolutely* true/false if it is true/false at every occupied circumstance, and I will say that a belief or assertion *has its truth-value relatively* if its truth-value differs across occupied circumstances. There are lots of interesting truth relativistic theses, but the one I want to focus upon is the following:

² This assumption is common; see *e.g.* Frege (1918[1956]) and Horwich (1998).

Relative Truth. Some beliefs and assertions have their truth-values relatively.

Hereafter, I take Relative Truth to be the thesis that divides truth relativists and truth absolutists.³

§2/ Correctness and the Objective Epistemic Ought

There is a familiar distinction drawn between two moral oughts: a subjective moral ought, which is sensitive to the agent's evidence, and an objective moral ought, which isn't. The distinction is often illustrated with mineshaft cases.⁴

A similar distinction can be drawn between epistemic oughts.⁵ There is a subjective epistemic ought, which is sensitive to the agent's evidence, and an objective epistemic ought, which isn't. Consider the following two examples:

White Wall. Lily, staring at a wall that looks to her to be red, believes and asserts that the wall is red. (In fact, Lily is staring at a white wall bathed in red light.)

³ Using 'truth relativism' in this way is not idiosyncratic; see *e.g.* Glanzberg (2007), Kölbel (2002; 2004a; 2004b), Nozick (2001), and Zimmerman (2007). But some authors reserve the term 'truth relativism' for a thesis more akin to what I call 'correctness relativism'; see *e.g.* MacFarlane (2014: 49-68).

⁴ See *e.g.* Regan (1980).

⁵ See *e.g.* Gibbard (2005).

Winning Ticket. Billy knows that a fair lottery was drawn yesterday. The results are not yet announced. Billy knows that besides his one ticket, there are 999 others. Nevertheless, he believes and asserts that his ticket has won. (And, in fact, perchance, his ticket has won.)

Evaluated along a subjective epistemic dimension, Lily does better. Her belief and assertion are supported by her evidence; Billy's are not. But evaluated along an objective epistemic dimension, Billy does better. Billy, despite his irrationality, is right—his ticket has won. And Lily, despite her rationality, is wrong—the wall before her is not red.

The words 'correct' and 'incorrect' have various senses in English, but my use of the terms will be semi-technical.⁶ I take both to be defined in terms of the objective epistemic ought. A belief or assertion is *correct* if it satisfies the demands of the objective epistemic ought, and *incorrect*, otherwise. Billy's belief and assertion are correct, and Lily's are incorrect.⁷

Correctness is sometimes introduced by an analogy to archery.⁸ It is said that beliefs and assertions have aims, and that they are correct just when they hit their aims. Since I think beliefs and assertions have aims only metaphorically, I prefer the deontic characterization. But nothing turns on this preference. Those who prefer to talk in terms of aims can take the objective epistemic ought to demand that beliefs and assertions hit their aims.

⁶ My use of 'correct' is not idiosyncratic; *cf.* Rosen (2001) and Thomson (2008).

⁷ Note that I do not think that correctness applies to propositions or representations, *pace* Kölbel (2015a). We do not want contents, themselves, to bear normative properties; *cf.* Boghossian (2003). Rather we want things we do with contents—beliefs and assertions—to be the bearers of those normative properties.

⁸ See *e.g.* Sosa (2011) and Williams (1970).

When it comes to beliefs and assertions that have their truth-values absolutely, as both Lily's and Billy's do, correctness is straightforward. Let the *correctness-value* of a belief or assertion be its status as correct or incorrect, and let's assume that every circumstance decides the correctness-value of every belief and assertion. I will say that a belief or assertion is *absolutely* correct/incorrect if it is correct/incorrect at every occupied circumstance, and I will say that a belief or assertion *has its correctness-value relatively* just if it differs in correctness-value across occupied circumstances. Every plausible conception of correctness entails:

Absolute Minimum. If a belief or assertion is absolutely true/false, it is absolutely correct/incorrect.⁹

If Relative Truth were false, then we could accept both of the following theses:

Absolute Correctness. No belief or assertion has its correctness-value relatively.

Normative Truth. A belief or assertion is correct just if it's true.

Indeed, if Relative Truth were false, Absolute Minimum would entail both Absolute Correctness and Normative Truth. But if Relative Truth is true, then correctness is not nearly so straightforward; for we then have an inconsistent triad. Relative Truth, Absolute Correctness, and Normative Truth cannot all be true.

⁹ How exactly the connection between absolute truth and absolute correctness should be formulated is a matter of some controversy; see *e.g.* Bykvist and Hattiangadi (2007). But the main concerns of this essay are orthogonal to those controversies.

This inconsistent triad sets the stage for our inquiry. What it takes for beliefs and assertions to be correct is one of the most important metaepistemological questions, and truth relativists have a choice. Either they accept Normative Truth, reject Absolute Correctness, and take beliefs and assertions to aim at truth, even when truth is a relative matter, or they accept Absolute Correctness, reject Normative Truth, and take beliefs and assertions to aim at something other than truth. In the next section, I will put forward three conceptions of correctness that truth relativists might defend. The first accepts Normative Truth and rejects Absolute Correctness. The other two accept Absolute Correctness and reject Normative Truth. All three entail Absolute Minimum.

§3/ Three Conceptions of Correctness

3.1 The simple view

The first option is *the simple view*. According to the simple view, what the objective epistemic ought demands of beliefs and assertions is truth. Every true belief and assertion is correct. Every false belief and assertion is incorrect. Some beliefs and assertions have their truth-values relatively, and every belief and assertion that has its truth-value relatively has its correctness-value relatively.¹⁰

Of the three conceptions of correctness I'll be considering, the simple view is the only one that verifies Normative Truth. Many philosophers think that

¹⁰ The simple view closely resembles what MacFarlane (2005; 2007a; 2007b; 2009; 2011; 2014) calls "truth relativism" and what others, in reference to assessment sensitivity, call "assessor-relativism" or "assessmentism." The simple view is not formulated in the framework in which MacFarlane operates, so it's difficult to determine how exactly similar the simple view is to MacFarlane's.

Normative Truth is obviously, perhaps even analytically, true,¹¹ so verifying Normative Truth is a powerful point in favor of the simple view.

But the simple view is a staggering thesis. Truth relativists who accept the simple view are not just truth relativists; they're *correctness relativists*. Moral relativism is relativism about the objective moral ought, and correctness relativism is relativism about the objective epistemic ought. If the simple view is true, then epistemology is, deep down, at its core, relative.

3.2/ The reflexive view

The second option is *the reflexive view*. According to the reflexive view, the objective epistemic ought demands, not that beliefs and assertions be true, but that they be *reflexively true*, i.e., true at the circumstance occupied by the agent to whom they belong.¹²

The reflexive view falsifies Normative Truth, so it predicts that there are *truth-correctness gaps*: beliefs and assertions that are either both false and correct or both true and incorrect. The reflexive view does not predict truth-correctness gaps among *your* beliefs and assertions. If you occupy circumstance *c*, then, according to the reflexive view, your beliefs and assertions are correct just if true-at-*c*. Relative to *c*, beliefs and assertions are true just if true-at-*c*. So, relative to *c* (your circumstance), your beliefs and assertions are correct just if true. But the reflexive view does predict truth-correctness gaps among the beliefs and assertions

¹¹ See *e.g.* Frege (1918[1956]), Gibbard (2005), Humberstone (1992), and Williams (1970).

¹² The reflexive view resembles what MacFarlane (2009; 2014) calls “nonindexical contextualism,” although so too does the targeted view, which I introduce in the next section. The reflexive view also resembles the sorts of truth relativism that have been defended by Brogaard (2008; 2010; 2012), Egan (2007; 2009; 2010), Egan, Hawthorne, and Weatherson (2005), Kölbel (2002; 2004a; 2004b), Recanati (2007; 2008; 2017), and Richard (2004; 2015).

of others. Suppose (i) that you occupy circumstance c_1 , (ii) that agent A occupies circumstance c_2 , (iii) that A believes and asserts that p , and (iv) that p is true-at- c_2 and false-at- c_1 . Then, according to reflexive view, relative to c (your circumstance), A 's belief and assertion are correct, on account of being true-at- c_2 , but false, on account of being false-at- c_1 .

Unlike the simple view, the reflexive view verifies Absolute Correctness. Suppose that A occupies circumstance c , and suppose that A believes and asserts that p . It may be a relative matter whether A 's belief and assertion are true: p may be true at some occupied circumstances and false at others. But it cannot be a relative matter whether A 's belief and assertion are reflexively true: p cannot be true-at- c at some occupied circumstances and false-at- c at others. The reflexive view thus entails that every belief and assertion has its correctness-value absolutely.

The reflexive view could be called 'the self-locating view', since reflexive truth is the crux of self-location: agents locate themselves in logical space by determining whether a proposition is reflexively true.

3.3 The targeted view

The simple and reflexive views are perhaps the first two conceptions of correctness that truth relativists might consider, but I think that there is reason to prefer a third alternative, *the targeted view*.

According to the targeted view, every belief and assertion has a *targeting function*, which maps each possible world to a set of circumstances that have that possible world as their worldly coordinate. If t_1 and t_2 are different targeting functions, then two agents who believe that p might, strictly speaking, take different propositional attitudes: one might believe-for- t_1 that p , and the other might believe-for- t_2 that p . Similarly, two agents who assert that p might, strictly

speaking, engage in different propositional activities: one might assert-to- t_1 that p , and the other might assert-to- t_2 that p .

If t is the targeting function of some belief or assertion and $w@$ is the actual world, then $t(w@)$ is the *actual target* of the belief or assertion. And according to the targeted view, beliefs and assertions aim at their actual targets. An agent who believes-for- t that p takes p to be true at every circumstance in $t(w@)$, and does correctly just if p is true at every circumstance in $t(w@)$. An agent who asserts-to- t that p puts p forward as true at every circumstance in $t(w@)$, and does so correctly just if p is true at every circumstance in $t(w@)$. What the objective epistemic ought demands of beliefs and assertions, according to the targeted view, is that they be true at (i.e. throughout) their actual targets.¹³

A targeted conception of assertion is natural; indeed, it's somewhat surprising that the received conception is untargeted. There are rare cases when a speaker addresses the entire universe in making an assertion. (The hubristic villain says, looking skyward, "Nothing will stop me from becoming king!") But an

¹³ The targeted view most closely resembles the views defended by Recanati (2007), Richard (2015), and Spencer (2016). The distinction between the content of an attitude and its target is much like the distinction that Perry (1986), drawing on Austin (1950), draws between what beliefs are *about* and what they *concern*. Also see Köbel (2017), Recanti (2017), and Szabó (2017).

Lasersohn (2005) gives a semantics for 'belief' that, if interpreted metaphysically, resembles the targeted view. Lasersohn (2016) rejects the view defended in Lasersohn (2005) on the grounds that 'belief' does not appear, syntactically, to be ternary. Lasersohn's change of mind helps illustrate the difference between the proposals. The targeted view is offered as a metaphysics for belief and assertion, not a semantics, so syntactic considerations cut very little ice against it. The targeted view predicts that the belief-like relation we bear to proposition is ternary, but it does not predict that 'belief' is. It could be, for example, that an agent is truly said to "believe that p " just if the agent bears some belief-for attitude toward p .

assertion is usually addressed, not to the entire universe, but to some proper part thereof. The targeting function represent this directedness.

A targeted conception of belief is less immediately natural, but we can get the idea on the table by comparing it to an untargeted conception.¹⁴

On an untargeted conception of belief, an agent who believes that p takes p to be true, and this fact is borne out in their dispositions. Some of the dispositions are intrapersonal. Someone who believes that p tends to be disposed to use p as a premise in practical and theoretical reasoning, tends to believe obvious deductive consequences of p , and tends to be disposed to abandon the attitude in the face of strong evidence that $\sim p$. Some of the dispositions are interpersonal. Someone who believes that p tends to be disposed to assert p , tends to regard others who believe that $\sim p$ as mistaken, and tends to take themselves to disagree with those who believe that $\sim p$.

On a targeted conception of belief, an agent who believes-for- t that p takes p to be true at $t(w@)$, and *this* fact is borne out in their dispositions. If the actual target of the agent's belief includes the circumstance the agent occupies, as it often will, the intrapersonal dispositions will be the same. The agent will tend to be disposed to use p as a premise in practical and theoretical reasoning, will tend to believe obvious deductive consequences of p , and will tend to abandon the attitude in the face of strong evidence that $\sim p$. The interpersonal dispositions will be the same in kind, but different in scope. A disposition to assert p to A , or to regard A 's belief that $\sim p$ as mistaken, or to take oneself to disagree with A 's belief that $\sim p$, tends to be associated with a belief of p whose actual target includes A 's circumstance.¹⁵

¹⁴ This explication of the targeted view draws on Spencer (2016).

¹⁵ Let *me_danger* be the *de se* proposition that I am in danger. If I am sure that JS is in danger and sure that I am JS, then I will exhibit a certain range of dispositions: I will seek protection and cry

An example might help make this clearer.¹⁶ Suppose that we've been convinced that bitterness is relative. Where x is some food, we think that there is such a thing as *bitter_x*, the proposition that x is bitter, and we think, moreover, that this proposition is true at some occupied circumstances and false at others. Now consider a series of three cases involving a pot of soup.

First case: Abe is at a picnic. One of the items on the table is a pot of soup. Abe tastes the soup and recoils from the bitter taste. Ben walks up and asks Abe whether the soup is bitter.

Abe likes soup, but dislikes bitter soup. Abe believes that the soup is bitter—his belief is part of what explains why he does not pour himself a bowl. Upon hearing Ben's question, Abe wants to be helpful. He expects that Ben, too, like soup and dislikes bitter soup. No odd or extraordinary pragmatic considerations are in play. The natural thing for Abe to do is respond in the straightforward way, "Yes, Ben. The soup is bitter."

Second case: Abe, still at the picnic, comes to learn that the soup is flavored with phenylthiocarbamide (PTC, for short). Depending on one's genetic makeup, PTC is either bitter or tasteless. People with gene X taste PTC as bitter; people without gene X taste PTC as tasteless. Abe also learns that he's in genetically mixed company: that some of the people at the picnic have gene X, and that others do not. Cam walks up and asks whether the soup is bitter.

out for help; the circumstance I occupy will clearly be a part of the actual target of my attitude. But consider a case in which I am sure that JS is in danger, but unsure whether I am JS. In this case, it will be indeterminate whether JS is in part of the actual target of my attitude. I will be sorta-kinda inclined to tell JS that they should seek protection and sorta-kinda not: after all, I am not telling myself to seek protection or cry for out for help. Thanks to Andy Egan for discussion on this point.

¹⁶ This example is from Spencer (2016: 527-8).

There is a felt difference between the first case and the second. If given three options—assert *bitter_soup*, assert *~bitter_soup*, or refuse to make either assertion—Abe may choose the third. (I would choose the third option in his shoes.) Abe still believes that the soup is bitter—his belief is part of what explains why he (still) does not pour himself a bowl. But Abe may be unwilling to be on *Cam*'s hook with regard to the soup being bitter. Abe may minimally change the subject and say something like, “Well, it tastes bitter to me.”

Third case: Abe, still at the picnic, comes to learn that the people at the picnic are dressed according to their genetics. People with gene X are wearing shirts of one color; people without gene X are wearing shirts of another color. Dan walks up to Abe and asks whether the soup is bitter. Abe notices that Dan's shirt is the same color as his own. A bit later, Ed walks up and asks Abe whether the soup is bitter. Abe notices that Ed's shirt is a different color from his own.

In this case, Abe's interpersonal dispositions may be differential. He may be willing to assert *to Dan* that the soup is bitter, but unwilling to assert *to Ed* that the soup is bitter.

If there really is such a thing as *bitter_soup*, then, in all three cases, Abe takes some belief-like attitude toward it. But Abe's attitude differs in the three cases, and, according to the targeted view, the attitudinal difference lies in the targeting function. In the first case, Abe has an expansive attitude. The actual target of his belief includes the circumstances occupied by each of his fellow picnic goers. In the second case, Abe has a narrower attitude. The actual target of his belief includes his own circumstance, but may not include the circumstances occupied by the other picnic goers. In the third case, Abe has an intermediate

attitude. The actual target of his belief includes just the picnic goers wearing shirts of the same color as his own.¹⁷

Lots of questions about the targeted conceptions of belief and assertion remain. (One might ask about vagueness. Actual agents are messy and finite, so it's presumably almost always vague what the targeting function of a belief or assertion is. One might ask about nonhuman animals. Insofar as I am inclined toward the targeted view, I think that every belief-like attitude and every assertion-like activity, even those belonging to the simplest of creatures, have targeting functions.¹⁸ One might ask about ordinary language.¹⁹ I'm not sure whether any word in English refers to believing-for-*t* or asserting-to-*t*, although, to be frank, my confidence in the targeted view is not very sensitive to there being such a word. The targeted view is offered as a metaphysics of belief and assertion, not a semantics.) Enough has been said to get the idea on the table, however, so I want to turn to evaluation. Of these three conceptions of correctness, which should truth relativists prefer?

¹⁷ Considering other cases might be instructive. Suppose that Abe learns that people are dressed according to their genetics, but does not know whether the genetics are indicated by having shirts of the same color or trousers of the same color. The targeting function associated with Abe's belief-for attitude then will be more complicated: it will map shirt-worlds to a set of circumstances that includes all of the picnic goers with the same color shirt as his own, and it will map trouser-worlds to a set of circumstances that includes all of the picnic goers with the same color trousers as his own. The dispositional profile of this attitude is straightforward: Abe has various conditional interpersonal dispositions. Upon learning that he is in a shirt-world, he is disposed to assert that the soup is bitter to someone wearing a shirt of the same color as his own. Thanks to Andy Egan for discussion.

¹⁸ Thanks to Daniel Hoek for discussion on this point.

¹⁹ Cf. MacFarlane (2014: 157).

§4/ Plan

Building a cumulative case for the targeted view, I am going to point out various *discrepancies*: ways in which our evaluation of beliefs and assertions differ from the predictions that the simple and reflexive views make. Some of the discrepancies I'll note have been observed before. But the targeted view has been mostly ignored, and the possibility of the targeted view alters and enhances the philosophical significance of these discrepancies. So revisiting these discrepancies is worth our while.

I have divided the discussion into parts. In §5, I consider *de se* and *de nunc* beliefs and assertions, whose truth-values are sensitive to the agential and temporal coordinates of circumstances, respectively, and in §6, I consider *de gustibus* and *de possibile* beliefs and assertions, whose truth-values are sensitive to the gustatory and informational coordinates of circumstances, respectively. The division is dialectical: as we'll see, the reflexive view is more plausible when applied to the *de se* and the *de nunc*, and the simple view is more plausible when applied to the *de gustibus* and the *de possibile*. But we should not tolerate a disunified theory of correctness. What the objective epistemic ought demands should not depend on whether the beliefs and assertions have their truth-values absolutely or relatively and should not depend on the domain of discourse to which the beliefs and assertions belong. The targeted view—the view that beliefs and assertions are correct just when they are true at their actual targets—is defensible in full generality, I believe, but the simple and reflexive views are not. The accumulation of discrepancies together with the reasonable demand for a unified theory of correctness thus amounts to a powerful case in favor of the targeted view.

§5/ De Se and De Nunc

Arguments that there are *de se* and *de nunc* beliefs and assertions are various. Some belong to action theory. Ninan (2016) and Torre (2018), for example, argue that one strong reason to think that there are *de se* beliefs is that we then can preserve certain law-like generalizations between propositional attitudes and motivational profiles.²⁰ Some belong to philosophy of language. The way we talk often suggests that the contents of our assertions are *de se* or *de nunc*. If tomorrow is Alfred's 54th birthday and Burt says, "Alfred is 53 years old," it seems that Alfred can say truly, in response, "That's true (today), but false tomorrow."²¹ Some of the arguments belong to epistemology. There are philosophers who are convinced that we must countenance *de se* and *de nunc* beliefs in order to understand the dynamics of rational credence,²² and there is, of course, Lewis's famous argument from ignorance:

Consider the case of the two gods. They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. [...] Still I can imagine them to suffer ignorance: neither one knows which of the two he is. They are not exactly alike. One lives on top of the tallest mountain and throws down manna; the other lives on top of the coldest mountain and throws down thunderbolts. Neither one knows whether he lives on the tallest mountain or the coldest mountain; nor whether he throws down manna or thunderbolts. (Lewis 1979: 520)

²⁰ See Ninan (2010; 2016) and Torre (2018).

²¹ See *e.g.* Brogaard (2012) and Richard (1981).

²² Anyone who denies what Titelbaum calls "the Relevance-Limiting Thesis" (2008: 556).

If an agent is partially ignorant, then there is some proposition that is true (at the circumstance they occupy) about which they are uncertain. If Lewis is right, an agent can be partially ignorant, despite knowing every proposition that is absolutely true. So, according to Lewis, there must be some propositional remainder: some proposition whose truth-value differs between occupied circumstances.²³

I will not try to defend the claim that there are *de se* and *de nunc* beliefs and assertions.²⁴ The question I am concerned with is conditional. I want to know what the objective epistemic ought demands of *de se* and *de nunc* beliefs and assertions, supposing that there are such things. But it is important to keep the arguments above in mind because there is a coherence requirement. An account of what it is for *de se* and *de nunc* beliefs and assertions to be correct must cohere with the motivations for thinking that there are such things.

5.1 The simple view

The simple view handles the *de se* and *de nunc* poorly. There are at least five problems.

The first problem concerns supervenience. The simple view predicts that correctness of belief and assertion supervenes on content. But that prediction seems wrong. Let *me_coldest* be the *de se* proposition that I am top of the coldest mountain, and suppose that both of Lewis's gods come to believe *me_coldest*. These two beliefs have the same content, but they differ in correctness-value. The god on top of the coldest mountain correctly believes *me_coldest*, and the god on top of the tallest mountain incorrectly believes *me_coldest*.

²³ For a strengthened version of this argument from ignorance, see Shaw (forthcoming).

²⁴ For arguments that there are no *de se* or *de nunc* beliefs and assertions, see *e.g.* Boer and Lycan (1980), Cappelen and Dever (2013), and Magidor (2015).

The second problem concerns truth-correctness gaps. The simple view predicts that there are none. But that prediction seems wrong. Suppose again that both of Lewis's gods come to believe *me_coldest*, and let c_{cold} be the circumstance occupied by the god on top of the coldest mountain. Beliefs and assertions inherit their alethic properties from their propositional contents, so every belief of *me_coldest* is true-at- c_{cold} . Truth is stable: relative to any circumstance c , a proposition is true just if it's true-at- c . So, relative to c_{cold} , every belief of *me_coldest* is true. But the god on top of the tallest mountain believes *me_coldest*, and that belief is, relative to every (occupied) circumstance, incorrect. So we have a truth-correctness gap: a belief that is, relative to some (occupied) circumstance, both true and incorrect. And once we get the hang of it, we see truth-correctness gaps everywhere. Let $\sim now_t$ be the *de nunc* proposition that t is not the present moment, and let c_t be a circumstance that has t as its temporal coordinate. Many past and future agents believe and assert $\sim now_t$. Relative to c_t , those past and future beliefs and assertions are false. But it may be that every past and future belief and assertion of $\sim now_t$ is absolutely correct.

The third problem arises from the connection between incorrectness and retraction. One can retract an assertion for a variety of reasons,²⁵ but the kind of retraction that has featured prominently in philosophy is *correctness-based*: the retraction of an assertion on the grounds that it fails to satisfy the demands of the objective epistemic ought. The correctness-based retraction of an assertion is *appropriate* just if the assertion is incorrect—the correctness-based retraction of an assertion is, in that way, akin to asserting that the assertion is incorrect. Hereafter, I'll mean correctness-based retraction by 'retraction'.

²⁵ Cf. MacFarlane (2014: 109).

According to the simple view: if, relative to c , an assertion is false, then, relative to c , the assertion is appropriate to retract. But this prediction seems wrong. Let *pigs_fly* be the proposition that pigs fly, and let *~now_Thurs* be the *de nunc* proposition that it's not now Thursday. Yesterday I made two assertions: I asserted both *pigs_fly* and *~now_Thurs*. It's now Thursday. Relative to the circumstance I occupy, both assertions are false. So the simple view predicts that, relative to the circumstance I occupy, both assertions are appropriate to retract. But that's wrong. The assertion of *pigs_fly* is appropriate to retract, but the assertion of *~now_Thurs* isn't.

The fourth problem arises from the connection between correctness and disagreement. Disagreement is said in many ways. (Oysters disagree with my stomach.) But disagreement, in its central sense, is competition for correctness.²⁶ Beliefs and assertions disagree when they cannot both be correct, and agents disagree when their beliefs or assertions do.²⁷

The simple view predicts that beliefs and assertions compete for correctness whenever they compete for truth. But that prediction seems wrong. Suppose that the god on top of the coldest mountain believes *me_coldest* and that the god on top of the tallest mountain believes *~me_coldest*. These beliefs compete for truth. But they don't compete for correctness. Both beliefs are absolutely correct.

The fifth problem arises from the connection between correctness, on the one hand, and reliability and trustworthiness, on the other. Reliability and trustworthiness are measures of correctness. To a close enough approximation, an

²⁶ I say, "There are tables." Peter van Inwagen (1990) says, "There are no tables." I wonder whether we disagree, and what I want to know is whether van Inwagen's belief and assertion compete for correctness with my own.

²⁷ This conception of disagreement is familiar; see e.g. MacFarlane (2014: 123-9) and the discussion of preclusion of joint accuracy and preclusion of joint reflexive accuracy.

agent's reliability (about some matter) is the proportion of their beliefs (about that matter) that are correct, and an agent's trustworthiness (about some matter) is the proportion of their assertions (about that matter) that are correct.

The rate of falsity among *de se* and *de nunc* beliefs and assertions is high. Of course, the rate is not constant across (occupied) circumstances, since some beliefs and assertions have their truth-values relatively. But relative to every (occupied) circumstance, the rate of falsity is high. The simple view predicts that false beliefs and assertions are always incorrect. So the simple view predicts, absurdly, that agents are unreliable and untrustworthy when it comes to *de se* and *de nunc* matters.

Take the days of the week, for example. People on Monday believe and assert *now_Mon*, the *de nunc* proposition that it's Monday. People on Tuesday believe and assert *now_Tues*—and so on, similarly, for the other five days of the week. Relative to any (occupied) circumstance, only 1/7 of these beliefs and assertions are true. The simple view thus predicts, absurdly, that people are only 1/7 reliable (and only 1/7 trustworthy) about which day it is.²⁸

We should not be willing to accept a theory that predicts that we are unreliable and untrustworthy about *de se* and *de nunc* matters.

²⁸ Suppose it's Tuesday, and suppose that *M*, a person on Monday, believes and asserts *now_Mon*. Are *M*'s belief and assertion true? My informants have been all over the map in their responses. I suspect the variance is due to the polysemy of 'true' in English. (See Khoo and Phillips (2019).)

One could introduce a notion of truth that made the simple view true. If we say that a belief or assertion is *normatively true* just if it's correct, then beliefs and assertions are correct just when they're normative true. But this theory sheds no light on correctness. The interesting simple view is the one formulated using the notion of truth that we have some antecedent grasp of, the notion that is constitutively tied to meaning, modality, and knowledge. And *that* notion of truth, I claim, satisfies both Inheritance and Stability.

5.2 The choice between the reflexive view and the targeted view

Once the simple view is set aside, the reflexive view might seem the obvious choice. After all, as we've seen, there's a strong correlation between reflexive truth and correctness among *de se* and *de nunc* beliefs and assertions, and the reflexive view is well-positioned to explain this correlation.

But this correlation also can be explained by the targeted view. The targeted view says that beliefs and assertions are correct just if they are true at their actual targets, and it places no constraints on actual targets: it says that any set of actual circumstances can be the actual target of a belief or assertion of any proposition. The targeted view, taken by itself, then, does not predict any strong correlation between reflexive truth and correctness among *de se* and *de nunc* beliefs and assertions. But if, as a matter of fact, the actual targets of *de se* and *de nunc* beliefs and assertions are narrow, including few circumstances, if any, besides the one occupied by the agents to which they belong, then the targeted view *does* predict the correlation we see between reflexive truth and correctness. And a strong case can be made for the requisite narrowness.

If I believe *now_t*, I am disposed to assert *now_t* to co-present agents. But I am not disposed to assert *now_t* to past or future agents. If someone lost in time called me on the phone and asked whether *t* is present, I would be like Abe is in regard to Cam: unwilling to take a stand and looking for some way to minimally change the subject. Similarly, if I believe *now_t*, I am disposed to disagree with co-present agents who believe \sim *now_t*, and I take co-present agents who believe \sim *now_t* to be mistaken. But I do not take myself to disagree with past or future agents who believe \sim *now_t*, and I do not take past or future agents who believe \sim *now_t* to be mistaken. The actual target of my belief thus appears to be narrow: it appears to include just those actual circumstances that have *t* as their temporal coordinate. And there are cases where the actual target is even narrower. I believe

me_JS, the *de se* proposition that I am JS. But I do not take myself to disagree with anyone who believes $\sim me_JS$. I am not sure whether *me_JS* can be asserted by uttering an English sentence—I do not think it can be asserted by uttering, “I am JS.” But even if there were some easy way to assert *me_JS*, I wouldn’t—unless I was soliloquizing or trying to mislead someone who was confused about their identity.

Since the reflexive view and the targeted view both predict a strong correlation between reflexive truth and correctness, we need some other way to choose between them, and I think that the best way forward is to look at the *divergent cases*: the cases in which the views make rival predictions about correctness. There are, in principle, two sorts.

The first sort is one in which a belief or assertion is reflexively true but, according to the targeted view, incorrect, on account of being unduly wide. For example, suppose that I believe *now_t* differently: that I am disposed to assert *now_t* to anyone, past, present, or future, and that I take myself to disagree with anyone who believes $\sim now_t$, irrespective of where they are in time. The targeted view then predicts that the actual target of my belief is wide, perhaps including every actual circumstance. The belief is reflexively true, so the reflexive view predicts that the belief is correct. But the belief is not true at its actual target, so the targeted view predicts that it’s incorrect.

A similar thing happens with assertion. Suppose that it’s Tuesday, and suppose that I assert *now_Tues* twice, once to a Tuesday addressee and once to a Wednesday addressee. Both assertions are reflexively true, so the reflexive view predicts that both are correct. The targeted view says that the assertion to the Tuesday addressee is correct, but that the assertion to the Wednesday addressee is incorrect.

I think the targeted view gets it right in both cases. My (wide) belief of *now_t* is reflexively true, but it seems criticizable along an objective epistemic dimension. And I think the targeted view is right to draw the distinction it does between my two assertions of *now_Tues*. My assertion to the Tuesday addressee is correct: it's inappropriate to retract and counts in favor of my overall trustworthiness. But my assertion to the Wednesday addressee is incorrect: it's appropriate to retract and counts against my overall trustworthiness.

The second sort of divergent case is one in which a belief or assertion is reflexively false but, according to the targeted view, correct, nevertheless. Say that a belief or assertion is *autocentric* if its actual target includes the speaker's circumstance, and *exocentric*, otherwise.²⁹ The second sort of divergent case arises only for exocentric beliefs and assertions. A proponent of the targeted view is not forced to think that exocentric beliefs and assertions are possible, but the case for exocentric beliefs and assertions is, I think, rather strong. Suppose that I work for the Department of Temporal Location. I get a call on Tuesday from someone lost in time, who, judging by their voice, is in a terrible hurry. My caller ID reveals that the person on the other end of the line is calling from Thursday. "What day is it?" they ask. "It's Thursday," I say, thereby asserting *now_Thurs*.³⁰ They thank me and hang up. My assertion is reflexively false, but the assertion seems to be absolutely correct, just as the targeted view predicts. (The assertion seems to count in favor of my overall trustworthiness, for example.) And the same goes for belief. If we accept exocentric assertions, then it's natural to accept exocentric beliefs as the attitudes canonically expressed by exocentric assertions. If I exocentrically

²⁹ Cf. Lasersohn (2005; 2009).

³⁰ The claim that *now_Thurs* is the proposition expressed is bolstered by an eavesdropper variant. My colleague, who is also confused about what day it is, overhears me say, "It's Thursday." He asks me, "Is that true?" In response, it's natural for me to say 'no'.

believe *now_Thurs*, then my belief is reflexively false. But that belief, like my exocentric assertion that expresses it, seems absolutely correct, despite being reflexively false, as the targeted view predicts.

5.3 Correctness *de se* and *de nunc*

The simple view handles *de se* and *de nunc* beliefs and assertions poorly. It wrongly predicts that correctness supervenes on content; it makes false predictions about appropriate retraction and disagreement; and it predicts that people are unreliable and untrustworthy when it comes to *de se* and *de nunc* matters. This poor performance is due to the correctness relativity. *Pace* the simple view, it seems that *de se* and *de nunc* beliefs and assertions have their correctness-values absolutely.

The reflexive view performs better; for the most part, *de se* and *de nunc* beliefs and assertions are correct just when they are reflexively true. But reflexive truth is neither necessary nor sufficient for correctness. Unduly wide *de se* and *de nunc* beliefs and assertions can be both reflexively true and incorrect, and exocentric *de se* and *de nunc* beliefs and assertions can be both reflexively false and correct. *Pace* the reflexive view, it seems that *de se* and *de nunc* beliefs and assertions sometimes aim at something other than reflexive truth.

The view that performs best is the targeted view, which ensures that *de se* and *de nunc* beliefs and assertions have their correctness-values absolutely, while allowing *de se* and *de nunc* beliefs and assertions to aim at any collection of actual circumstances, not just the agent's own.

§6/ *De Gustibus* and *De Possibile*

Some think that the only beliefs and assertions that have their truth-values relatively are *de se* and *de nunc*, but others disagree. There are domains, like

matters of taste and epistemic possibility, that “fall short of being fully objective,” as MacFarlane (2014: v) puts it, and some think that beliefs and assertions associated with these less-than-fully objective domains have their truth-values relatively, as well.³¹

I will not try to defend the claim that there are *de gustibus* and *de possibile* beliefs and assertions. My question is again a conditional one. I want to know what the objective epistemic ought demands of *de gustibus* and *de possibile* beliefs and assertions, supposing that there are such things. But the coherence requirement mentioned above applies here, too. An account of what it is for *de gustibus* and *de possibile* beliefs and assertions to be correct must cohere with the motivations for thinking that there are such things. So I want to start with *the linguistic argument*—the main motivation for thinking that there are *de gustibus* and *de possibile* beliefs and assertions—following MacFarlane’s (2014) presentation fairly closely.

6.1 The linguistic argument

Let’s start with ‘tasty’. An *absolute* account of ‘tasty’ claims that beliefs and assertions associated with ‘tasty’ always have their truth-values absolutely. The linguistic argument puts forward two desiderata for an account of ‘tasty’ and alleges that no absolute account satisfies both.

The argument begins with an observation. According to MacFarlane, our use of ‘tasty’ adheres to the following solipsistic rule:

³¹ See *e.g.* Brogaard (2008; 2010), Egan (2007; 2010); Egan, Hawthorne, and Weatherson (2005), Köbel (2002), Lasersohn, (2005), and MacFarlane (2007a; 2011; 2014).

Pleasure Rule. If you know first-hand how something tastes, call it ‘tasty’ just in case its flavor is pleasing to you, and ‘not tasty’ just in case its flavor is not pleasing to you. (MacFarlane 2014: 4)

As evidence that our use of ‘tasty’ adheres to the pleasure rule, MacFarlane notes that speeches that offend against the pleasure rule sound odd. Consider:

- (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. (MacFarlane 2014: 4)

Assuming that our use of ‘tasty’ adheres to the pleasure rule, we have our first desideratum, *the reasonability desideratum*. An account of ‘tasty’ should render our adherence to the pleasure rule reasonable.

Not every theory satisfies the reasonability desideratum. Consider *objectivism*, the view that there is some objective property of foods invariably denoted by ‘tasty’.

Objectivism renders the pleasure rule unreliable, as a toy case brings out. Suppose that there are ten people in a gymnasium. All ten taste some alien food, *z*. Five of the people—*the Yum-people*—find the taste of *z* pleasing and, adhering to the pleasure rule, sincerely say, “Food *z* is tasty.” The other five—*the Yuk-people*—find the taste of *z* displeasing and, adhering to the pleasure rule, sincerely say, “Food *z* is not tasty.” If objectivism is true, then the Yum-people believe and assert *tasty_z*, the proposition that food *z* is tasty, and the Yuk-people believe and assert \sim *tasty_z*, the negation of *tasty_z*. If objectivism is true, then one of these propositions—either *tasty_z* or \sim *tasty_z*—is absolutely true. So, if objectivism is

true, then, in this toy case, the pleasure rule proves to be just 50% reliable. And the lesson of this toy case generalizes. All of us know that there are vast differences in the foods people find pleasing, so all of us are in a position to know that objectivism entails that the pleasure rule is unreliable.

If objectivism is true, a competent speaker should be in a position to know as much. And if our use of 'tasty' adheres to the pleasure rule, a competent speaker should be in a position to know that, too. So, if objectivism is true, a competent speaker should be in a position to know that their use of 'tasty' adheres to an unreliable rule. But that threatens to render their adherence to the pleasure rule unreasonable.

The problem takes the form a dilemma. Let's say that the *reliability* of a gustatory propensity is the rate of correctness among the beliefs and assertions it produces in accordance with the pleasure rule, and let's say that an agent is *chauvinistic* if the agent thinks that their gustatory propensity is considerably more reliable than the pleasure rule. Either an agent is chauvinistic or not.

If the agent is not chauvinistic, then their adherence to the pleasure rule is unreasonable. The agent knows (or anyway is in a position to know) that the pleasure rule is unreliable. So, if the agent is not chauvinistic, then the agent cannot coherently think that the beliefs and assertions they have formed in accordance with the pleasure rule are probably correct. But if an agent cannot coherently think that their beliefs and assertions are probably correct, then the agent is not reasonable.

However, on the other horn, chauvinism, itself, seems unreasonable.

One way to argue that chauvinism is unreasonable is by pointing out that it requires that agents think that they are epistemically lucky. As MacFarlane says:

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 [the pleasure rule] 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. (MacFarlane 2014: 5)

A second way to argue that chauvinism is unreasonable is by noting that it leads to foreseeable exploitation. Consider an agent who knows that their gustatory propensity will change—the agent hates prunes, let's say, but knows that they'll love them. The agent adheres to the pleasure rule now, chauvinistically thinking that the gustatory propensity they now have is more reliable than the gustatory propensity they'll later have, and the agent knows that they will adhere to the pleasure rule later, chauvinistically thinking that the gustatory propensity they then have is more reliable than the gustatory propensity they now have. Agents should be willing to put their money where their judgments are, so let's suppose that we play a game of IOU. When the agent makes an assertion concerning tastiness, we owe them \$1 if their assertion is correct, and they owe us \$2 if their assertion is incorrect. Today the agent chauvinistically asserts that prunes are not tasty, and later, as they knew they would, they chauvinistically assert that prunes are tasty. Whether prunes are tasty or not, the agent, on account of their chauvinism, owes us \$1.

If it is reasonable for an agent to adhere to the pleasure rule, then it is reasonable for the agent to play our game of IOU. So, if it's reasonable for agents

to be chauvinistic, then it's reasonable for agents to take this sure-loss sequence of IOU-wagers. But taking this sure-loss sequence of IOU-wagers seems clearly unreasonable.

That, then, complete the dilemma. If an agent is in a position to know that that the pleasure rule is unreliable, then it appears to be unreasonable for the agent to adhere to the pleasure rule, whether the agent is chauvinistic or not. So objectivism, it appears, does not satisfy the reasonability desideratum.

There are absolute accounts of 'tasty' that satisfies the reasonability constraint. One such is *solipsistic contextualism*: the view that, when *A* says, "Food *x* is tasty," *A* asserts the proposition that food *x* is tasty to *A*, a proposition that is true just if *A* finds the flavor of *x* pleasing. An agent's gustatory reaction to a food is a good guide to whether the agent finds the food pleasing, so solipsistic contextualism renders the pleasure rule reliable. If solipsistic contextualism is true, then almost every belief and assertion formed in accordance with the pleasure rule is absolutely correct. The threat of unreasonableness thus never arises.

But this brings us to the second desideratum, *the incorrectness desideratum*. Our patterns of disagreement and retraction suggest that the rate of incorrectness among the beliefs and assertions formed in accordance with the pleasure rule is high and much higher than solipsistic contextualism predicts.

We do not often disagree about which foods are tasty to which people. The following dialogue sounds terrible:

Alice: Licorice is tasty to me.
Bob: # No. / I disagree. / That's false.

But we often disagree about which foods are tasty. A dialogue like this one sounds fine and familiar:

Alice: Licorice is tasty.

Bob: No. / I disagree. / That's false. (*cf.* MacFarlane 2014: 9)

Disagreement, in its central sense, is competition for correctness. If solipsistic contextualism is true, beliefs and assertions formed in accordance with the pleasure rule almost never compete for correctness. (They're almost always absolutely correct.) So solipsistic contextualism predicts less disagreement than there appears to be.

A similar problem arises for appropriate retraction. MacFarlane gives the examples of fish sticks:

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 are tasty," I would retract the earlier assertion. I wouldn't say, "They were tasty then, but they aren't tasty anymore," 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. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (2014: 13-4, emphasis original)

Retracting correct assertions is inappropriate. Solipsistic contextualism predicts that almost every assertion formed in accordance with the pleasure rule is absolutely correct. So solipsistic contextualism predicts that it's almost never appropriate to retract an assertion formed in accordance with the pleasure rule. But

that seems wrong. Even if fish sticks were tasty to MacFarlane as a kid, it seems appropriate for him to retract his assertion now that his tastes have changed.

Objectivism and solipsistic contextualism are not the only absolute accounts of ‘tasty’. There are moderate forms of contextualism, for example, which say that ‘tasty’ claims are relativized, not to specific individuals, but instead to groups. Some of these moderate forms of contextualism satisfy the incorrectness desideratum. But, as MacFarlane notes, moderate forms of contextualism that satisfy the incorrectness desideratum struggle with the reasonability desideratum:

If claims about what is “tasty” are often claims about what is tasty to a group, then it is unclear why finding the food pleasing oneself should be sufficient warrant for claiming that it is “tasty.” There might, admittedly, be some cases where one has good grounds for thinking that others’ tastes are relevantly similar to one’s own. But since that is often not the case, we should expect there to be many cases where [the pleasure rule] fails, and sentences like (1)-(3) sound okay. I submit that there are not. Even when I know that some of the people I am talking to do not like the taste of orange juice, it would be bizarre for me to deny that orange juice (which I like a lot) is tasty, or to express skepticism about whether it is tasty. (MacFarlane 2014: 12)

We thus arrive at the conclusion of the linguistic argument: that no absolute account of ‘tasty’ satisfies both the reasonability and incorrectness desiderata.

And it’s easy to see how the argument can be extended to ‘might’. First, we identify a solipsistic rule to which our use of ‘might’ adheres. Perhaps:

Information Rule. Apply ‘might’ to p just if p is consistent with your information, and apply ‘not might’ to p just if p is inconsistent with your information.

Supposing that our use of ‘might’ adheres to the information rule, the first desideratum is to render our adherence to the information rule reasonable. One theory that satisfies the reasonability desideratum is solipsistic contextualism: the view that, when A says, “It might be that p ,” A asserts that p is consistent with A ’s information. But solipsistic contextualism struggles with the incorrectness desideratum; for our patterns of disagreement and retraction suggest that the rate of correctness among the beliefs and assertions formed in accordance with the information rule is higher than solipsistic contextualism predicts. Consider the following vignette:

Sally: Joe might be in China.
George: No, he can’t be in China. He doesn’t have his visa yet.
Sally: Oh, really. Then I guess I was wrong.
(MacFarlane 2014: 240)

As MacFarlane says, *pace* solipsistic contextualism,

It seems that George is contradicting Sally and rejecting her claim. It also seems that, having learned something from George, Sally concedes that she was wrong. Finally, it seems appropriate for her to retract her original claim, rather than continuing to stand by it. (MacFarlane 2014: 240)

There are other absolute accounts of ‘might’, of course. But, according to the linguistic argument, none of them satisfies both the reasonability and incorrectness desiderata.

6.2 Does the linguistic argument motivate a relativist account of ‘tasty’?

Even if we are convinced that the linguistic argument refutes absolute accounts of ‘tasty’ and ‘might’, we might wonder whether it motivates relativistic accounts. Where x is some food, let’s suppose that there is such a thing as *tasty_x*, the proposition that food x is tasty, and let’s suppose that this proposition is true at a circumstance just if the agent who occupies the circumstance finds the flavor of x pleasing. Similarly, where p is a proposition, let’s suppose that there is such a thing as *might_p*, the proposition that p might be the case, and let’s suppose that this proposition is true at a circumstance just if p is compatible with the information of the agent who occupies the circumstance. Our question is this: can we append a conception of correctness to these relativistic accounts of ‘tasty’ and ‘might’ and thereby satisfy both the reasonability and incorrectness desiderata? I’ll consider our three options—the reflexive view, the targeted view, and the simple view—taking them in that order.

6.3 The reflexive view

The reflexive view is not an attractive option. It faces two main problems.

The first problem is exocentrism. Some *de possibile* assertions—like the following example, from Egan, Hawthorne, and Weatherson (2005: 140)—are exocentric:

Ann is planning a surprise party for Bill. Unfortunately, Chris has discovered the surprise and told Bill all about it. Now Bill and Chris are

having fun watching Ann try to set up the party without being discovered. Currently Ann is walking past Chris' apartment carrying a large supply of party hats. She sees a bus on which Bill frequently rides home, so she jumps into some nearby bushes to avoid being spotted. Bill, watching from Chris' window, is quite amused, but Chris is puzzled and asks Bill why Ann is hiding in the bushes. Bill says, "I might be on that bus."

There are also exocentric *de gustibus* assertions. Of course, most *de gustibus* assertions are autocentric. Suppose that Sara and Clara are roommates. Sara tastes both Alpo and Purina, and says, with some surprise, "Alpo is tastier than Purina." Sara's assertion is autocentric, as evidenced by the fact that it would be infelicitous for her to say that she thinks the opposite of what she said:

Sara: Alpo is tastier than Purina.
Clara: Do you really think that?
Sara: # No. I actually think that Purina is tastier.

But some *de gustibus* beliefs and assertions are exocentric. Suppose that Justin and Dustin have a dog, Fido. Justin says, "I'll get Fido some Purina at the grocery." Dustin, looking Fido, says, "Alpo is tastier than Purina. (Let's get it. It's worth the extra five bucks.)" Dustin's assertion is exocentric, as evidenced by the fact that it would be felicitous for him to say that he thinks the opposite of what he just said. Suppose that Dustin has tasted both Alpo and Purina and prefers the taste of Purina. In such a context, the following dialogue sounds fine:

Dustin: Alpo is tastier than Purina.
Justin: Do you really think that?

Dustin: No. I actually think that Purina is tastier. (But what I like is irrelevant. Let's get the Alpo.)

The reflexive view predicts that reflexively false beliefs and assertions are always absolutely incorrect. But that prediction seems wrong. It seems that exocentric beliefs and assertions, like Bill's and Dustin's, can be both reflexively false and correct.

The second problem, which is of greater moment, concerns incorrectness. Like solipsistic contextualism, the reflexive view predicts too little incorrectness.

The reflexive view and solipsistic contextualism are not equivalent. They disagree both about identity and about truth-value, and I think that both disagreements favor the reflexive view.

Take the disagreement about identity. There seems to be widespread same-saying and same-thinking when it comes to tastiness. Lots of people think that licorice is tasty. Most people agree that chocolate is tasty. The reflexive view is easy to reconcile with widespread same-saying and same-thinking when it comes to tastiness, but solipsistic contextualism is not. According to solipsistic contextualism, what I believe and assert when I say, "Licorice is tasty," is distinct from what you believe and assert when you say, "Licorice is tasty." And that's odd. It's odd to think that I was the first person ever to have said what I said when I first said, "Licorice is tasty."

The reflexive view also has an advantage with regard to truth-value. We often ascribe falsity to beliefs and assertions formed in accordance with the pleasure rule. Upon *A* saying, "Licorice is tasty," *B*, who finds licorice displeasing, might say, "No. That's false." These ascriptions of falsity are hard to reconcile with solipsistic contextualism, but they accord nicely with the reflexive view. According to the reflexive view, although *A*'s assertion is reflexively true and

hence absolutely correct, A 's assertion is false at the circumstance B occupies, so B 's claim that A 's assertion is false is reflexively true and hence (also) absolutely correct.

But the reflexive view and solipsistic contextualism both predict too much correctness. In fact, they give rise to almost the same pattern of correctness. Take the case of 'tasty'. According to solipsistic contextualism, A believes and asserts the proposition that licorice is tasty to A , and A 's belief and assertion are absolutely correct just if that proposition is absolutely true. According to the reflexive view, A believes and asserts *tasty_licorice*, and A 's belief and assertion are absolutely correct just if that proposition is true at the circumstance A occupies. But the proposition that licorice is tasty to A is absolutely true just if *tasty_licorice* is true at the circumstance A occupies, so the two views give rise to the same pattern of correctness. And the same goes for 'might': the proposition that p is consistent with A 's information is absolutely true just if *might_p* is true at the circumstance A occupies. The two views thus have the same shortcoming: they predict too much (absolute) correctness, and hence too little disagreement and too little appropriate retraction.

6.4 The targeted view

The targeted view avoids the problems that beset the reflexive view. The targeted view better handles exocentrism, allowing that exocentric beliefs and assertions can be both correct and reflexively false. But the more important advantage concerns incorrectness. The targeted view satisfies the reasonability desideratum in a different way than the reflexive view does, and as a result it also satisfies the incorrectness desideratum.

According to our relativistic account of 'tasty', *tasty_x* is true at the circumstance an agent occupies just if the agent finds the flavor of x pleasing.

Thus, relative to the circumstance an agent occupies, the proposition that the agent finds the flavor of x pleasing entails $tasty_x$. That's what explains why speeches that offend against the pleasure rule sound so odd. An agent who is sure that they find the flavor of x pleasing but unsure whether x is tasty exhibits semantic incompetence, failing to appreciate the relativistic entailment relations between those contents.

Our relativistic account of 'tasty' predicts that almost every belief and assertion formed in accordance with the pleasure rule is reflexively true. The reflexive view predicts too much correctness because it says that beliefs and assertions are absolutely correct whenever they are reflexively true. The targeted view takes a different course.

According to the targeted view, there are many autocentric belief-for attitudes one could take toward $tasty_x$, and there are likewise many autocentric assertion-to activities. The fact that an agent finds the flavor of x pleasing ensures that the agent can correctly take some (narrow) autocentric belief-for attitude toward $tasty_x$, and likewise ensures that the agent can correctly assert $tasty_x$ to some (narrow) actual target. To satisfy the reasonability desideratum, a theory must predict that an agent is in a position to autocentrically believe and assert $tasty_x$ upon finding the flavor of x pleasing, and the targeted view does so. As MacFarlane says,

Even when I know that some of the people I am talking to do not like the taste of orange juice, it would be bizarre for me to deny that orange juice (which I like a lot) is tasty, or to express skepticism about whether it is tasty. (2014: 12)

The targeted view explains why that would be bizarre. If an agent knows that they find the flavor of x pleasing, then the agent is in a position to know that they cannot autocentrically believe or assert $\sim\textit{tasty}_x$ correctly. But the targeted view does not predict that the rate of correctness among the beliefs and assertions formed in accordance with the pleasure rule is high; for those beliefs and assertions could be incorrect, on account of being unduly wide, even if they are reflexively true.

For example, Alice finds the flavor of licorice pleasing and says to Bob, “Licorice is tasty.” The actual target of her assertion includes Bob’s circumstance. If Bob finds the flavor of licorice displeasing and takes some autocentric belief-for attitude toward $\sim\textit{tasty}_{\textit{licorice}}$, as we may suppose that he does, then Alice’s assertion competes for correctness with Bob’s belief. It’d be right for Bob to say,

“No. / That’s false. / I disagree. / That’s mistaken.”

After all, according to the targeted view: Alice’s assertion *is* false at Bob’s circumstance; Alice and Bob, in fact, *do* disagree; and Alice’s assertion *is* mistaken, i.e., incorrect, despite being reflexively true, since it is not true throughout its actual target.

The same goes for appropriate retraction. Return to MacFarlane’s example and focus on the passage I’ve underlined for emphasis:

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. So, if someone said, “But you said years ago that fish sticks were tasty,” I would retract that assertion. [...] I wouldn’t say,” “When I said that, I only meant that they were tasty to me

then.” I *didn't* mean that. At the time I took myself to be disagreeing with adults who claimed that fish sticks weren't tasty. (MacFarlane 2014: 14)

According to the targeted view, kid-MacFarlane's assertion is incorrect, despite being reflexively true, since it is not true throughout its actual target.

The targeted view forces agents to strike a balance. On the one hand, there are advantages of beliefs and assertions with wider actual targets. If I want to coordinate with you vis-à-vis tastiness, the actual target of my beliefs and assertions must include your circumstance. But beliefs and assertions with wider actual targets are at greater risk of incorrectness, for the objective epistemic ought demands more of them.

There is an important point of agreement between the simple and reflexive views. Both say that an agent's belief or assertion is correct at the agent's circumstance if the belief or assertion is true at the agent's circumstance. But the targeted view does not share in this agreement. According to the targeted view, what the objective epistemic ought demands of beliefs and assertions depends on what the actual targets of those beliefs and assertions are. Thus, according to the targeted view, an agent's belief and assertions might be incorrect at the agent's circumstance even if they are true at the agent's circumstance.

6.5 The simple view

The simple view might appear to satisfy both the reasonability and incorrectness desiderata, and to do without forcing us to abandon the familiar, untargeted conceptions of belief and assertion. Of course, the simple view entails the staggering normative thesis that beliefs and assertions sometimes have their correctness-values relatively. But, at first blush, that correctness relativity might appear to be an advantage. Take the case of 'tasty', again. The simple view

ensures that almost all of the beliefs and assertions formed in accordance with the pleasure rule are correct relative to the circumstance occupied by the agent at the time of formation. But the simple view also predicts lots of disagreement and relatively appropriate retraction. It predicts that whenever one agent believes or asserts *tasty_x* and another believes or asserts \sim *tasty_x*, the two agents disagree. And it predicts that just about every retraction of an assertion concerning tastiness is correct relative to the circumstance in which it is retracted. For example, it predicts that, relative to the circumstance that adult-MacFarlane occupies, the retraction of kid-MacFarlane's assertion is appropriate.

I'm convinced, however, that the simple view does not adequately handle *de gustibus* and *de possibile* beliefs and assertions. I'll briefly mention four problems it faces, and then I'll argue that, contrary to initial appearances, the simple view does not satisfy the reasonability desideratum.

The first problem concerns exocentrism. Recall Bill's exocentric assertion that he might be on the bus. According to the simple view, relative to the circumstance Bill occupies, the objective epistemic ought demands that Bill's (exocentric) assertion be true at the circumstance Bill occupies. But that seems wrong. Bill's assertion seems to be correct at every (occupied) circumstance, and it seems to be so on account of it being consistent with Ann's information that Bill is on the bus.³² Bill's information seems absolutely irrelevant.

The second problem concerns mistakenness. According to the simple view, if *p* is false-at-*c*, then, relative to *c*, every belief and assertion of *p* is mistaken. But that seems wrong. When my niece says that dirt is tasty, I judge that she's mistaken. I say, "Stop eating dirt, Finley. Dirt is not tasty. It's gross and not to be eaten." But when worm-like aliens, who eat dirt as the main staple in their diet,

³² That exocentrism makes trouble the simple view is familiar; *cf.* MacFarlane (2014: 155-6).

say that dirt is tasty, I do not judge that they are mistaken. Indeed, if a hungry alien asked me where it could find some tasty food, I might (exocentrically) point them to the dirt pile in my back yard.³³

The same problem arises for epistemic possibility. Mary is at the store already. I know that she knows whether the store is open, so I know that she either believes *~might_open* or *~might_closed*. I am unsure whether the store is open, so I believe both *might_open* and *might_closed*. According to the simple view, Mary is, relative to the circumstance I occupy, mistaken: either she knows that the store is open, in which case, relative to the circumstance I occupy, she mistakenly believes *~might_closed*, or she knows that the store is closed, in which case, relative to the circumstance I occupy, she mistakenly believes *~might_open*. But that prediction seems wrong. Knowing that Mary has some belief that's false at the circumstance I occupy does not incline me to regard that belief of hers as mistaken (relative to my circumstance). What's consistent with my information seems absolutely irrelevant to what the objective epistemic ought demands of Mary's beliefs.³⁴

The third problem concerns disagreement. The simple view predicts that any competition for truth is a competition for correctness. But that prediction seems wrong. I take myself to disagree with Finley, but I do not take myself to disagree with the aliens, even though the aliens believe and assert the very proposition that Finley does. There are disagreements about what might be the case. But I do not take myself to disagree with Mary, even though I know that she believes the negation of something I believe.

³³ Cf. Lasersohn (2005; 2009), Richard (2015), and Spencer (2016).

³⁴ This point about mistakenness connects to Dietz (2008) and draws on the same sorts of symmetric uncertainty that features in Ross and Schroeder (2013). Also see MacFarlane (2014: 260-1).

The fourth problem concerns retraction. According to the simple view, if A previously asserted p and p is false at the circumstance A presently occupies, then, relative to the circumstance A presently occupies, it's appropriate for A to retract the previous assertion, as opposed to standing by it. But there are cases, like the following one, from von Fintel and Gillies, that seem to run contrary to this prediction:

- Alex: The keys might be in the drawer.
Billy: (*Looks in the drawer, agitated.*) They're not. Why did you say that?
Alex: Look, I didn't say that there *were* in the drawer. I said they *might* be there—and they might have been. Sheesh.
(von Fintel and Gillies 2008: 258)

Let c_{after} be the circumstance Alex occupies after the keys have been found not to be in the drawer. If, relative to c_{after} , it is appropriate for Alex to retract his assertion, then, relative to c_{after} , it is not appropriate for Alex to stand by his assertion. But, relative to c_{after} , it seems appropriate for Alex to stand by his assertion.³⁵

The targeted view avoids all four of these problems. (1) It avoids the exocentrism problem because it allows exocentric beliefs and assertions to be both absolutely correct and reflexively false. (2) It avoids the mistakenness problem because it allows correctness to vary independently of content. According to the targeted view: Finley might incorrectly (and hence mistakenly) believe and assert the very proposition that the alien correctly (and hence unmistakably) believes and

³⁵ Also see Marques (2014; 2018).

asserts; Mary might correctly (and hence unmistakably) believe and assert *~might_closed*, even though I correctly (and hence unmistakably) believe and assert *might_closed*. (3) It avoids the disagreement problem because it allows for cases of differential disagreement: that is, cases in which *A* and *B* both believe or assert *p*, *C* believes or asserts *~p*, *A* and *C* disagree, but *A* and *B* do not.³⁶ And (4), it avoids the retraction problem. The targeted view predicts that some ‘might’ claims formed in accordance with the information rule are incorrect and appropriate to retract. Recall this vignette:

Sally: Joe might be in China.
George: No, he can’t be in China. He doesn’t have his visa yet.
Sally: Oh, really. Then I guess I was wrong.
(MacFarlane 2014: 240)

If the actual target of Sally’s assertion includes George’s circumstance, as it almost certainly does, then, according to the targeted view, Sally’s assertion is incorrect and hence appropriate to retract, despite being reflexively true. But consider von Fintel and Gillies’ vignette again:

Alex: The keys might be in the drawer.
Billy: (*Looks in the drawer, agitated.*) They’re not. Why did you say that?
Alex: Look, I didn’t say that there *were* in the drawer. I said they *might* be there—and they might have been. Sheesh.
(von Fintel and Gillies 2008: 258)

³⁶ Cf. Spencer (2016).

The actual target of Alex's assertion almost certainly includes the circumstance Billy occupies before looking into the drawer. But if the purpose of Alex's assertion is to brainstorm places to look for the keys, then its actual target probably doesn't include the circumstances that Billy and he occupy after looking in the drawer. And if Alex's assertion does not include the post-looking circumstances, then, according to the targeted view, it's appropriate for Alex to stand by his assertion, as opposed to retracting it.

Having mentioned these four problems, however, I want to turn to the reasonability desideratum; for I think the main problem facing the simple view is its failure to satisfy the reasonability desideratum.³⁷

The problem starts with reliability. The simple view renders the pleasure rule and the information rule unreliable.

To see the unreliability of the pleasure rule, return to the gymnasium. If the simple view is true, then the Yum-people believe and assert *tasty_z*, the Yuk-people believe and assert \sim *tasty_z*, the negation of *tasty_z*, and all of these beliefs and assertions have their correctness-values relatively. Relative to a circumstance occupied by a Yum-person, the Yum-people believe and assert correctly, and the Yuk-people believe and assert incorrectly. Relative to a circumstance occupied by a Yuk-person, the Yum-people believe and assert incorrectly, and the Yuk-people

³⁷ MacFarlane (2014: ch. 12) considers an argument—akin to the one put forward by Ross and Schroeder (2013)—that his form of truth relativism is irrational, on account of violating Reflection. My argument that the simple view fails to satisfy the reasonability desideratum is not primarily concerned with Reflection. It's couched in terms of reliability and foreseeable exploitability, which, in my view, makes the argument considerably stronger. But it's an argument with a similar spirit, and it purports to have the same ramifications. Also see Evans (1985) and Marques (2014).

believe and assert correctly. But relative to every (occupied) circumstance, half of these beliefs and assertions are incorrect. So, if the simple is true, then, relative to every occupied circumstance, the pleasure rule is, in this toy case, just 50% reliable. And the lesson of this toy case generalizes. There are widespread differences in the foods different people find pleasing. So, if the simple view is true, then people are, on the whole, very unreliable in their judgments of tastiness.

When we turn to the information rule, things get even worse. Suppose that there are one hundred fair die. Let *might_n:x* be the proposition that die *n* might land *x*, and let *c_n* be the circumstance the agent occupies just prior to rolling die *n*. At *c₁*, the agent, adhering to the information rule, believes and asserts six propositions: *might₁:1*, *might₁:2*, *might₁:3*, *might₁:4*, *might₁:5*, and *might₁:6*. Relative to *c₁*, all six propositions are true, so the agent's reliability with regard to how the fair die might land is 100%. The agent rolls die 1 and sees it land (say) six. At *c₂*, the agent, adhering to the information rule, believes and asserts six propositions: *might₂:1*, *might₂:2*, *might₂:3*, *might₂:4*, *might₂:5*, and *might₂:6*. Relative to *c₂*, all six of the agent's beliefs and assertions concerning how die 2 might land are true and correct, but five of the agent's beliefs and assertions concerning how die 1 might land are false and incorrect. So, relative to *c₂*, the agent's reliability with regard to how the fair die might land is 7/12. The agent then rolls die two and sees it land (say) one. At *c₃*, the agent, adhering to the information rule, believes and asserts six propositions: *might₃:1*, *might₃:2*, *might₃:3*, *might₃:4*, *might₃:5*, and *might₃:6*. Relative to *c₃*, all six of the agent's beliefs and assertions concerning how die 3 might land are true and correct, but ten of the agent's beliefs and assertions concerning how die 1 and die 2 might land are false and incorrect. So, relative to *c₃*, the agent's reliability with regard to how the fair die might land is 8/18. And this pattern continues: as the agent rolls ever more of the fair die, the agent's reliability approaches 1/6. And

that, I think, is absurd. The claim that people heretofore have been very unreliable about which outcomes of chance processes might obtain seems obviously false.

If we say that a rule is *invariably* unreliable if it is unreliable relative to every (occupied) circumstance, then the reliability problem can be summed up succinctly: the simple view renders the pleasure rule and the information rule invariably unreliable.

If the simple view is true, then competent speakers should be in a position to know as much. And if our use of ‘tasty’ and ‘might’ adhere to the pleasure rule and the information rule, respectively, a competent speaker should be in a position to know that, too. So, if the simple view is true, a competent speaker should be in a position to know that their use of ‘tasty’ and ‘might’ adhere to invariably unreliable rules. And that threatens to render their adherence to these solipsistic rules unreasonable.

The problem takes the form a dilemma. As before, let’s say that the *reliability* of a gustatory propensity (information state) is the rate of correctness among the beliefs and assertions it produces in accordance with the pleasure rule (information rule), and let’s say that that an agent is *chauvinistic* if the agent thinks that their gustatory propensity (information state) is considerably more reliable than the pleasure rule (information rule). Either the agent is or isn’t chauvinistic.

If the agent is not chauvinistic, then their adherence to the solipsistic rules is unreasonable. The agent knows (or anyway is in a position to know) that these solipsistic rules are unreliable. So, if the agent is not chauvinistic, then the agent cannot coherently think that the beliefs and assertions they have formed in accordance with the solipsistic rules are probably correct. But if an agent cannot coherently think that their beliefs and assertions are probably correct, then the agent is not reasonable.

However, on the other horn, chauvinism, itself, seems unreasonable.

Now there is a crucial difference between objectivism and the simple view. If objectivism is true, then some chauvinistic agents are absolutely incorrect. They think that their gustatory propensity (information state) is more reliable than the pleasure rule (information rule), and that belief is absolutely incorrect. If the simple view is true, however, then no chauvinistic agent is absolutely incorrect. Relative to many (occupied) circumstances, the pleasure rule is more reliable than your gustatory propensity, and the information rule is more reliable than your information state. But, if the simple view is true, then, relative to the circumstance you occupy, your gustatory propensity is way more reliable than the pleasure rule, and your information state is way more reliable than the information rule. It might seem, then, that chauvinism is unreasonable if objectivism is true, but reasonable if the simple view is true. But the same considerations that convince me that chauvinism is unreasonable if objectivism is true also convince me that chauvinism is unreasonable if the simple view is true.

Chauvinistic agents must think that they are epistemically lucky, irrespective of whether objectivism or the simple view is true. If chauvinistic agents know that we all came by our gustatory propensities by way of a (genetic) lottery, then they must think that others got unlucky, receiving unreliable gustatory propensities, and that they, themselves, got lucky, winning the lottery, receiving one of the few reliable gustatory propensities that the lottery made available.

And irrespective of whether objectivism or the simple view is true, chauvinistic agents are foreseeably exploitable. This is perhaps easiest to see in the case of 'might'. Suppose that we have a chauvinistic agent play the IOU game from above. Whenever the agent makes an assertion about how a fair die might or might not land, we owe them \$1 if their assertion is correct, and they owe us \$2 if their assertion is incorrect. If the simple view is true, assertions concerning how a

fair die might or might not land have their correctness-values relatively, so the balance of IOUs will itself be relative, a matter that can differ across occupied circumstances. Before rolling the die, the chauvinistic agent makes six assertions: *might_1*, *might_2*, *might_3*, *might_4*, *might_5*, and *might_6*. The die lands (say) six. The chauvinistic agent then makes six more assertions: \sim *might_1*, \sim *might_2*, \sim *might_3*, \sim *might_4*, \sim *might_5*, and, reiterating, *might_6*. Which of these assertions is correct is a relative matter. But relative to every (occupied) circumstance, the agent owes us \$3 because, relative to any (occupied) circumstance, five of the assertions are false and incorrect. And this sure-loss was foreseeable from the start: the chauvinistic agent knew in advance that their chauvinism would lead them to make twelve assertions, and that the rate of correctness among those twelve assertions would be, relative to every (occupied) circumstance, 7/12.

The same foreseeable exploitation arises in the case of 'tasty'. In fact, we can use the same example from above. The chauvinistic agent hates prunes, but knows that they'll later love them. The agent adheres to the pleasure rule now, chauvinistically thinking that the gustatory propensity they now have is more reliable than the gustatory propensity they'll later have, and the agent knows that they will adhere to the pleasure rule later, chauvinistically thinking that the gustatory propensity they have then is more reliable is more reliable than the gustatory propensity they now have. Today the agent chauvinistically asserts that prunes are not tasty, and later, as they knew they would, they chauvinistically assert that prunes are tasty. So, relative to every (occupied) circumstance, they owe us \$1.

If it's reasonable for agents to be chauvinistic, then it's reasonable for agents to think that they are lucky to have the gustatory propensity (information state) that they do, and it's reasonable for agents to take sure-loss sequences of

IOU-wagers. But it does not seem to be reasonable for agents to think that they are epistemically lucky in that way, and it does not seem to be reasonable for agents to take sure-loss sequences of IOU-wagers. So it seems that chauvinism is unreasonable, even if the simple view is true. But the simple view does not satisfy the reasonability desideratum if it does not render chauvinism reasonable. So it appears that the simple view does not satisfy the reasonability desideratum.

It's worth noting that the targeted view does not have these same normative defects. Unlike the simple view, the targeted view is consistent with the pleasure rule and the information rule being reliable, even given widespread differences in gustatory propensities and information states. Thus, unlike the simple view, the targeted view does not need to render chauvinism reasonable to satisfy the reasonability desideratum. And unlike the simple view, the targeted view does not doom agents who adhere to these solipsistic rules to foreseeable exploitability. Agents should be willing to bet on the correctness of their beliefs and assertions. If the simple view is true, then an agent who adheres to these solipsistic rules is doomed to form (relatively) incorrect beliefs and doomed to make (relatively) incorrect assertions, and thus doomed to foreseeable exploitability if they undergo any foreseeable change in their gustatory propensity or information state. But the targeted view is different. If the targeted view is true, then beliefs and assertions formed in accordance with these solipsistic rules *can* be incorrect, but they don't have to be. If the actual targets are appropriately narrow, then, in principle, every belief and assertion formed in accordance with these solipsistic rules could be absolutely correct, even if the agents undergo all sorts of foreseeable changes in their gustatory propensities and information states. So the targeted view does not doom agents to foreseeable exploitability.

§7/ Conclusion

Our inquiry began with an inconsistent triad. Relative Truth, Normative Truth, and Absolute Correctness cannot all be true. Truth absolutists escape inconsistency by denying Relative Truth, but we set truth absolutism aside and asked how truth relativists should escape inconsistency.

The inconsistent triad reveals the fundamental metaepistemological division among truth relativists. Truth relativists either accept the simple view—the view that accepts Normative Truth and rejects Absolute Correctness—or they accept Absolute Correctness, reject Normative Truth, and say that what the objective epistemic ought demands of beliefs and assertions is something other than truth. I put forward two views that accept Absolute Correctness and reject Normative Truth: the reflexive view, which says that the objective epistemic ought demands that beliefs and assertions be reflexively true, and the targeted view, which says that the objective epistemic ought demands that beliefs and assertions be true at their actual targets.

I argued that the simple view makes false predictions about mistakenness, disagreement, and retraction, and leads to unacceptable normative conclusions. If the simple view is true, then agents who adhere to solipsistic rules, such as the pleasure rule and the information rule, are unreliable and foreseeably exploitable. The simple view is exciting: it entails correctness relativism—the claim that some beliefs and assertions have their correctness-values relatively—and correctness relativism deserves philosophical attention, no less so than does moral relativism. But I think that correctness relativism is false, that *everyone*—proponents and opponents of Relative Truth, alike—should accept Absolute Correctness.

The reflexive view and the targeted view are more similar than they are different. But there are differences between them, and those differences favor the targeted view. The targeted view better handles exocentric beliefs and assertions,

better handles disagreement, and better handles retraction. So I champion the targeted view: the view that beliefs and assertions are correct just when they are true at their actual targets.

References

- Austin, J. L. (1950): "Truth," in his *Philosophical Papers*, 1961, Oxford University Press, pp. 117—3.
- Boer, Steven. E and William G. Lycan (1980): "Who, Me?" *Philosophical Review*, 89, pp. 427-66.
- Boghossian, Paul (2003): "The Normativity of Content," *Philosophical Issues*, 13, pp. 31-45.
- Brogaard, Berit (2008): "In Defense of a Perspectival Semantics for 'Know'," *Australasian Journal of Philosophy*, 86, pp. 439-59.
- (2010): "Strong Representationalism and Centered Content," *Philosophical Studies*, 151, pp. 373-92.
- (2012): *Transient Truths: An Essay in the Metaphysics of Propositions*, Oxford University Press.
- Bykvist, Krister and Anandi Hattiangadi (2007): "Does Thought Imply Ought?" *Analysis* 67, pp. 277-85.
- Cappelen, Herman and Josh Dever (2013): *The Inessential Indexical*, Oxford University Press.
- Dietz, Richard (2008): "Epistemic Modals and Correct Disagreement," in García-Carpintero and Kölbel, *Relative Truth*, Oxford University Press, pp. 239-62.
- Egan, Andy (2007): "Epistemic Modals, Relativism, and Assertion," *Philosophical Studies*, 133, pp. 1-22.
- (2009): "Billboards, Bombs, and Shotgun Weddings," *Synthese*, 166, pp. 251-79.

- (2010): “Disputing About Taste,” in Feldman and Warfield (eds.), *Disagreement*, Oxford University Press, pp. 247-86.
- Egan, Andy, John Hawthorne, and Brian Weatherson (2005): “Epistemic Modals in Context,” in Preyer and Peter (eds.), *Contextualism in Philosophy*, Oxford University Press, pp. 131-68.
- Evans, Gareth (1985): “Does Tense Logic Rest on a Mistake?” in Phillips (ed), *Collected Papers*, Oxford University Press, pp. 341-63.
- Frege, Gottlob (1918[1956]): “The Thought: A Logical Inquiry,” *Mind* 65, pp. 289-311.
- Gibbard, Allan (2005): “Truth and Correct Belief,” *Philosophical Issues*, 15, pp. 338-50.
- Glanzberg, Michael (2007): “Context, Content, and Relativism,” *Philosophical Studies*, 136, pp, 1-29.
- Horwich, Paul. (1998): *Truth* (2nd rev. ed.), Clarendon Press.
- Khoo, Justin and Jonathan Phillips (2019): “New Horizons for a Theory of Epistemic Modals,” *Australasian Journal of Philosophy*, 97, pp. 309-24.
- Kölbel, Max (2002): *Truth Without Objectivity*, Routledge.
- (2004a): “Faultless Disagreement,” *Proceedings of the Aristotelian Society*, 104, pp. 53-73.
- (2004b): “Indexical Relativism Versus Genuine Relativism,” *International Journal of Philosophical Studies*, 12, pp. 297-313
- (2015a): “Relativism 1: Representational Content,” *Philosophy Compass*, 10, pp. 38-51
- (2017): “About Concerns,” in Depraetere and Salkie (eds.), *Seamtics and Pragmatics: Drawing a Line*, Springer, pp. 197-214.
- Lasersohn, Peter (2005): “Context Dependence, Disagreement, and Predicates of Personal Taste,” *Linguistics and Philosophy*, 28, pp. 643-86.

- (2009): "Relative Truth, Speaker Commitment, and Control of Implicit Arguments," *Synthese*, 166, pp. 359-74.
- (2016): *Subjectivity and Perspective in Truth-Theoretic Semantics*. Oxford University Press.
- Lewis, David (1979): "Attitudes *De Dicto* and *De Se*," *Philosophical Review*, 87, pp. 513-45.
- Liao, Shen-yi (2012): "What Are Centered Worlds?" *Philosophical Quarterly*, 62, pp. 294-316.
- MacFarlane, John (2005): "Making Sense of Relative Truth," *Proceedings of the Aristotelian Society*, 105, pp. 321-39
- (2007a): "Relativism and Disagreement," *Philosophical Studies*, 132, pp. 17-31.
- (2007b): "Semantic Minimalism and Nonindexical Contextualism," in Preyer and Peter (eds.), *Context-Sensitivity and Semantic Minimalism*, Oxford University Press, pp. 240-50.
- (2009): "Nonindexical Contextualism," *Synthese*, 166, pp. 231-50
- (2011): "Epistemic Modals are Assessment Sensitive," in Egan and Weatherson (eds.), *Epistemic Modals*, Oxford University Press, pp. 144-78.
- (2014): *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- Magidor, Ofra (2015): "The Myth of the *De Se*," *Philosophical Perspectives*, 29, pp. 249-83.
- Marques, Teresa (2014): "Relative Correctness," *Philosophical Studies* 167, pp. 361-73.
- (2018): "Retractions," *Synthese*, 195, pp. 3335-59.
- Ninan, Dilip (2010): "De Se Attitudes: Ascriptions and Communication," *Philosophy Compass*, 5, pp. 1-57

- (2016): “What is the Problem of De Se Attitudes?” in Torre and García-Carpintero (eds.), *About Oneself: De Se Thought and Communication*, Oxford University Press, pp. 86-120.
- Nozick, Robert (2001): *Invariances: The Structure of the Objective World*, Harvard University Press.
- Perry, John (1986): “Thought without Representation,” *Proceedings of the Aristotelian Society*, 60, pp. 137-51.
- Recanati, François (2004): *Literal Meaning*, Oxford University Press.
- (2007): *Perspectival Thought: A Plea for (Moderate) Relativism*, Oxford University Press.
- (2008): “Moderate Relativism,” in García-Carpintero and Kölbel, *Relative Truth*, Oxford University Press, pp. 41-62.
- (2017): “About the *Lekton*: Response to Kölbel,” in Depraetere and Salkie (eds.), *Seamtics and Pragmatics: Drawing a Line*, Springer, pp. 215-26.
- Regan, Donald (1980): *Utilitarianism and Cooperation*, Oxford University Press.
- Richard, Mark (1981): “Temporalism and Eternalism,” *Philosophical Studies*, 39, pp. 1-13.
- (2004): “Contextualism and Relativism,” *Philosophical Studies*, 119, pp. 215-42.
- (2015): “What Is Disagreement?” in Richard, *Truth and Truth Bearers: Meaning in Context, Volume II*, Oxford University Press.
- Rosen, Gideon (2001): “Brandom on Modality, Normativity, and Intentionality,” *Philosophy and Phenomenological Research* 63, pp. 611-23.
- Ross, Jacob and Mark Schroeder (2013): “Reversibility or Disagreement,” *Mind*, 122, pp. 43-84.
- Shaw, James (forthcoming): “*De Se* Exceptionalism and Frege Puzzles,” *Ergo*.
- Sosa, Ernest (2011): *Knowing Full Well*, Princeton University Press.

- Spencer, Jack (2016): “Disagreement and Attitudinal Relativism,” *Mind* 125, pp. 465-96.
- Stephenson, Tamina (2007): “Judge Dependence, Epistemic Modals, and Predicates of Personal Taste,” *Linguistics and Philosophy*, 30, pp. 487-525.
- Szabó, Zoltán Gendler (2017): “Dyadic Truth,” in Gyuris, Mady, and Rescski (eds.), *KK120: A Festschrift for Laszó Kálmán and András Kornai*, Hungarian Academy of Sciences, pp. 499-528.
- Thomson, Judith Jarvis (2008): *Normativity*, Open Court Publishing.
- Titelbaum, Michael G. (2008): “The Relevance of Self-Locating Beliefs,” *Philosophical Review* 117, pp. 555-606.
- Torre, Stephan (2018): “In Defense of De Se Content,” *Philosophy and Phenomenological Research*, 97, pp. 172-89.
- Van Inwagen, Peter (1990): *Material Beings*, Cornell University Press.
- von Fintel, Kai and Anthony S. Gillies (2008): “CIA Leaks,” *Philosophical Review*, 117, pp. 77-98.
- Williams, Bernard (1970): “Deciding to Believe,” in Kiefer and Munitz (eds.), *Language, Beliefs, and Metaphysics*, SUNY Press, pp. 95-111.
- Zimmerman, Aaron (2007): “Against Relativism,” *Philosophical Studies*, 133, pp. 313-48.