

The paper was presented at the Epistemic Value Conference, August 2006, Stirling, Scotland. The papers from the conference and the epistemic value workshop of the fall of 2005 are supposed to appear in two volumes on epistemic value published by OUP.

The Trivial Argument for Epistemic Value Pluralism. Or How I Learned to Stop Caring about Truth

Berit Brogaard

February 18, 2006

Relativism offers a nifty way of accommodating most of our intuitions about epistemic modals, predicates of personal taste, color expressions, future contingents, and conditionals. But in spite of its manifest merits relativism is squarely at odds with epistemic value monism: the view that truth is the highest epistemic goal. I will call the argument from relativism to epistemic value pluralism *the trivial argument for epistemic value pluralism*. After formulating the argument, I will look at three possible ways to refute it. I will then argue that two of these are unsuccessful, and defend the third, which involves denying that there are any genuinely relative truths.

It is generally agreed that truth is an epistemic goal. Some think it is the highest epistemic goal. It is what we strive towards when we engage in inquiry. Let us follow Michael Lynch in taking ‘inquiry’ to mean ‘the range of epistemic practices we engage in when asking and answering questions, whether banal (“where did I put my other sock”) or sublime (“can something come from nothing?”)’.¹ And let us call the view that truth is the highest epistemic goal ‘epistemic value monism’. Epistemic value pluralism then is the view that truth is only one of many epistemic goals, and that its supremacy can be overridden.²

Epistemic value monists are not oblivious to the fact that there are other things besides truth that we might be after when we engage in the practices associated with asking and answering questions. For instance, in asking ‘how do I look?’ we might be seeking a

¹ Michael Lynch, “The Value of Truth and the Truth of Values”, Epistemic Value Workshop, Stirling, Oct. 2005, p. 1. See also Lynch, “Expressivism and the Value of Truth”, forthcoming, and *True to Life* (MIT Press, 2004).

² The term is borrowed from Wayne Riggs, “Insight, Open-Mindedness and Understanding”, manuscript. Riggs takes epistemic value pluralism to be consistent with truth being the overarching aim of inquiry. But this is not how I will use the term.

compliment, and in asking ‘did you ever feel this way when you were my age?’ we might be looking for support and encouragement. But epistemic value monists think that most of these other goals are not *epistemic* goals but goals of a different kind. Nor are epistemic value monists oblivious to the fact that there are other epistemic goals besides truth, for instance, avoiding the gambler’s fallacy, forming beliefs while you are sober, being justified in believing that it is raining, being open-minded, and so on.³ But they think that these other epistemic goals are intermediary. They are goals only insofar as they facilitate our getting to the truth.

Which of the two theses, epistemic value pluralism or epistemic value monism, is correct is thought to be a difficult epistemological question: it can be decided only by reflection on the practices we engage in when asking and answering questions. In particular, it is not something that can be decided by looking at the semantics of a particular group of expressions.

This, anyway, is how things are thought to be. But if a certain very plausible semantic hypothesis is correct, then this is not how things are. The plausible hypothesis is that there are genuinely relative expressions – expressions whose extension depends not only on a context of use but also on a context of assessment. If there are genuinely relative expressions, then epistemic value monism is false. I will call this *the trivial argument for epistemic value pluralism*. After formulating the argument, I will look at three possible ways to refute it. I will then argue that two of these are unsuccessful, and defend the third, which involves denying that there are any genuinely relative expressions.

I. The Trivial Argument for Truth Value Pluralism

Epistemic value monists take truth to be our highest epistemic goal. What they mean by that is that it is *prima facie* good to believe what is true. So, if I go outside in order to find out whether it is raining, then the aim of my inquiry is to believe that it is raining if it is raining and believe that it is not raining if it is not. But, of course, I haven’t satisfied the epistemic goal if I believe

³ See Igor Douven, “The Pragmatics of Belief”, Epistemic Value Work Shop, Stirling October 2005.

that it is raining and also believe that it is not raining. The epistemic goal is not merely to believe what is true, but rather to believe all and only what is true. Where a proposition p is simply true iff it is true at the actual world, we can formulate this as follows:⁴

(EG) For any proposition p , you should believe that p iff it is simply true that p .

(EG) seems prima facie plausible. However, it is not entirely happy as it stands. For on the assumption that for any proposition p we ought to believe either p or its negation, (EG) turns out to entail a particular “Fregean” semantic thesis, namely:

(FT) For any proposition p , the truth-value of p is relative only to the world of evaluation.

Of course, it is open to deny the inference from (EG) to (FT) by rejecting the assumption that for any proposition p , we ought to believe either p or not- p . But there is a weaker connection between (EG) and (FT). If the Fregean thesis (FT) is false, then there are propositions which are not simply true or simply false. And if there are such propositions, then it is very plausible that we ought to believe some of them. But (EG) implies that if there is some proposition p such that p is not simply true, then you shouldn't believe p . So, if there is just one proposition p such that p isn't simply true and you ought to believe p , then (EG) is false. So, even if we reject the assumption that for any proposition p , we ought to believe either p or not- p , the plausibility of (EG) stands and falls with the plausibility of (FT).

(FT) gains prima facie plausibility from the fact that propositions are supposed to be whatever it is that we believe when we believe something. That is, propositions are supposed to be the objects of propositional attitudes such as beliefs. But the objects of belief seem to possess

⁴ See Michael Lynch, “The Values of Truth and the Truth of Values”, p. 1.

their truth-values eternally.⁵ Quite plausibly I do not simply believe that John is sitting. My belief seems more specific than that. For example, I might believe that John is sitting at 3 p.m. on March 8, 2006. But the proposition that John is sitting at 3 p.m. on March 8, 2006 possesses its truth-value eternally. If the proposition that John is sitting at 3 p.m. on March 8, 2006 is true at 3 p.m. on March 8, 2006, then it is also true at 4 p.m. on March 8, 2006. If entities which do not possess their truth-value eternally are not specific enough to be the objects of belief, and propositions are objects of belief, propositions are true or false only relative to a world of evaluation, which is to say that (FT) is true.

But, as David Kaplan and others have made vivid,⁶ there are also arguments against (FT). Propositions are thought to be the semantic value of sentences in contexts. For example, the semantic value of an utterance of the sentence ‘John is currently a firefighter’ at 3 p.m. on March 8, 2006 is supposed to be the proposition that John is a firefighter at 3 p.m. on March 8, 2006. But the view that propositions are the semantic value of sentences in context makes trouble for (FT). Consider, for instance, my utterance at 3 p.m. on March 8, 2006 of ‘John is firefighter’ and your utterance at 3 p.m. on March 8, 2006 of ‘it has been that John is firefighter’. In standard semantics, ‘it has been that’ is treated as a sentential operator. It takes us from our current circumstance of evaluation to a circumstance in the past. Your utterance of the sentence ‘it has been that John is firefighter’ is true just in case ‘John is firefighter’ is true at a time before the time of your utterance. So, if my utterance is false and yours is true, then the semantic value of ‘John is firefighter’ has different truth-values depending on whether the sentence occurs in a temporal context or not. But this strongly suggests that ‘John is a firefighter’ does not express a

⁵ Or locations.

⁶ David Kaplan, “Demonstratives”, in J. Almog, H. Wettstein, and J. Perry (eds), *Themes from Kaplan*, New York: Oxford University Press, 1989, 233-249. See also David Lewis, “Index, Context and Content”, 1980, reprinted in Lewis 1998, *Papers in Philosophical Logic*, Cambridge University Press, Cambridge, UK, and Jeffrey C. King, “Tense, Modality, and Semantic Values”, *Philosophical Perspectives* 17 (2003): *Language and Philosophical Linguistics*, John Hawthorne and Dean Zimmerman (eds.), 195-246.

Fregean proposition. For sentences are supposed to have the same semantic value regardless of the linguistic environment in which they occur.

So, if propositions are both the objects of belief and the semantic values of sentences in context, then it would seem that they both possess and do not possess their truth-values eternally.

David Lewis and others have proposed to resolve this tension by distinguishing between the compositional semantic value of a sentence in context and the proposition expressed by that sentence.⁷ ‘It has been that John is firefighter’ expresses a proposition that is the object of your belief that John was a firefighter, but it also has a distinct compositional semantic value that is composed of the semantic values of the parts.

The distinction between the propositional semantic values and the compositional semantic values of sentences allows us to preserve (FT). But there is a different way to preserve (FT). Jeffrey King has recently argued that temporal expressions (including the tenses) are not operators on sentences but object-language quantifiers over times.⁸ On this view, ‘John is a firefighter’ has a hidden indexical variable whose value is supplied by context. This variable is bound when ‘John is a firefighter’ is embedded in a temporal context, as in ‘John was a firefighter’. The latter expresses, relative to context, the proposition that there is a time t such that t is earlier than the time of speech, and John is a firefighter at t . If the tenses are object-language quantifiers, then there is no need to recognize a special compositional semantic value

⁷ Lewis, “Index, Context and Content”, Salmon, Nathan Salmon, *Frege’s Puzzle*, The MIT Press/Bradford Books, Cambridge, MA, 1986, Salmon, “Tense and Singular Propositions” in *Themes From Kaplan*, Oxford University Press, New York, Almog, Perry, Wettstein (eds.), 1989, Michael Dummett, *The Logical Basis of Metaphysics*, Harvard University Press, Cambridge, 1991, Jason Stanley, 1997a, “Names and Rigid Designation”, in *A Companion to the Philosophy of Language*, ed. B. Hale and C. Wright, Oxford: Blackwell Publishers, 1997b, “Rigidity and Content”, in *Logic, Language and Reality: Essays in Honor of Michael Dummett*, R. Heck (ed.), Oxford University Press, Oxford.

⁸ Jeffrey C. King, “Tense, Modality, and Semantic Values”. King backs up his proposal with linguistic evidence. However, I think the evidence he presents does not actually support a quantificational treatment of the tenses. One example King gives is this: ‘yesterday, I turned of the stove’. King thinks that the tense operator approach is bound to say that ‘yesterday’ is an operator. Treating ‘yesterday’ as an operator does indeed give us the wrong result. However, there is good reason not to treat ‘yesterday’ as an operator. ‘Yesterday’ is better treated as a component part of the composite tense operator ‘it has been yesterday that’. This treatment gives us the right result. For the details of this argument, see my forthcoming book *Transient Truths – An Essay in the Metaphysics of Propositions*, chap. 2.

for tense operators to operate on. Treating tense operators as object-language quantifiers thus makes it possible to maintain the view that propositions are both the objects of belief and the semantic values of sentences.

There is much to be said for both of these proposals. However, I think that there is independent reason for rejecting (FT).⁹

One reason is this: when I say ‘John is a firefighter’ and you say ‘Nuh-uh, John is not a firefighter. He is a police officer’, it seems we are disagreeing. But this is not so if (FT) is correct. For when I say ‘John is a firefighter’ at t_1 , what I am saying is that at t_1 John is a firefighter, and when you say ‘John is a firefighter’ at t_2 , what you are saying is that at t_2 John is not a firefighter. But *at t_1 John is a firefighter*’ and *at t_2 John is a firefighter* do not contradict each other. Note, how, unnatural-sounding the following exchange is.

(A and B are talking on the phone. B is standing outside the door of an office where a conversation is taking place between John and his superior).

A: ... John is a firefighter.

(Behind closed doors the superior shouts: ‘you are fired’)

B: I guess you are right. But John is not a firefighter.

If A asserted that John is a firefighter at t_1 , then we should expect B to reply exactly as he does. For it is still true at t_2 that John is a firefighter at t_1 . But it would make no sense for B to respond in this way. It would make much more sense for him to reply with ‘no, you are wrong. I am standing outside the superior’s office, and his superior just told him that he was fired’.

A critic may rejoinder that the argument is too quick. Perhaps the semantic value of the hidden indexical variable in ‘John is a firefighter’ is not always the time of speech. ‘John went to

⁹ For a lengthier defense of temporalism, see my forthcoming book *Transient Truths – An Essay in the Metaphysics of Propositions*, especially chapters 1 and 3.

a local bar' need not be interpreted as 'John went to a bar that is local to the speaker' but may be interpreted as 'John went to a bar that is local to John', 'John went to a bar that is local to the hearer', 'John went to a bar that is local to John's grandmother', and so on. The speaker can more or less freely fix the hidden indexical variable associated with 'local'. Perhaps the same goes for 'John is a firefighter'.

I think, however, that this version of eternalism is even less plausible. Consider the following dialogue between A and B:

AS I SAID

A: What does John do for a living?

B: He is a firefighter.

C: Nuh-uh, he was fired two years [weeks/days/minutes/seconds] ago. He is a lawyer now.

B: [Aggravated] So, as I said, John is a firefighter.

B's last remark is exceedingly odd. If, however, the value of the alleged indexical variable can be freely chosen, then we should expect it to be felicitous.

A second reason to doubt that (FT) is correct is that it seems to yield an inadequate account of belief revision. If I learn that John was fired, intuitively I will not simply add one more belief to my belief system; I will dispose of my belief that John *is* a firefighter and add the belief that John *was* a firefighter but is no longer.

But this is not what we get if (FT) is true. Suppose I believe that John is a firefighter at 3 p.m. on April 5, 2004. If (FT) is true and you inform me that John is not a firefighter at 3 p.m. on August 8, 2004, then this would give me no reason to discard my old belief. After all, I have been given no reason to believe that my old belief is false. If eternalism is true, then all of our beliefs are temporally specified in this way. So if they are true at all, then they remain true when

the world changes. Contrary to appearances, then, changes in the world will give us no reason to revise our original beliefs.

But changes in the world do give us reason to revise our beliefs. If I am told on August 8, 2006 that John has been fired, then there is something I cease to believe: I cease to believe that John is a firefighter. The following exchange illustrates this point:

RANKINGS

(A and B are talking in the hallway)

A: Our department is number two on the Leiter Report!

(One week later)

B: What you said last week isn't true anymore. I just saw the rankings.

A: Of course, it is. Haven't you read Frege? Maybe that's why we dropped.

If (FT) is true, then the proposition A believes on the earlier occasion cannot change its truth-value over time. So A's last remark should be perfectly fine.

A third reason to question (FT) is that it yields an unrealistic account of belief retention. If (FT) is true, then I do not believe that John is a firefighter, period. I believe that John is a firefighter at 15:00 on July 5, 2004, that John is a firefighter at 15:01 on July 5, 2004, and so on. So, to "continue" to believe what I would express with 'John is a firefighter' for five years is to believe an indefinite number of propositions of the form 'John is a firefighter at t ', one for each moment of time during the five years. But that can't be right. Surely, there is just a single belief that I continue to have for the five years, namely the belief that John is a firefighter.

Relatedly, when we say 'I still believe that p ', the advocate of (FT) is required to say that we mean that we still believe the same time-indexed proposition. But if that were right, then it would be hard to explain cases of the following sort:

DECEIT

Wife: When I married John I thought he was a police officer. Thirty years later I still believed he was a police officer. Turns out that he was fired two years into our marriage.

or:

LOST LOVE

Friend: Yes, Barbara did love you ten years ago. So you were right back then. But you *still* believe that she loves you, don't you Peter?

or:

DEFENSE

Student: I think my dissertation is done.

Supervisor: You do? Well, I think you are wrong. Work on it for a few more weeks.

Then read it again. If you *still* think that it's done, then we'll talk.

If (FT) is true, then what we believe when we believe something is a time-indexed proposition. But it is hardly the case that the wife in DECEIPT means that she still believed the same time-indexed proposition after thirty years, viz. the proposition *my husband is a police officer at t*, where *t* is some time thirty years ago, that the friend in LOST LOVE means that Peter still believes the proposition *Barbara loves me at t*, where *t* is some time ten years ago, or that the supervisor in DEFENSE is asking S to return if S still believes the proposition *S's dissertation is done at t*, where *t* is the time of their exchange. To my mind, such cases raise one of the most pressing kinds of problems for advocates of (FT).

But if there is at least one proposition p such that p is true at one time but false at another, then (FT) is false. A more plausible candidate for being the object of my belief that John is a firefighter is what we might call a *tensed proposition*. Tensed propositions change their truth-values across time. In fact, as Arthur Prior has argued,¹⁰ there is overwhelming evidence for thinking that present-tensed propositions just are temporal propositions, i.e., propositions that determine functions from worlds and times to truth-values. If he is right, then (FT) is false. The existence of tensed propositions, of course, does not rule out that there could not also be tenseless propositions, for instance, the proposition that John is a firefighter at 3 p.m. on July 5, 2006. But (FT) states that for any proposition p , the truth-value of p is relative only to the world of evaluation. And this is false, if some tensed propositions are worthy of belief.

Let us return the above formulation of the epistemic goal. Above we formulated the epistemic goal as follows:

(EG) For any proposition p , the epistemic goal is to believe that p iff it is simply true that p

As we saw, (EG) is plausible only if (FT) is:

(FT) For any proposition p , the truth-value of p is relative only to the world of evaluation.

¹⁰ *Time and Modality* (Oxford: Oxford University Press 1957), pp. 9-10, *Past, Present and Future* (Oxford: Clarendon, 1967), pp. 8-10, 14-15, *Papers on Time and Tense* (Oxford: Oxford University Press, 1968), 17-23, "Now", *Nous* 2 (1968), 101-119. Prior points out that there is a difference between the present tensed sentence 'John is a firefighter' and the indexical sentence 'John is currently a firefighter'. Unlike tense operators like 'it has been that', 'it will be that' and 'it is the case that', 'it is currently the case that' and 'it is now the case that' take primary scope with respect to other operators.

So, if (FT) is questionable, so is (EG). To conclude that epistemic value monism is doubtful would be too quick, however. For there is a semantic thesis in the neighborhood that may be true, namely:¹¹

(KT) For any proposition p , and any time t , p has a truth-value relative to t .

Keeping the world fixed, the epistemic goal can then be reformulated as follows:

(EG') For any proposition p and any time t , the epistemic goal is for you to believe that p at t if and only if p is true at t , and the truth of p does not depend on any further parameters.

A problem arises, however, if there are genuinely relative expressions in the language, as a number of relativists have argued. John MacFarlane, for example, has argued that 'know' and 'might' are plausible candidates for being relative expressions, whereas future contingents (e.g. 'there will be a sea battle tomorrow') are plausible candidates for being relative sentences. Others have offered a relativistic semantics for predicates of personal tastes (e.g. 'fun' and 'tastes great'), color expressions (e.g. 'is the same color as'), and non-standard uses of gradable adjectives (e.g., 'huge' and 'rich').¹²

¹¹ 'KT' stands for 'Kaplan's thesis'. The world is kept fixed.

¹² See e.g. Max Kölbel, *Truth Without Objectivity* (Routledge: London, 2002), Kölbel, "Faultless Disagreement", *Proceedings of the Aristotelian Society* 104 (2003), 53-73, John MacFarlane, "Future Contingents and Relative Truth", *The Philosophical Quarterly* 53 (2003), 321-36, "Making Sense of Relative Truth", *Proceedings of the Aristotelian Society* 105 (2005), 321-39, "The Assessment Sensitivity of Knowledge Attributions", T. Szabo-Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology* 1, 2005, Mark Richard, "Contextualism and Relativism", *Philosophical Studies* 119 (2004), 215-42, Andy Egan, John Hawthorne, and Brian Weatherson, "Epistemic Modals in Context", in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy* (Oxford: Oxford University Press, 2004), and Weatherson, "Conditionals and Relativism", manuscript.

A relative expression is one whose *extension* varies with the context of use *and* a context of assessment. A context of assessment is a context in which a sentence is considered or evaluated for truth. Any context of use is a context of assessment, but not every context of assessment is a context of use. In a context in which I evaluate John's utterance of the sentence 'I am hungry', for example, 'I am hungry' is evaluated for truth in my context but, according to the relativists, it is not used in that context.

Relativists claim that the existence of relative expressions in the language requires us to revise standard semantics. The content of relative expressions is context invariant. A relative expression has the same semantic value at any context of use and any context of assessment. Its semantic value is (or determines) a function from circumstances of evaluation to extensions. But the relativists' circumstances of evaluation are non-standard. They contain a world and a time parameter determined by the context of use, and a judge parameter determined by the context of assessment. So, if I utter the sentence 'this chili tastes great' at some time t in the actual world $@$, and you evaluate my utterance at some time t^* in the actual world $@$, then our contexts together determine a $\langle @, t, \text{you} \rangle$ circumstance. The sentence 'this chili tastes great' is true relative to my context of use and your context of assessment iff the proposition expressed by 'this chili tastes great' is true at $\langle @, t, \text{you} \rangle$.

Standard semantics does not allow for circumstances that are determined partially by a context of use and partially by a context of assessment. The circumstances of standard semantics are fully determined by the context of use. So, if there are any relative expressions in the language, then standard semantics is in need of revision.

The relativists' proposed revision of standard semantics does not affect the truth-value and content of just any old sentence. For instance, if I utter the sentence 'I am human' at a time t in the actual world $@$, my context of use and my context of assessment determine a $\langle @, t, \text{me} \rangle$ circumstance. If you then evaluate my utterance for truth at t^* , then our contexts together determine a $\langle @, t, \text{you} \rangle$ circumstance. But if none of the expressions in 'I am human' is a

relative expression (which is very plausible), the truth-value of my utterance of the sentence ‘I am human’ will be the same regardless of who evaluates the utterance.

But if relativism is right, then there are some propositions which have no truth-value relative only to a world and a time. For instance, the proposition expressed by my utterance of the sentence ‘this chili tastes great’ has no truth-value only relative to my context of use. Rather, it is true or false only relative to a context of assessment. So, if John considers my assertion and dislikes the chili, then the proposition is false, and if Mary considers my assertion and likes the chili, then the proposition is true.

But this makes trouble. If I really enjoy this chili at time t , I will be inclined to believe at t that it tastes great. But if ‘tastes great’ is a genuinely relative expression, then ‘this chili tastes great’ does not have a truth-value only relative to a world and a time. My believing that this chili tastes great thus violates (EG’). So, I should not believe that this chili tastes great. But if ‘tastes great’ is a genuinely relative expression, then ‘this chili does not taste great’ does not have a truth-value only relative to a world and a time either. So, my believing that this chili doesn’t taste great also violates (EG’). So, I should not believe that this chili does not taste great. But this seems just plainly false. Vagueness to one side, I should believe either that this chili tastes great or that it does not taste great. But then truth cannot be our only epistemic goal. Hence, epistemic value monism is false.

II. First Reply: Revising the Epistemic Goal

Both epistemic value monists and epistemic value pluralists should be worried about this argument. For the conclusion of the argument concerns the aim of the practices that we engage in when asking and answering questions, whereas the main premise is an empirical hypothesis about English (or some other language). But intuitively, the question of what the ultimate aim of the practices we engage in when asking and answering questions cannot be decided by looking at the

semantics of a particular group of expressions. There is thus good reason to think that the trivial argument is flawed. The question is where it goes wrong.

It might perhaps be thought that the existence of relative expressions in English does not prove that epistemic value monism is wrong but only that we need to revise epistemic value monism to take account of assessment-sensitivity. Arguably, there are independent problems with our above formulations of the epistemic goal. As Lynch has argued, we cannot believe everything that is true, and it is highly doubtful that our epistemic goal is to be God.¹³ Lynch suggests that we revise the standard formulation of the epistemic goal to take account of our human deficiencies. The epistemic goal, he suggests, may be formulated as follows. For any p that we are able to consider our epistemic goal is to believe p if and only if p is (simply) true. Lynch admits that might be difficult to unpack the modality of ‘is able to’, but adds that whatever it means, it is not supposed to mean that we should believe only the true propositions that we happen to consider.

Lynch’s reformulation of our epistemic goal seems to fit the relativist’s purposes better than our original formulation, for it concerns only propositions that we are able to consider. It thus appears to be about believers in contexts of assessment rather than believers in contexts of use. It is natural, therefore, to suggest that we formulate the epistemic goal as follows:

(EG") For any time t and any proposition p that you are able to consider, the epistemic goal is for you to believe that p at t if and only if it is true that p relative to the standards in play in your context of assessment C_A at t

So, if John is able to consider the proposition that this chili tastes great, then he should believe that this chili tastes great if and only if this chili tastes great relative to the standards in play in his context of assessment.

¹³ “The Values of Truth and the Truth of Values”, p. 2.

Unfortunately, (EG'') is mistaken if relativism is right. For relativists say that the extension of a proposition is a function of the world and time of the context of use, and the standards in play in the context of assessment. A proposition is not simply true relative to context of assessment; it is true relative to a context of use and a context of assessment. The context of use determines a time and a world at which to evaluate the proposition and the context of assessment determines the standards with respect to which the proposition is to be evaluated. If John said last year before graduating, 'I am a student', then the proposition expressed is true relative to the time determined by John's context of use, but not true relative to the time determined by my current context of assessment. So, the truth-value of the proposition that John is a student will depend on the context of use.

The following principle might be thought to do better as a formulation of our epistemic goal:

(EG''') For any time t , any proposition p that you are able to consider and any context of use C_U , the epistemic goal is for you to believe that p at t if and only if it is true that p relative to the time t^* determined by C_U and the standards in play at your context of assessment at t

But this principle can't be right either. For the present-tensed proposition that John is a student may be true relative to the time determined by Peter's context of use but false relative to the time determined by Mary's context of use. By (EG'''), we should and shouldn't believe that John is a student. So, (EG''') instructs us to have inconsistent beliefs.

Our best option, it seems, is to formulate the epistemic goal as follows:

(EG''''') For any time t and any proposition p that you are able to consider, the epistemic goal is for you to believe that p at t if and only if it is true that p relative to the

time determined by your context of use C_U at t and the standards in play at your context of assessment C_A at t .

But (EG''') is not entirely happy either. One problem with (EG''') is that your context may not be a context of use with respect to the proposition that you are considering. If, for example, John utters the sentence 'this chili tastes great' and you are considering the proposition expressed by John's utterance, then the context of use is distinct from your context. We may, of course, treat any sequence of parameters of the right sort as a context of use, as David Kaplan argued,¹⁴ and there is much to be said for Kaplan's notion of a context; but there is a further, and more serious, problem with (EG''').

(EG''') says that our highest epistemic goal is to believe what is true in the special case where the context of use C_U and context of assessment C_A are identical. But this seems like a refutation of the claim that there are relative truths. We have a notion of being true at a context of use C_U and a context of assessment C_A , but it is not one that has any value to us except when $C_U = C_A$.¹⁵ Truths at a context of use C_U and a context of assessment C_A are not truths that are worthy of belief; they are not truths we should believe. In fact, they are truths that we should *not* believe. To put the point another way: the only truths worthy of belief are those that might count as truths by non-relativist lights. So, if the highest epistemic goal is to believe what is true in the special case where the context of use C_U is identical to the context of assessment C_A , then our highest epistemic goal is to believe all and only non-relative truths; genuine relative truth is of no value to us. But if there are genuinely relative truths, then this is unacceptable. So, it seems doubtful that reformulating the epistemic goal will help to block the trivial argument.

III. Second Reply: Rejecting Relativism

¹⁴ D. Kaplan, "Demonstratives".

¹⁵ MacFarlane raises a related objection to a related formulation of the norm of assertion. See "Making Sense of Relative Truth", p. 331.

A second way to respond to the trivial argument is to insist that the fact that relativism cannot take truth to be our prime epistemic goal shows that the assumption that there are genuinely relative expressions is mistaken. The problem with this reply, however, is that the evidence for there being genuinely relative expressions is strong.

The most obvious candidates for being relative expressions are predicates of personal taste (e.g. ‘fun’ and ‘tastes great’). As Peter Lasersohn points out, ‘with predicates of personal taste, we actually operate from a position of epistemic *privilege*, rather than the opposite. If you ride the roller coaster, you are in a position to speak with authority as to whether it is fun or not; if you taste the chili, you can speak with authority as to whether it is tasty’.¹⁶ But there is equally good evidence for thinking that epistemic expressions (e.g. ‘might’, ‘know’ and ‘justified’), color expressions (e.g. ‘the same color as’), and so on are relative expressions. For expository simplicity, I shall here focus on predicates of personal taste. But my argument can easily be extended to take account of other expressions thought to be genuinely relative.

The first piece of evidence for treating predicates of personal tastes as relative expressions comes from the fact that there is widespread disagreement about what tastes great, tastes horrible, is fun and so on. From the fact that I find this chili tasty I cannot infer that you will find it tasty as well. Semantic invariantists think expressions like ‘taste great’ have the same value and extension in any context of use and any context of assessment. So, they cannot account for the insolubility of disputes over what is tasty without ascribing an implausible systematic semantic incompetence to ordinary speakers.

Contextualism seems more plausible. For example, it might be suggested that predicates of personal taste contain a hidden indexical variable which has the speaker as a default value. So, when I say ‘this chili tastes great’, I am in effect saying that this chili tastes great to me. And when you say ‘this chili tastes great’, you are saying that this chili tastes great to you’. So

¹⁶ Lasersohn, “Context Dependence, Disagreement, and Predicates of Personal Taste”, *Linguistics and Philosophy* 28 (2005), 643-86, p. 655.

construed, contextualism can explain why the truth-value of sentences containing predicates of personal taste seems to vary from speaker to speaker. But it cannot explain why disagreements about what tastes good, what is fun, and so on, often seem faultless. The following disagreement, for example, seems faultless:

(I)

John: This chili tastes great.

Mary: I disagree. It does not taste great at all.

The following exchange, on the other hand, points toward semantic incompetence on Mary's part:

(II)

John: I am hungry

Mary: I disagree. I am not hungry at all.

If the contextualists are right, then (I) and (II) both indicate that Mary is semantically incompetent. But most of us fail to see the alleged semantic mistake made by Mary in (I). So, if contextualism is right, then most of us are ignorant of the semantics of predicates of personal taste.

A related problem for contextualism is that when our internal affective states change, we often retract our earlier judgments. Consider the following exchange:

(III)

John: ... chocolate ice cream tastes great

Mary: but a few months ago you said that it didn't taste great.

John: I was mistaken. I didn't really appreciate the taste back then.

If John means different things by ‘tastes great’ on the two different occasions, as contextualism tells us, we should expect John to respond differently, for instance, with ‘yes, but I was right back then as well. I just didn’t mean the same thing by “taste great”’. But this is not how we behave at all.

Other expressions typically taken to be context-sensitive, such as ‘tall’ and ‘flat’, seem to function differently.¹⁷ If we are talking about the height of men in general, I might say ‘Michael Jordan is tall. He is 6’6"!’. If we are talking about the height of NBA basketball players, on the other hand, I might say ‘Michael Jordan is not tall. He is only 6’6”’. If reminded of the apparent discrepancy, I will not say: ‘I was mistaken. I didn’t realize how small he really is’. Instead, I will say something like: ‘I only committed myself to Jordan being tall for a man. The average height for a man is 5’10”. But the average height for an NBA player is 6’6”. So, Michael is tall for a man, but not for an NBA player’. This sort of difference between predicates of personal taste and other expressions generally regarded as context-sensitive count against treating the content of the latter as contextually variable. It seems that we do not, as a matter of fact, take the properties expressed by ‘tastes great’ and ‘fun’ to vary across contexts of use.

A third problem for contextualism turns on the behavior of predicates of personal taste in propositional attitude reports. Suppose that after her conversation with John Mary meets Peter who is throwing a party. Peter wants to know how John likes chocolate ice cream. Mary replies:

(1) He thinks it tastes great

(1) seems to be an adequate report of what John believes, provided, of course, that John complied with the norms of assertion on the earlier occasion. The report seems adequate even if Mary believes it doesn’t taste great. But if she believes it doesn’t taste great, then the contextualist is

¹⁷ See e.g. John MacFarlane, “The Assessment Sensitivity of Knowledge Attributions”.

bound to say that John and Mary mean different things by ‘tastes great’. Since Mary is the speaker of (1), the content of ‘tastes’ depends on her internal affective states. So, (1) represents John as believing that chocolate ice cream tastes great, by Mary’s lights. But evidently this is not what John believes.¹⁸

Both invariantism and contextualism apparently have difficulties accounting for the linguistic data involving predicates of personal taste. Relativism fares considerably better.

First, relativism explains why there is widespread and insoluble disagreement about what is tasty, what is fun, and so on. There is widespread disagreement about what is tasty, what is fun, and so on, because the truth of sentences containing predicate of personal taste is relative to the standards in play in the context of assessment. So, the proposition that this chili is tasty is true when assessed by John only if John finds this chili tasty, and it is true when assessed by Mary only if Mary finds this chili tasty.

Second, relativism explains why the disagreement appears faultless. For relativists about ‘fun’ and ‘tastes great’ hold that the semantic content of these expressions is constant across contexts. So, if John and Mary disagree about whether this chili tastes great, then there is a proposition the truth-value of which is the object of their disagreement, namely the proposition that this chili tastes great.

Third, relativism explains why we are strongly inclined to retract our earlier judgments when our internal affective states change. We are inclined to do this, because sentence truth is relativized to a context of use and a context of assessment. ‘Chocolate ice cream is disgusting’ might be true when assessed from John’s earlier perspective but false when assessed from his later perspective.

Fourth, relativists claim to be able to explain why we are quick to disquote. Since the content of ‘tastes great’ is contextually invariant, Mary’s report ‘John believes that chocolate ice

¹⁸ See e.g. “Context-shifting Arguments”.

cream tastes great' simply expresses the proposition that John stands in the *believe that* relation to the proposition that chocolate ice cream tastes great.

Relativism about predicates of personal taste thus provides a powerful explanation of the linguistic data involving predicates of personal taste without having to resort to such dubious phenomena as semantic blindness or widespread semantic incompetence. Rejecting the thesis that there are relative expressions without offering an alternative that explains the linguistic evidence equally well would be ad hoc and desperate.

IV. Third Reply: Offering an Alternative

A third thing one might do when faced with the trivial argument is reject the assumption that there are relative expressions but offer an alternative theory that explains all the same linguistic data as relativism. This is the option I will defend. I propose the following theory as a replacement for relativism. Predicates of personal taste have the same content across context; the content determines a function from circumstances of evaluation to extensions. The default circumstance is fully determined by the context of use which, like Kaplan, I will take to be a sequence of a speaker, an addressee, a world, a time, and a location. But unlike Kaplan's default circumstances my default circumstances contain a judge parameter whose default value is the speaker. So, my assertion of 'this chili tastes great' is true only if it is true at a circumstance containing me as the judge. To set it apart from standard contextualism I shall call this view *perspectivalism*.¹⁹ Perspectivalism, I will now argue, explains all the same data as relativism without relativizing sentence truth to a context of use and a context of assessment.

The prime reason for relativizing sentence truth to contexts of use and contexts of assessment is that there seem to be cases of sentences that are true relative <world, time, judge>

¹⁹ MacFarlane calls this brand of contextualism 'nonindexical contextualism'. See "Semantic Minimalism and Nonindexical Contextualism", forthcoming in *Content and Context: Essays on Semantics and Pragmatics*, ed. G. Preyer and G. Peter and "Non-Indexical Contextualism", forthcoming in Brogaard, ed. *Synthese: Relative Truth*. As we will see, my account differs from MacFarlane's nonindexical contextualism with respect to how it deals with reported sentences.

circumstance, where the value of the world and time parameters are determined by the context of use and the value of the judge parameter is determined by the context of assessment. Consider, for instance, John's utterance at *t* of:

- (2) The sentence 'this chili will taste great tomorrow' was false when Mary uttered it on the earlier occasion

(2) ascribes falsehood to a sentence-in-context, namely the sentence reported in (2). The truth-value of the reported sentence depends on the time and world at which Mary uttered it, and the standards in play in John context. Thus, Mary's utterance of the sentence 'this chili will taste great tomorrow' is true relative to John's context of assessment iff the chili in question tastes great by John's standards, the day after Mary uttered the sentence.

But, as I will now argue, we do not need to distinguish between a context of use and a context of assessment to get this result. We only need to treat 'tastes great' as a perspectival, i.e. an expression whose extension is a function of a circumstance containing a judge parameter whose default value is the speaker.

Cases like (2) seem to pose a problem for all brands of contextualism because the embedded sentence is mentioned but does not appear to be used. This appearance is illusory, however. The quoted sentence in (2) is a direct speech report. But, as François Recanati has argued,²⁰ there is good reason to think that a reported sentence is used, not simply mentioned. One reason is that the quoted material in a direct speech report may be available for copying, as in:²¹

- (3) 'I'm going to see the dean', she said; and she did.

²⁰ "Indexicality and context-shift", Conference Paper, Workshop on Indexicals, speech acts, and logophors, Harvard University, (2004, 11-20), § 3.2.

²¹ The first argument is Recanati's (§ 3.2, example 4); the second is my own.

‘And she did’ is elliptical for ‘and she did see the dean’. So, the elided material ‘see the dean’ is available for copying. But this suggests that the quoted material is used. For if the sentence were merely mentioned, the quoted material would not be available for copying, witness ‘“I’m going to see the dean” is a sentence; and she did’.

Another reason to think a quoted sentence in a direct report is used not simply mentioned is that expressions in the matrix clause can depend anaphorically on expressions in the quoted sentence. Consider, for instance:

(4) ‘Give me your money_i, or I’ll shoot’, he said, but I didn’t give it_i to him.

The pronoun ‘it’ in the matrix clause is anaphoric on ‘your money’. But this requires that ‘your money’ picks out an individual for ‘it’ to refer to.

Recanati suggests that direct speech reports create a shifted context with respect to which the expressions that occur in the scope of the report are assigned a semantic value. For example, if I assert ‘and then John said, “I am leaving now” ’ at t , then the direct speech report creates a shifted context in which John is the value of the speaker parameter, and some past time t^* is the value of the time parameter. So, the value of ‘I’ is John, and the value of ‘now’ is t^* . Direct speech reports can thus be seen to function as context-shifters, or monsters, as Kaplan called them.²² Kaplan thought there weren’t any context-shifting *operators* in English, but didn’t rule out direct speech report shifts.²³

Let us return to John’s utterance at t of:

²² D. Kaplan, “Demonstratives”, p. 510. Kaplan: ‘no operator can control the character of the indexical within its scope, because they will simply leap out of its scope to the front of the operator. I am not saying we could not construct a language with such operators, just that English is not one. And such operators could not be added to it’.

²³ As he puts it, ‘there is a way to control an indexical, to keep it from taking primary scope, and even to refer it to another context. Use quotation marks. If we mention the indexical rather than use it, we can, of course, operate directly on it’.

- (2) The sentence ‘this chili will taste great tomorrow’ was false when Mary uttered it on the earlier occasion.

If direct speech reports create shifted contexts, we can explain the intuitive truth-value of (2) without relativizing sentence truth to context of use and contexts of assessment. John’s context is a context of use with respect to which the expressions in (2) are assigned a semantic value. The quoted sentence ‘this chili will taste great tomorrow’ is not assigned a semantic value with respect to John’s context. Instead, it is assigned a semantic value with respect to a shifted context in which John is the speaker, and the time is some time t^* before t . So, the semantic value of any indexical expressions in the quoted sentence will depend on the parameters of the shifted context. The semantic value of ‘tomorrow’ is thus the day after t^* , not the day after t . But ‘tastes great’ is not an indexical, its content does not vary across contexts of use; so it has the same semantic value relative to John’s and Mary’s contexts. Its content is a function from $\langle \text{world, time, judge} \rangle$ triples to extensions.

Besides fixing the semantic values of the expressions in (2), John’s context determines a $\langle w, t, \text{John} \rangle$ circumstance with respect to which (2) is evaluated. The quoted sentence in (2) is not evaluated with respect to this circumstance. But the reason is not that it is quoted. For direct speech reports do not by themselves shift the parameters of the circumstance of evaluation. To shift the parameters of the circumstance we need a circumstance shifter. For example, the proposition expressed by ‘I am hungry’ is evaluated with respect to the circumstance determined by the context of use in an assertion of ‘John says, “I am hungry” ’ but with respect to a shifted circumstance in an assertion of ‘John *said*, “I am hungry” ’. The matrix of (2) does contain a circumstance shifter, namely the past tense operator. But tense operators are selective; they shift only the time feature of the circumstance of evaluation. The past tense operator shifts the time t to some time t^* before t . So, the circumstance with respect to which the quoted sentence is

interpreted is a $\langle w, t^*, \text{John} \rangle$ triple. The quoted sentence is true relative to this circumstance iff this chili will taste great the day after t^* by John's standards. So, (2) is true iff Mary's assertion of 'this chili will taste great tomorrow' is false by John's standards, which is as it should be.²⁴

Our account of (2) can be extrapolated to account for any case of a sentence being evaluated for truth. For on Kaplan's notion of context any sequence of a speaker, addressee, world, time and location count as a context. A fortiori, a context in which the speaker is not saying anything counts as a context. Hence, a context at which a sentence is quietly evaluated for truth will count as a context. The present account can thus handle any case where a sentence is being evaluated for truth without relativizing sentence truth to a context of assessment.

Before returning to our discussion of epistemic value I would like to consider a couple of possible objections to the present proposal. One might be worried about the non-standard circumstances I am positing. One cannot just posit the existence of whatever circumstances one would like. The reason Kaplan included a time and world parameter in his circumstances was that those parameters could be shifted by circumstance shifting operators. But, it may be said, there is no operator that can shift the value of the judge parameter.²⁵

By way of reply, I think that there *are* circumstance shifters that can shift the value of the judge parameter. 'According to so-and-so, p ' is true iff for all worlds w compatible with the exact content of what so-and-so would assent to at the world of speech @ and the time of speech t^* , p is true at $\langle w, t^*, \text{so-and-so} \rangle$. But consider now my utterance of:

(5) According to John, this chili tastes great

²⁴ I assume that there is a difference between the truth-predicate for utterances when restricted to the meta-language and when it is not so restricted. When the truth-predicate for utterances is restricted to the meta-language, it functions in just the way suggested by Kaplan. See B. Brogaard, "Sea-Battle Semantics", *Philosophical Quarterly*, forthcoming.

²⁵ Jason Stanley, "Against Knowledge Relativism", <http://pwp.netcabo.pt/0154943702/Handout2.pdf>. This is reprinted (with some revisions) as Chapter 7 of Stanley's recent book *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005).

(5) is true iff for all worlds w compatible with the exact content of what John would assent to at $\langle @, t^* \rangle$, the proposition *this chili tastes great* is true at $\langle w, t, \text{John} \rangle$. In other words, for the proposition *This chili tastes great* to be true at a circumstance compatible with the exact content of what John would assent to here and now, someone who would assent to the very things John would actually assent to must be the judge in that circumstance. Propositional attitude verbs function in the same way. ‘John believes that this chili tastes great’ is true only if ‘this chili tastes great’ is true at a shifted circumstance that has John as the judge.

Another problem is posed by factive propositional attitude verbs, such as ‘realize’, ‘know’, ‘regrets that’, etc. Consider, for instance, my utterance at t of:

(6) John realizes that this chili tastes great

(6), it seems, is true only if it is true at a circumstance that has me as the judge and a circumstance that has John as the judge. But it is straightforward to explain why we get this kind of double-indexing. Propositional attitude verbs create shifted circumstances. ‘John believes p ’ is true only if it is true at a circumstance at which John is the judge. But where O is a factive operator, Op is true at a circumstance only if p is true at that circumstance. So, ‘John realizes that’ creates a shifted circumstance at which John is the judge. But ‘John realizes that’ is factive. So, (6) is true at the $\langle @, t, \text{me} \rangle$ circumstance only if ‘this chili tastes great’ is true at that circumstance.

A third problem is that there seem to be cases where the default value of the judge parameter is not the speaker.²⁶ There are two kinds of problematic cases. The first kind of problematic case is that of generics’.²⁷ On my account, it seems, a generic such as ‘if the wine is disgusting, people should spit it out’ is false if I dislike the wine, hardly a satisfactory result. The second kind of problematic case is constituted by cases of persuasion and cases of reporting. For

²⁶ Lasersohn, “Context Dependence, Disagreement, and Predicates of Personal Taste”, pp. 672ff.

²⁷ Thanks here to Kent Bach.

instance, if you refuse to taste the chili I just made, I might say ‘c’mon it’ll taste great’. What I say, it seems, is true if it is true by your standards, or perhaps by both of our standards. Or if someone asks me how Mary liked the rides at the theme park, I might reply ‘the roller coaster was fun, but the free fall was a bit scary’.²⁸ Here again, it seems that what I said is true just in case the roller coaster was fun by Mary’s standards, and the free fall was a bit scary by Mary’s standards. So, in these cases it seems that the default value of the judge parameter is not the speaker.

However, I think these cases are instances of what Recanati has called ‘free circumstance shift’. A circumstance shift is free just in case it is not operator controlled but is controlled instead by the speaker’s intentions. Here is a test of free shifts of the judge parameter. There is free circumstance shift just in case inserting an appropriate circumstance shifter in the front of the sentence does not change the truth-value of the sentence. As ‘For all x , if the wine is disgusting according to x , x should spit it out’, ‘c’mon it’ll taste great, according to you’ and ‘according to Mary, the roller coaster was fun, but the free fall was a bit scary’ have the same truth-values as the originals, these cases are plausible cases of free circumstance shifts.

Circumstance shifts are quite common in narratives. Consider, for instance:²⁹

- (7) John is completely nuts. Nothing tastes great anymore, not even chocolate ice cream.

Inserting ‘according to John’ at the front of the second sentence will not affect the truth-value in this case either.

My working hypothesis then is that sentences containing perspectivals are true relative to circumstances of evaluation that contain not only a world and a time parameter but also a judge

²⁸ Lasersohn, “Context-Dependence”, p. 672.

²⁹ This example is based on Recanati’s example: ‘John is totally paranoid. Everyone wants to kill him, including his own mother’. *Oratio Obliqua, Oratio Recta*. Cambridge, (MA: MIT Press, 2000).

parameter. The default value of the judge parameter is the speaker, but circumstance shifters that operate on the value of the judge parameter (e.g., ‘according to John’) can shift the value.

Since the default value of the judge parameter is determined by the speaker parameter of the context of *use*, perspectivals are context-sensitive. But they are not context-sensitive in the same way as indexicals. Where indexicals have variable values, perspectivals have variable extensions.

The proposed semantics explains the same linguistic data as relativism but without relativizing sentence truth to a context of use and a context of assessment. The relativist takes the truth of any sentence, including unembedded sentences, to be relative to a context of use and a context of assessment. So, if I assert ‘this chili tastes great’, what I asserted has no truth-value, except relative to a context of assessment. Perspectivalism, on the other hand, takes unembedded sentences to have a truth-value relative to the context of use. So, if I assert ‘this chili tastes great’, then what I said has a truth-value relative to my context of use. Since the revisions to standard semantics required by perspectivalism are far less radical than those required by relativism, perspectivalism is preferable to relativism on methodological grounds.

Though I will be unable to argue for it here, it is quite plausible that any candidate for being a genuinely relative expression can be adequately accounted for as a perspectival. This is good news for the epistemic value monist, or anyone who is unconvinced by the trivial argument. For the hypothesis that there are perspectivals in the language does not create a demand for epistemic value pluralism. Perspectivalism is consistent with the following formulation of the epistemic goal:

(EG''''') For any proposition p that you are able to consider and any time t , the goal is for you to believe that p at t if and only if it is true that p at the circumstance determined by your context of use C_U at t

As (EG''''') is a very meek revision of (EG'), the hypothesis that there are perspectivals is compatible with truth being the highest epistemic goal.

V. Conclusion

Relativism provides a compelling explanation of linguistic data involving predicates of personal taste (e.g. 'tasty' and 'fun'), moral expressions (e.g. 'right' and 'wrong'), future contingents (e.g. 'there will be a sea battle tomorrow'), color expressions (e.g. 'is the same color as'), epistemic modals (e.g. 'might'), and so on. But relativism is incompatible with epistemic value monism—the view that truth is the highest epistemic goal. The argument from relativism to epistemic value pluralism I called *the trivial argument for epistemic value pluralism*, because it threatens to show all too swiftly that epistemic value monism is false. I have argued that there is no way of reformulating epistemic value monism in such a way as to take account of the assessment-sensitivity of the target expressions. Nor can we simply dismiss relativism without providing an alternative. I have argued, however, that a version of non-indexical contextualism can account for the same data as relativism without relativizing sentence truth to contexts of assessment. On this account, sentence truth is relative only to a context of use, but the environments in which sentences are evaluated create shifted contexts at which to interpret the evaluated sentences.³⁰

³⁰ Thanks to Kent Bach, Michael Glanzberg, Jon Kvanvig, Ernie Lepore, Michael Lynch, John MacFarlane, Julien Murzi, Joe Salerno, Jonathan Schaffer, Jason Stanley, and the audience at the Epistemic Value Conference in Stirling for discussion of these and related issues. Special thanks to my commentator Mikkel Gerken for comments that much improved the paper and to Duncan Pritchard for organizing a great conference.