A Case For Anti-Realism About Knowledge

Alexander Jackson

18th August 2009

Abstract: Existing responses to the lottery and preface puzzles are deeply problematic. We need a new metaphysical notion to solve the puzzles. As I explain, whether a knowledge-attribution is 'metaphysically acceptable' depends in part on what is psychologically salient to the judge. This notion is also directly validated by the following two intuitions: that the 'right' thing for you to judge about whether curry is delicious depends on whether curry tastes good to you; and that you 'can' go either way on a borderline case. I outline an approach to metaphysics on which what matters is what makes it metaphysically acceptable to judge that p, not whether (or why) it is true that p. I show how this metaphysics is compatible with a standard truth-conditional approach to compositional semantics. An appendix extends my approach to deal with puzzles about 'pragmatic encroachment' into our knowledge-attributions.

1. Introduction.

We need a new metaphysical notion to solve the 'lottery' and 'preface' puzzles about knowledge. Let me introduce the notion using the case of deliciousness. My Dad judges that curry is not delicious, and I disagree. I think my Dad's judgement is false – curry is delicious! But that's irrelevant to the metaphysics of deliciousness. In the sense that's important to metaphysics, I think he is 'right' to judge that curry is not delicious, because it does not taste good to him. I am 'right' to judge that it is delicious, because it does taste good to me. Give this sense a label: though they are incompatible, both judgements are 'metaphysically acceptable'. That's possible because deliciousness is subjective. Whether it is metaphysically acceptable for someone to judge that curry is delicious depends on whether it tastes good to them, i.e. on the mental state and dispositions of the judge. I will argue that the lottery and preface puzzles show that contrary judgements about knowledge can both be 'right'. Whether it is metaphysically

acceptable for someone to attribute knowledge depends partly on what's psychologically salient to the judge. Knowledge is partially (and only partially) subjective.

I set up the puzzles about knowledge in sections two and three. We normally think that people know a lot of things. We can be pushed to deny almost any particular knowledge attribution by 'lottery' or 'preface' considerations. We cannot plausibly reject the skeptical judgements reached in those ways; nor should we have a stable theoretical view that people know very little. Both the initial knowledge attribution, and the denial of knowledge made once the judge has been subjected to the lottery or preface considerations, must be 'right' in the theoretically important sense. Existing attempts to spell this out are deeply problematic. Apart from the argument against truth-relativism, my sketch of the dialectical terrain summarizes the current literature with a broad brush. The main goal of this paper is to show how to satisfy the familiar constraints on a solution to our puzzles.

Section four argues that there is the required sense in which a judgement can be 'right', that's different to the judgement's being true and to its being epistemically justified. The notion gets direct intuitive support from the case of deliciousness, and from our 'permissiveness' about judgements regarding a borderline case. Section five puts this 'metaphysical acceptability' at the heart of metaphysics. I reject the ideology of grounding the truth of p, replacing it with what makes it metaphysically acceptable for J to judge that p. This allows us to understand the commitments of realism and anti-realism about a subject-matter. I have given an anti-realist theory of knowledge. This metaphysics doesn't require revisions to our theory of meaning. It is compatible with a standard implementation of truth-conditional semantics.

Section six applies the metaphysical framework, giving the detailed claims about knowledge that resolve the lottery and preface puzzles. An appendix to the paper extends this theory to solve puzzles arising from 'pragmatic encroachment' into our knowledge attributions.

2. The Lottery and Preface Puzzles.

This section gives working examples of the lottery and preface puzzles. I describe the common-sense reactions, and formulate principles that capture how they generalize. This

section and the next introduce and briefly motivate intuitive constraints on our theory familiar from the work of Timothy Williamson (2000), John Hawthorne (2004) and others.[1] My goal is not to give an exhaustive defence of those constraints. My main goal is to show how to respect them. The unsympathetic reader can construe my main claim as conditional: given that the intuitions are as I describe, my theory is the best way to accommodate them. Many epistemologists end up rejecting some of those intuitions because they think they can't all be respected. It should not be controversial that we should try to accommodate all the intuitions. This paper shows how to do so.

We normally think people know lots of things. For example, if Sally drove to work this morning, we think that at 5pm she knows that her car is parked round the corner from her office. But when we consider it, it is natural to think that Sally doesn't know that her car hasn't been stolen since she parked it. As a result, we judge that Sally doesn't know that her car is parked round the corner. Take this change in judgement about whether Sally knows as our working example of the lottery puzzle. (The lottery puzzle is so-called because of the instance where you give up the claim that Irene knows she won't be able to afford a vacation in Barbados next year, on the grounds that she doesn't know she won't win the lottery in the meantime. Sally's example makes it more obvious that the template generalizes.)

Plausibly, the retraction manifests our commitment to Single Item Closure (SIC).

(SIC) If S knows that P, and S knows that if P then Q, then S is in a position to know that Q.[2]

We accept that Sally can't know her car hasn't been stolen, and it's obvious that if it is parked round the corner then it hasn't been stolen. We react by judging that Sally doesn't know her car is round the corner. Think of SIC as labeling that intuitive response.

---

[1] As will be obvious, my understanding of the puzzle bears a special debt to John Hawthorne 2004.

[2] This traditional formulation of closure is neutral on issues of epistemic priority, unlike those of Hawthorne (2004 p. 34) and Williamson (2000 p. 117). It does not say that [P] and [if P then Q] are premises from which S can come to know that Q. Hence I do not call it Single *Premise* Closure.

Suppose I resist the conclusion by insisting that Sally does know that her car hasn't been stolen. The following manoeuvre forces me to change my mind. I can't accept that: there's a chance the train will be late and I know it won't be. Admitting that there's a chance the train will be late forces me to accept I don't know it will be on time. An analogous argument can be run in the third-person case. How likely is it, given Sally's evidence, that her car has been stolen? We can't say that there is zero chance that her car has been stolen, given her evidence. After all, Sally knows that cars are stolen in such circumstances from time to time. But if there is a chance that her car has been stolen, given her evidence, then she does not know that her car hasn't been stolen. And so she doesn't know that it is parked round the corner.

The Epistemic Possibility Contraint (EPC) labels our acceptance of the first inference in this kind of argument.

> (EPC) If there is a non-zero chance that p is false, given S's evidence, then S does not know that p. [3]

We don't need a theory of the relevant epistemic notion of chance to feel the intuitive force of instances of EPC. Many epistemologists give up these intuitions, because EPC seems to lead straight to skepticism. Yet the intuition is real, and should be accommodated if possible. David Lewis (1996) aims to preserve a similar infallibilist intuition (p. 550). John Hawthorne emphasizes the attractiveness of EPC (2004 pp. 24-28). Timothy Williamson (2000) is committed to such a principle on the basis that all knowledge is evidence (pp. 205-6), and defends its plausibility (pp. 249-251). Many philosophers working on epistemic possibility respect the intuition that if the train might be late, then I don't know it won't be.[4] The threat of skepticism does not deprive that intuition of its prima facie force.

The attack on our initial verdict that Sally knows her car is round the corner generalizes to many things we ordinarily think people know.[5] So we have a puzzle. Does the above reasoning bring out the fact that people actually know very little? Do people

---

[3] Compare Hawthorne 2004 p. 111.

[4] Amongst others: DeRose 1991; Egan, Hawthorne and Weatherson 2005; Yalcin 2007.

[5] See Vogel 1990 and Hawthorne 2004 pp. 1-7.

know lots of things, and we go wrong when we are led to deny this? Those unpalatable options seem to be the only possibilities. That's what is puzzling about our changing verdicts about whether Sally knows.

I give two versions of the 'preface puzzle': the 'epistemic probability' version, and then the 'closure' version.[6] It is November. John has the list of the 100 people who have paid to attend the philosophy conference he has organized for February. We are normally happy to judge that John knows that Ally will be at the conference, because he sees her name on the list. On similar grounds, John knows that Bert will be at the conference; etc., for the 100 people on the list. But we don't think that John can know that all 100 people on the list will be at the conference – it is very likely given his evidence that at least one person will get sick in February. (At past such conferences, two or three people have always cancelled.) The natural reaction is to conclude that John doesn't actually know that Ally will be at the conference, etc. After all, the large chance that someone will get sick is the result of the small chances that Ally will get sick, that Bert will get sick, etc. (Suppose John has no special information about the health of the people on his list.) As we noted above, even a small non-zero chance that p is false seems incompatible with knowing that p.

Here's the 'closure' argument that we should retract the claim that John knows that Ally will be at the conference. John cannot assert: "It is likely someone will get sick; Ally won't." The former claim should force him to withdraw the latter. He should admit he doesn't know that Ally won't get sick. More carefully, suppose Ally makes it to John's conference in February. We can't then report: Last November it was likely given John's evidence that at least one person would get sick and drop out, and he already knew Ally was not going to. Intuitively, the first claim is correct and forces us to deny the second.

Plausibly, this manifests our commitment to Multiple Item Closure (MIC).[7]

---

[6] I follow Hawthorne 2004 pp. 46-50. Compare Christensen 2004, chapter 3.

[7] You won't get the intuition if you merely think first that John knows that Ally will be at the conference, etc., and then think that John doesn't know that everyone will be. If one fails to reconsider whether John knows Ally will be at the conference, it will look like a

(MIC) If S knows that P1, S knows that P2, …, and S knows that Pn, and S knows that (if P1&P2&…&Pn then Q), then S is in a position to know Q.

Suppose John's knowing the conjuncts were perfectly compatible with his not knowing the conjunction. Then it would be true that: it was likely given John's evidence that at least one person would get sick, and he already knew Ally wouldn't. But that isn't right. Given that John had no differentiating information about the health of the people on his list, those claims are incompatible. So knowing the conjuncts is not compatible with being unable to know the conjunction. In other words, MIC is true. John can't know the conjunction, so he doesn't know all the conjuncts. By symmetry, he doesn't know any of them.

The preface puzzle is so-called because of the case where the preface of a book asserts that the main text probably contains several errors. How can the author go on to make the assertions that constitute the body of the book? As John's example makes clear, the puzzle generalizes. Take a large number of things that you think someone knows in similar ways. Usually, you'll think that they can't know the conjunction. So (by either of the above two arguments) they don't know all of the conjuncts, in contrast to your initial verdicts. If the subject is in a similar epistemic position to all the conjuncts, then by symmetry they don't know any of them. As with the lottery puzzle, rejecting the skeptical reasoning is implausible, but so is embracing widespread skepticism. I propose a new way to sail between Scylla and Charybdis, avoiding the problems with the approaches seen in the literature.

3. Three kinds of Response.

This section sketches the space of possible responses to the lottery and preface puzzles. I show prima facie problems with the options currently articulated. (I do not

---

counter-example to MIC. That's the reaction of David Christensen 2004, chapter 3. It's the same mistake as thinking first that Sally knows her car is round the corner, then that she doesn't know it hasn't been stolen, and concluding that SIC is false. In both cases, one must reconsider one's initial knowledge-attribution to get the skeptical intuition, manifesting commitment to SIC or MIC.

pretend to have given the last word on their tenability.) Further constraints on a solution emerge, which will motivate my positive view. The main purpose of this paper is to explain how the desiderata I describe can all be respected.

Responses to our puzzle fall into three broad categories. The first concludes that nobody knows very much. Call this 'skepticism'. The second is that people do know a lot, and there is something wrong with the reasoning that leads us to judge that they don't. Call this 'dogmatism'. (This has nothing to do with James Pryor's 'dogmatist' theory of perceptual justification (2000).) The third kind of response is that you were in the most important sense 'right' to assert that Sally knows that her car is parked round the corner, but then 'right' to deny it in response to the reasoning outlined above. Let's examine these responses in order.

3a. Skepticism.

The skeptical response is that Sally simply does not know that her car is parked round the corner. We were just wrong to think that she did. This conclusion generalizes as much as the puzzle does. People know very little.

Accepting the moral generalized from a specific case is a significant step. Just noticing in the abstract that one could apply similar reasoning to Heather's putative knowledge that she will do her laundry tomorrow does not make it clear that she does not know it. Actually thinking of ways in which the future could go differently is necessary for generating the judgement that she does not know. And of course, we soon go back to attributing lots of knowledge to people, even the very claim that we rejected in response to the skeptical reasoning. Temporarily accepting that Sally doesn't know does not automatically commit us to a stable theoretical position that nobody knows anything.

What's wrong with skepticism? One might complain that it is intuitively obvious that people know a lot of things. This bare appeal to intuition is weak when applied to a specific example. When we go through the reasoning, the intuitive verdict is that Sally does not know.[8]

[8] This is not the case with Brain-In-a-Vat skepticism, which unfortunately is often run together with the lottery puzzle.

I explain the repulsion from skepticism by appeal to the point of the concept of knowledge, namely the normative importance of knowing. Knowledge is the epistemic Norm of Assertion (KNA). (See Unger 1975 chapter 6, DeRose 1996, and Williamson 2000 chapter 11.)

> (KNA) One must: assert that p only if one knows that p. (Williamson 2000 p. 243)

John Hawthorne endorses KNA and the Practical Reasoning Constraint (2004 pp. 21-31).

> (PRC) One ought only to use that which one knows as a premise in practical reasoning. (Hawthorne 2004 p. 30)

Jason Stanley holds,

> One should act only on what one knows. (Stanley 2005 p. 9)

Hawthorne and Stanley (2008) defend the Action-Knowledge Principle (AKP).

> (AKP) Treat the proposition that p as a reason for acting only if you know that p.

Interpret these principles as giving the epistemic standard we aim to meet in asserting, etc.[9] Knowing is epistemically necessary and sufficient for asserting or acting on a belief.

Here are two of Williamson's arguments for KNA. Firstly, you can't assert: "Though I don't know it, p." The best explanation is that KNA is true, so that assertion must be improper. Secondly, you can't assert your ticket won't win the lottery. Intuitively, that's because you don't know it won't. This explanation suggests that knowledge is the epistemic standard for assertion.

Here's an argument for PRC and AKP (from Hawthorne 2004 pp. 29-31, 147). Bill reasons: "My lottery ticket won't win; so it is worthless; so I should sell it even for one cent." That conclusion is unacceptable, yet every step is valid. Intuitively, the reasoning is bad because Bill does not know his ticket won't win. If he did know it, the reasoning would be good. This suggests that knowledge is the epistemic standard a premise should meet.

---

[9] I explain in section 8 of my Forthcoming why this does not render justified belief normatively irrelevant. We should generally distinguish the 'should' of success from the 'should' of justification. The knowledge-involving norms are those we aim to meet, and thus define the conditions for success, not for being justified.

Let's see how this affects the case at hand. We judge that Sally does not know her car is parked round the corner, because it might have been stolen. As Sally knows, it happens about once a year that someone she works with has their car stolen during the day. Sally and Jennifer discuss where to go for dinner. Sally says, "My car is parked round the corner, so let's drive to the diner on route 1." Sally should not flat-out assert to Jennifer that her car is parked round the corner, because she doesn't know it. She should say it's highly likely. Sally shouldn't really reason as she does, because she does not know her premise to be true. She should reason from the fact it is highly likely her car is round the corner. KNA and PRC get this exactly right. It is very plausible that those normative judgements go along with the view that Sally does not know that her car is parked round the corner. The reader might think, "Oh come on, Sally can assert her car is round the corner, and use it as a premise!" That's the same mood in which one responds, "Oh come on, she knows it is parked round the corner!" This supports KNA and AKP.

Obviously there is a lot more to say here, but let's work on the basis that AKP, KNA, etc, are true.[10] This broadens our task from just explaining the linguistic data about our use of the word "knows". It requires we understand normative issues faced by every agent and communicator. If those principles are true, then skepticism entails that nearly all our assertions are out of order, as is nearly all of our reasoning. That can't be right. Skepticism is bad because it cripples our understanding of people's normative situation. We feel unhappy enough having been put through one instance of the lottery puzzle and thus judging that Sally is violating the epistemic norms of assertion and action. Luckily we soon forget that we had been forced to say that Sally doesn't know that her car is parked round the corner. We should not accept the general and stable theoretical position that nearly all our assertions and reasonings are out of order. That's a bad picture of how humans should conduct their business.

[10] See Fantl and McGrath 2007 for more defence, and Brown 2008 for counter-arguments. In the last paragraph of section six, and in the appendix, I argue that knowledge behaves in exactly the way we would expect of the epistemic standard for premises and assertions.

3b. Dogmatism.

Dogmatism is the view that, for example, Sally does know that her car is parked round the corner, and we went wrong by reasoning to the conclusion that she doesn't. In other words, it is easy to get people into a mistaken skeptical frame of mind. This is an unattractive position. It must deny some intuitively compelling claims used to generate the puzzles. For example, there are two dogmatic ways to respond to the preface puzzle as generated by considerations of probability. The first holds there is zero probability given John's evidence that at least one of the people who have signed up for the conference will get sick. That's clearly wrong, so let's explore the second strategy. This holds that there is a good chance given John's evidence that at least one person will get sick. So on this view there is a non-zero chance that Ally will not be at the conference (etc), but John still knows that Ally will be there. It rejects the Epistemic Possibility Constraint. The natural extension to the lottery is that Sally knows that her car hasn't been stolen, even though there is chance it has been, given her evidence. Having made this concession, the dogmatist can keep Single Item Closure. But Multiple Item Closure must go as well as EPC. Clearly John cannot know that none of the 100 will get sick. For his evidence makes it quite likely that someone will get sick. (See Christensen 2004 chapter 3 for more on the costs of combining dogmatism with MIC.) But MIC captures a real intuition.

More bad news: dogmatism is incompatible with the normative role of knowledge. We retract our judgement that Irene knows she won't be able to afford a vacation in Barbados next year, because we judge that she doesn't know that she won't win the lottery in the meantime. Presumably, the dogmatist view is that she does know that she won't win the lottery, thus preserving SIC. But Irene should not reason: my ticket will lose, so it is worthless, so I should sell it even for one cent. The natural explanation is that she does not know that the ticket will lose. The dogmatist can't accept that explanation, as on that view Irene does know that her ticket will lose. That's the cost of denying that knowing that p is sufficient for it to be an acceptable premise in practical reasoning, as the dogmatist must.

3c. Have Your Cake and Eat It.

First we judged that Sally knows her car is parked round the corner. Then we judged that she doesn't. Skepticism and dogmatism agree that one of those judgements is simply right, and the other is simply wrong. The third kind of response to the puzzle is that both judgements were 'right', given the different situations in which they were made. Call this the 'have your cake and eat it' strategy. Given the implausibility of skepticism and dogmatism, it seems the most promising avenue to investigate.

The literature contains three attempts to implement this strategy, namely contextualism, Subject-Sensitive Invariantism (SSI), and truth-relativism. Sub-sections 3c.i-iii sketch problems with them.

3c.i. Contextualism.

Contextualism is the view that the word "knows" refers to different relations in different contexts of use.[11] Thus utterances of the sentence "Sally knows her car is round the corner" can express different propositions, and have different truth-values. First we utter, "Sally knows her car is round the corner." After the matter of car-theft has been made salient we utter, "Sally doesn't know that her car is round the corner." The contextualist diagnosis of the puzzle is that these utterances can both be true, because they do not contradict each other. Contextualism implements the 'have your cake and eat it' strategy by holding that both utterances are 'right' in the straightforward sense of being true.

There are big problems for contextualism if we want all utterances of the sentences MIC, EPC, KNA and AKP to be true.[12] It looks like all the bad things the dogmatist says about knowledge the contextualist says about the relation referred to by "knows" in the initial context. Suppose the contextualist accepts all utterances of MIC. Then given the initial referent of "knows", John is in a position to "know" that none of

[11] E.g. Cohen 1988, DeRose 1995, Lewis 1996.

[12] See Hawthorne 2004 pp. 80-98, 179. Any speaker should utter truly, "John is in the same epistemic position with regards to whether Ally will be at the conference as whether any other of the people on his list will be there." This rules out the strategy for accommodating MIC offered by Hawthorne pp. 97-8.

his 100 friends will get sick. Yet it is very likely given John's evidence that someone will get sick. How can that be a paradigmatic sense of the word "know"?

Let's put those problems aside, and turn to the most straightforward objection to the contextualist resolution of our puzzles. That solution says that we assert and then deny different propositions. That's how both utterances can be true. But this does not fit the data. If that solution were correct, we would respond correctly to lottery cases as follows.

"Sally knows her car is parked round the corner."

- "Does she know it hasn't been stolen since she left it there?"

"I suppose not. So Sally doesn't know that her car is round the corner. But that's not to say that I spoke falsely a moment ago."

That's not how we respond. As Williamson puts it (2005 p. 220), the intuitively correct reaction is:

"So I was wrong. Sally doesn't know that her car is round the corner after all."

According to contextualism, you weren't wrong. "After all" implies you are changing your mind, and contextualism says you aren't.

There is a strong intuition that in our puzzle cases, we assert and then deny the very same proposition. The contextualist solution to the puzzle denies this. So there is strong reason to reject the contextualist solution. From now on, I will assume that our puzzles concern contrary attitudes to the same proposition.

(Rejecting the contextualist solution to our puzzles does not entail rejecting contextualism about the word "knows". One could posit contextual variation only where it is not obvious that the two utterances contradict each other. Keith DeRose (2006) seems to take this path. This gives up the contextualist solution to the lottery and preface puzzles. For in these cases, the intuition that we changed our minds is robust.)

3c.ii. Subject-Sensitive Invariantism.

There is good reason to have our cake and eat it concerning contradictory attitudes to the same proposition. This constraint on a solution motivates Subject-Sensitive Invariantism (Hawthorne 2004 chapter 4). On this view, utterances of "knows" always refer to the same relation. The distinctive thesis is that whether the subject knows

something depends on facts about the subject that are not traditionally thought of as epistemic. In the case at hand, whether car-theft is salient to Sally can affect whether she is in a position to know that her car is parked round the corner (whether or not she continues to believe it). If car-theft becomes salient to Sally, she may cease to be in a position to know where her car is. This might help with self-ascriptions of knowledge. But it has no hope of dealing with third-person ascriptions of knowledge, like those I set out in section two. Nothing changes about Sally or John when we go through the lottery or preface reasoning and change our minds about whether they know. SSI will have to give up EPC and MIC, just like dogmatism.[13]

3c.iii. Truth-Relativism.

Truth-relativism is a more radical attempt to capture the idea that it was 'right' to assert, and then 'right' to deny, the very same proposition. Mark Richard 2004 and John MacFarlane 2005 give truth-relativist solutions to the lottery puzzle. I will argue that this strategy cannot be applied to the preface. On this view, truth is relativised to a possible world centered on a judge and a time.[14] Suppose Polly is subjected to our example of the preface puzzle. Let 'Polly1' designate the pair of Polly and the time at which she makes her initial judgement that John does know Ally will be at the conference. Let 'Polly2' designate the pair of Polly and the later time at which she judges that John doesn't know Ally will be there. Our truth-relativist holds that there is one proposition Polly changes her mind about, namely: <John knows that Ally will be at the conference>. This avoids the objection to contextualism. The puzzle is dissolved because that proposition is true-for-Polly1 but false-for-Polly2. The alleged sense in which both judgements are 'right' or 'correct' is that they are both true relative to the judge-time pair at which they were made.

Being true-for-that-judge is only a metaphysically serious sense of 'correct' if LINK is true.

[13] Hawthorne 2004 p. 162-6 tries to explain away the third-person intuitions. Williamson 2005 argues that this removes any dialectical advantage SSI has over dogmatism.

[14] See also Kolbel 2003; Egan, Hawthorne and Weatherson 2005; Yalcin 2007.

(LINK) If p is true-for-J, then p is not completely at odds with the underlying facts.

For example, it can't be true-for-someone that I know what numbers will win next week's lottery. That would be a crazy take on my epistemic situation. Again, there are different legitimate ways of categorizing things are red. But it must be true-for-everyone that London buses as red – that's the only assignment of relative truth that even minimally respects the underlying facts. (Don't worry that LINK is vague. It's just a label for a kind of step I will make in the argument below. The reader should be able to assess those steps on their own merit.)

Unfortunately, the truth-relativist solution to the preface puzzle is incompatible with LINK being true. The strategy is to endorse Polly's reasoning from the premise that John doesn't know everyone will be at the conference, to the conclusion that he doesn't know Ally will be there. In reasoning in that kind of way, Polly commits herself to CK.

(CK) If S knows that P1, …, S knows that Pn, and S is in a position to know that {P1,…, Pn} entails Q, then S is in a position to know that Q.

It would be incoherent of Polly to systematically accept any instance of the principle, but reject the generalization. It must always be 'correct' for her to accept CK, so CK is true-for-Polly1. (Why would you have to accept CK if it were false-for-you? Nor does it help to assign CK a third truth-value relative to Polly1.) I will now argue that this means it is true-for-Polly1 that John is in a position to know that no-one will get sick and miss the conference. But that's completely at odds with the underlying facts: John's evidence points strongly to at least one person getting sick. So LINK is false.

Relative truth is closed under entailment. Otherwise, the thing we have relativised is not truth.

(CLOSED) Suppose P1 is true-for-J, …, Pn is true-for-J, and {P1, …, Pn} entails Q. Then Q is true-for-J.

Given that relative truth is closed under entailment, the following argument is valid.

1. It is true-for-Polly1 that John knows that Ally will be at the conference.
2. It is true-for-Polly1 that John is in the same epistemic position with respect to everyone on his list.
3. It is true-for-Polly1 that John knows who is on the list.

4. CK is true-for-Polly1.

So,

5. It is true-for-Polly1 that John is in a position to know that no-one on his list will get sick and miss the conference.

Suppose for reductio that LINK is true. By LINK (roughly), 2 and 3 are true. For it would be completely at odds with the underlying facts to assign falsity (or a third truth-value) to the proposition that John is in the same epistemic position with respect to everyone on the list. To minimally respect the reality of John's epistemic situation, an assignment of relative truth must make 2 true, and similarly for 3. 1 is essential to the truth-relativist resolution of the puzzle, and I argued above that 4 would have to be true. The truth of 5 follows by CLOSED. That's a reductio of LINK, because it is completely at odds with the underlying facts for an assignment of relative truth to classify John as in a position to know no-one will get sick in February. So using truth-relativism to solve the preface puzzle commits you to LINK being false.

This undermines the truth-relativist resolution of the puzzle, which depends on the truth of LINK. If a crazy proposition is true-for-Polly1, then being true-for-Polly1 can't be what makes her judgement 'correct'. Without LINK, the truth-relativist hasn't given a serious sense in which both of Polly's judgements were 'right'. Further, relative truth is an honorific. I suspect it is incoherent to hold that *truth* that is relative, and yet that propositions that fail to minimally respect the underlying facts are true-for-someone.

Of course, one might apply truth-relativism to the lottery, but not the preface. Such a truth-relativist would bite the bullet on MIC and EPC. This strikes me as a dialectically weak position. Resorting to a bold metaphysics, but getting only half the job done, is an unsatisfactory combination. Further, we want to explain why we judge we can't know a given ticket won't win the lottery, and intuitively, MIC and EPC are the reasons why.

4. Metaphysical Acceptability.

Section three showed that there must be a sense in which our initial acceptance and subsequent denial that Sally knows are both 'right'. What's that sense of being 'right'? We saw in discussing contextualism that those two judgements contradict each

other. So they can't both be true. Truth is not the sense of a judgement's being 'right' that will resolve our puzzles.

Nor can 'right' mean that the judgement is epistemically justified. As we think about the lottery case, we are forced to judge that Sally doesn't know her car is round the corner. So we must accept that our first judgement was false. But the deeper insight, by which we avoid pernicious skepticism, is that the first judgement was 'right' nevertheless. We don't get this result by saying that skepticism is true, but people are often epistemically justified in making false knowledge-attributions. Similarly, it is no alternative to skepticism to hold that, while nobody knows very much, false knowledge attributions are often pragmatically assertible. We need the fact that both judgements are 'right' to render philosophically irrelevant the question of which is true. We have been forced to say that our first judgement was false. To escape skepticism, we must say that that doesn't matter – what matters is that both judgements were 'right'. Epistemic justification and pragmatic assertibility can't play this role.

> We need to be able to make the following speech.
>
> Sally doesn't know her car is round the corner. But that's merely what's 'right' for me to judge in my current situation. In my previous situation, it was 'right' for me to judge that Sally does know. Neither of those situations, and neither of my judgements, is metaphysically privileged.

The truth-relativist says that both judgements were 'right' in the sense of being true relative to the (judge, time) pair at which they were made. I argued against this characterization in section 3.c.iii. Instead, I propose that we take the relevant sense in which a judgement is 'right' as basic in our metaphysical theorizing. Label this sense the 'metaphysical acceptability' of someone's judging that p. There are three stages in my explanation of this notion and the theory it features in.

Firstly, I give two cases in which we have a common-sense intuition that such a sense of being 'right' is at work. Sections 2-3 were a long argument that there must be such a sense in which conflicting judgements about knowledge can both be 'right'. By contrast, there is a direct intuition that conflicting judgements on a matter of taste can both be 'right'. There is a direct intuition that conflicting judgements on a borderline case can both be 'right' or 'acceptable'. These cases show that we are dealing with a

widespread phenomenon. They give us an intuitive grasp on how the theory is meant to work. Further, if a notion appears in an intuition, that notion is intuitive. So metaphysically acceptable judgement is an intuitive notion (even if the label is unfamiliar).

Secondly, I argue that such a theory predicts the right results. This is easiest to see in the case of matters of taste, such as whether curry is delicious. The notion of metaphysically acceptable judgement can be grasped by understanding the theory it features in, as well as via intuitions about 'faultless disagreement'.

Thirdly, I show in section 5 how to make metaphysical acceptability the central metaphysical notion. Roughly, it replaces the notion of a truth being grounded in others. I define subjectivity and objectivity, and show the link with how we should think of reality-as-it-is-in-itself. I explain why my approach is compatible with truth-conditional and referential approaches to compositional semantics. Indeed, we need an approach like mine if we are to free linguistics from metaphysical considerations. Section 5 zooms out, showing us the attractive panorama built from the materials provided by section 4.

Let's start with the intuition that there can be 'faultless disagreements' about matters of taste. For example, my Dad does not like curry. He thinks that curry is not delicious. I disagree. I think curry is delicious, and that my Dad's judgement on that matter is false. Yet it is not a serious metaphysical position that I am locked on to the nature of deliciousness, and my Dad isn't. His judgement is false, but it is the 'right' verdict for him to give, because curry does not taste good to him. It does not do my Dad metaphysical justice merely to say that his false judgement is epistemically justified.

Contextualists about the word "delicious" reject this intuitive picture. They say that an utterance of "Curry is delicious" means that curry tastes very good to an appropriate proportion of a group of people determined by the context of utterance. On this view, if my Dad and I disagree, we disagree about whether curry tastes delicious to most people in some group. That's implausible. My Dad is painfully aware that most (British) people love curry. He bemoans the fact that most people these days like "foreign muck," rather than proper food like fish and chips or cottage pie. So our disagreement does not stem from my Dad having a false sociological view. It stems from the fact that my Dad doesn't like curry, and so thinks it is not delicious. (More generally, we can

suppose that my Dad and I have the same information about which people really like curry.) So it seems that contextualism is unsatisfactory. (See Cappelen and Hawthorne 2009 chapter 4 for an attempt to undermine this intuition.[15]) As with the data about knowledge, space restricts me to establishing the prima facie force of the intuitions. I am motivating and explaining my position, not decisively refuting the alternatives.

Intuitively, I can recognize that my Dad is 'right' to think that curry is not delicious, even though I disagree with him about that. If a notion appears in an intuition, it is intuitive. So that sense of a judgement's being 'right' is an intuitive notion. Label it 'metaphysical acceptability', to distinguish it from epistemic justification and pragmatic assertibility. The complete metaphysical theory about deliciousness is as follows. It is metaphysically acceptable for someone to judge that x is delicious iff x tastes very good to them; it is metaphysically acceptable for someone to judge that x is not delicious iff x does not taste very good to them.

We want to say that my perspective on deliciousness is not metaphysically privileged over my Dad's. Both our judgements are equally good from the metaphysical point of view. But I think my judgement is true, and my Dad's is false. If I am able to remain metaphysically 'evenhanded', it must be irrelevant to the metaphysics whose judgement is true.[16] It must be irrelevant to the metaphysics of deliciousness whether curry is delicious. The metaphysical theory says what makes it metaphysically acceptable for someone to judge something to be delicious. The theory is silent on whether curry is

[15] Cappelen and Hawthorne 2009 argue that there is an 'exocentric' use, on which someone who dislikes hotdogs can encourage a child who likes them to eat one by saying, "Hotdogs are delicious!" This is irrelevant to my theory, which concerns what judgements are acceptable. The speaker here does not judge that hotdogs are delicious, so there is no challenge to my theory. Cappelen and Hawthorne also worry about intuitions framed in generic terms. But we have the intuitions concerning my disagreement with my Dad about whether this batch of curry is delicious. Further, do not ask whether I really or strongly disagree with my Dad (tout court). Ask whether we disagree about whether this batch of curry is delicious.

[16] This manoeuvre resolves Wright's conundrum about evenhandedness (2008 section 6).

delicious, though I hold that theory and am not silent on the matter. The theory says that the 'right' thing for me to judge is that curry is delicious, but that's not the same thing as the theory saying that the proposition is true. After all, the theory also says that the 'right' thing for my Dad to judge is that curry is not delicious. When deciding what to have for dinner, the relevant question whether curry is delicious. When in the seminar room, it is not.

Curry is delicious – but that's merely the verdict that's metaphysically acceptable for me to give. My Dad's contrary verdict is the one that's metaphysically acceptable for him to give. So from the metaphysical point of view, our contrary verdicts are on a par. (From the gastronomic point of view, my verdict is the right one, because it is true.)

Let me show that this kind of view predicts the right results, by telling a Just So Story. Suppose a tribe plays a language-game whose basic rules are: if something tastes good to you, say it is delish; if it doesn't, say it is not delish. Because the tastes of the speakers largely overlap, this game allows them to coordinate what food to grow, prepare, and eat, by sharing information about what's delish. The tribe extend their language-games by minimalist truth. The basic rules for it are: treat [it is true that p] as equivalent to [p], and treat [it is false that p] as equivalent to [not-p]. So when the speakers' tastes coincide, one says that lasagna is delish, and another says that that's true. When the speakers' tastes diverge, one says that curry is delish, and another says that that's false. Both speakers are judging as the rules of the game require. Both their judgements are 'correct' or acceptable in that sense.

Suppose that the tribe come to understand how this language-game is played, i.e. its basic rules. They see that both sides of a disagreement about whether curry is delish can be 'correct' in the deep sense. You can ask them whether curry is delish, and they will answer – by playing the game according to its rules. A speaker might say:

> Curry is delish – but that's merely the verdict the rules require me to give. The rules require my Dad to give the opposite verdict. The rules of the game endorse both sides of our disagreement, even though they tell me which side to take.

Of course, I think this is how we actually play the game of judging things to be delicious. 'Understanding how the game is played' is one thing; 'playing the game' is another. Figuring out the rules of the game is all there is to doing the metaphysics of

deliciousness. Saying whether curry is delicious is irrelevant to a philosophical understanding of deliciousness. Section 5 defends the generalization: [whether p] is irrelevant to the metaphysics of p; all that's relevant is what makes it metaphysically acceptable for a given person to judge that p.

Part of playing the game is assessing other people's judgements as true or false. I think my Dad's verdict on curry is false, and in that sense wrong. But judging that [my Dad's verdict about curry is false] is merely to make a move within the game, and is not part of understanding how the game is played. We might extend this way of evaluating judgements by saying that they "aim at truth." In the sense that's relevant to gastronomy, my Dad's judgement aims at the truth, and falls short. That's not the evaluation that's relevant to understanding how the game is played, i.e. the metaphysics of deliciousness. In the sense that's relevant there, judgements aim at being metaphysically acceptable, and my Dad's verdict is successful.

Let's move on to the second kind of intuition that contrary verdicts can both be 'acceptable'. For a judgement to be 'right' is for it to be the only acceptable attitude for the judge to have to that proposition at that time. The case of deliciousness shows that it can be metaphysically acceptable for one person to judge that p, but not so for another. Here's a different kind of case that provides direct intuitive support for the notion. We are 'permissive' about judgements about borderline cases. Suppose a car is half-way in colour between paradigmatic red and paradigmatic orange. We think it is acceptable for Jane to judge that the car is red, and acceptable for her to judge that it is not red. Crispin Wright is,

> impressed by the datum that … one is *entitled*, if one is so moved, to a verdict in the borderline area. (Wright Forthcoming p. 5)

Mark Richard (2004) agrees: judges can permissibly disagree about whether winning $1,000,000 makes you rich for a New Yorker.

Either judgement is acceptable. That does not mean either would be true (on pain of contradiction). That does not just mean that either would be epistemically justified (which leaves it open that one of the verdicts misrepresents the fact of the matter). The intuition is that either verdict fits reality well enough. I take this intuition at face value: there is such a notion of acceptable judgement. Saying that both verdicts are acceptable is

meant to let us be evenhanded between them, so it must render irrelevant the question of which is true. Its truth or falsity is irrelevant to the metaphysics of the proposition that Quentin's car is red. Otherwise the metaphysics would not be evenhanded and approve of both.

(I extend the ideology to apply to pairs of attitudes. It is metaphysically unacceptable to: judge that Quentin's car is red and judge that it isn't. Either attitude is acceptable, but not both together. Similarly, it is a top-down fact that it is metaphysically unacceptable to: judge that Quentin's car is red and judge that it is orange. I extend the ideology so that the occurrent attitude of suspending judgement towards a proposition can be metaphysically acceptable or not. By contrast, ignoring a proposition is not to take an occurrent attitude towards it, and is not evaluable in that way. When one suspends judgement on a borderline case, intuitively one does not miss out on a fact of the matter. I capture this by saying that suspending judgement is metaphysically acceptable. I develop a theory of vagueness from these materials in my 2009.)

Plausibly, there is no variation between judges as to which objects it is metaphysically acceptable to judge to be red. Rather, it is metaphysically acceptable for any judge to affirm or to deny that Quentin's car is red. By contrast, there is variation between judges in what it is metaphysically acceptable to judge to be delicious. If curry tastes good to you, you can't go either way on whether it is delicious; affirming is the only metaphysically acceptable attitude. I think that knowledge is analogous to deliciousness, rather than borderline cases. However, vagueness provides a second intuitive case in which we can disagree about whether p, yet take both judgements to be 'acceptable'. I can make the following speech.

> Quentin's car is red – but that's merely the way I have chosen to go on a question that is permissive of either answer. If you judge that it is not red, your verdict is metaphysically acceptable too.

Borderline cases offer direct intuitive support for the notion of metaphysically acceptable judgement, and thus legitimize using the notion to give a theory about knowledge. That's so even though deliciousness provides a closer analogy.

I have given two cases in which it is intuitive that a judgement's being acceptable from the metaphysical point of view comes apart from its being true. One can be

metaphysically evenhanded between a judgement one thinks is true and a judgement one thinks is false. It is not intuitively obvious that this happens in the case of knowledge. Sections two and three form a strong argument that it does. That in a lottery case both judgements were 'right' is a theoretical conclusion, not a direct intuition. Before car-theft became salient, it was metaphysically acceptable for me to judge that Sally knows her car is round the corner; afterwards, it wasn't. (We don't end up in a situation that is permissive of either verdict about whether Sally knows.) This is analogous to its being metaphysically acceptable for me to judge that curry is delicious, but not for my Dad to do so. Whether it is metaphysically acceptable at t for J to judge that p can depend on the mental state of J at t. (One should think it is the same notion of acceptability in all three cases on theoretical grounds. The theory that different notions are in play lacks generality and hence explanatory power, and postulates unnecessary complexity. One should go for the unified theory of section five.)

It doesn't change things to consider whether it is *true* that Sally knows, or true that curry is delicious. Truth is 'minimal'. It is metaphysically unacceptable to have different attitudes towards [p] and [it is true that p]. They must be treated as equivalent. Judging that [it is true that p] has the same level of metaphysical acceptability as judging that [p].

The analogy between deliciousness and knowledge has its limits. The metaphysical acceptability of an attribution of deliciousness depends only on the mental state (or dispositions) of the judge. The metaphysical acceptability of an attribution of knowledge depends partly on the mental state of the judge, and partly on other things. Changes in what is psychologically salient to the judge explain changes in the metaphysical acceptability of judging that Sally knows her car is round the corner. However, if Sally does not believe that her car is round the corner then any judgement that she knows it is metaphysically unacceptable.

The theory does not choose between skepticism and dogmatism. The theory says which knowledge-attributions are metaphysically acceptable. Sometimes it is metaphysically acceptable to judge that Sally knows her car is parked round the corner, and sometimes it isn't. The theory is silent on whether Sally knows. It is silent on whether skepticism is true, and it is silent on whether dogmatism is true. The theory itself

does not choose between skepticism and dogmatism. (It is also silent on whether either skepticism or dogmatism is true – such matters are not relevant.)

Let me explain in two steps how this allows the person who holds the theory to evade the force of the dilemma between skepticism and dogmatism. Firstly, just consider being put through the lottery reasoning about Sally. You end up in a situation where the only metaphysically acceptable verdict is that Sally doesn't know. But you are still in a position to recognize that your previous judgement, that Sally does know, was metaphysically acceptable too. You see that your current judgement is not metaphysically privileged over your previous one. Your situation has changed, changing which 'move' in the game is the right one for you to make at the time. Similarly, the only metaphysically acceptable verdict for me to give is that curry is delicious. Yet I can recognize that my judgement that curry is delicious is not metaphysically privileged over my Dad's contrary judgement. Pressing the theorist on whether Sally knows is the same mistake as pressing the theorist on whether curry is delicious. The theorist can answer, but the answer will be irrelevant to philosophically understanding knowledge or deliciousness. The theorist will make the following speech.

> Sally doesn't know – but that's merely the verdict that's metaphysically acceptable for me to give now. Before car-theft was raised to salience, it was metaphysically acceptable for me to judge that Sally does know. All that's relevant to philosophically understanding knowledge is that I judged acceptably both times; both times I made the right move in the language-game.

(If this seems opaque, the reader might want to meditate on the Just So Story about 'delish'.)

Now suppose one generalizes from the case of Sally, and considers the lottery puzzle. The only metaphysically acceptable attitudes are that SIC, MIC, EPC and AKP are true. Maybe considering the puzzle puts a judge in a situation in which it is metaphysically acceptable to affirm that people know very little. Even if so, this is of little significance. Considering the lottery puzzle puts the theorist in an unusual situation, requiring an unusual 'move' in the language-game. Nevertheless, we recognize that the generous knowledge-attributions we make in more common situations are metaphysically acceptable.

Further, sometimes even the best available verdict is still quite bad. Introducing my theory, I sometimes talked about what the rules of the language-game require. I think it is better to take metaphysical acceptability to be analogous to the intrinsic goodness of an act. In a moral dilemma, all the options are bad (though some might be less bad than others). This can't be captured in terms of what's permitted or required. Our ideology has more expressive power if we take the basic notion to be analogous to goodness rather than permission or requirement. It is attractive to say that certain paradoxes are analogous to moral dilemmas: the least bad response to them is still metaphysically very bad, indeed unacceptable. I think this is a plausible characterization of the situation of someone who realizes the generality of the lottery and preface puzzles, and is presented with SIC, MIC, and EPC. Maybe the best thing for them to judge is that people know very little, yet that forced skeptical judgement is metaphysically worse than our normal generous attributions of knowledge. Considering the lottery puzzle makes us view our epistemic situation through a distorting lens. (The details become important in giving such a response to the sorites paradox, so I leave further discussion of these matters to my 2009. Note that truth-relativism can't capture these ideas.)

I have introduced and motivated the ideology of metaphysically acceptable judgements. Section five gives a broader view of the metaphysics and philosophy of language I propose. Section six details how it resolves the lottery and preface puzzles, including how it avoids my objection to truth-relativism.

5. Metaphysics and Meaning.

This section places the metaphysical acceptability of judgements at the heart of metaphysics. This framework implies no revisionary theory of meaning.

Nearly all judgements have a level of metaphysical acceptability.[17] Saying what makes judgements that p acceptable replaces saying what grounds the truth of p (e.g. Fine 2001). This allows us to understand subjectivity and objectivity as follows.

[17] Maybe we can characterize the 'open future' as the view that predictions don't have a level of metaphysical acceptability until the time the concern has passed. They haven't 'answered to reality' till then. They do already have a level of epistemic acceptability.

P is *completely subjective* iff only facts about the judge determine whether it is metaphysically acceptable for them to judge that p.

P is *partially subjective* iff facts about the judge (qua judge) play a partial role in determining whether it is metaphysically acceptable for them to judge that p.

P is *objective* iff one occurrent attitude is the only metaphysically acceptable one for anyone to take towards p.

Deliciousness is completely subjective. (An analogous moral anti-realism holds that there are moral beliefs and judgements. For them to be metaphysically acceptable is for them to reflect the judge's deepest commitments of a certain kind. Of course, that's not what it is for them to be morally acceptable.) I think that the lottery and preface puzzles show that knowledge is partially subjective. Whether a knowledge-attribution is metaphysically acceptable depends both on what's psychologically salient to the judge, and on the situation of the subject of the ascription. Consider a borderline case where the underlying precise facts are objective, such as whether winning $1,000,000 makes you rich for a New Yorker. On the definitions I have suggested, such a proposition is neither subjective nor objective. I don't mind if the reader prefers slightly different definitions. The ideology of metaphysical acceptability usefully distinguishes many cases, and it doesn't matter exactly how the rough categorizations of common-sense map on to them.[18]

Metaphysical acceptability helps us lock on to the idea of reality-as-it-is-in-itself. I accept Kit Fine's unanalyzed concept of a 'fact in metaphysical reality', abbreviated R(p) (Fine 2001, 2008). A good approximation of R(p) is: it is a fundamental fact that p. Some facts are metaphysically privileged. The others are still facts. Suppose that p. Realism about p holds that R(p). Anti-realism about p holds that not-R(p). It is not even a fact that there are no chairs, let alone a fundamental fact. Anti-realism about chairs is not the trivially false view that R(there are no chairs). It is the view that there are no R-facts about chairs. There are chairs, but they are not elements of metaphysical reality, i.e. not 'metaphysically real' (Fine 2008). Analogously, there is such a property as being delicious, and such a relation as knowing, but they are not 'metaphysically real'.

[18] I give a more detailed taxonomy in my 2009.

Realism is linked to metaphysical acceptability by the following thesis. (Here I differ from Fine.)

> It is a fact-in-reality that p iff what makes it metaphysically acceptable for everyone to judge that p is: that p.

The metaphysics of p consists in saying what makes judging that p metaphysically acceptable or not. Usually, the explanation will not appeal to whether p. For example, our accounts of deliciousness, vague properties, and knowledge did not appeal to facts about their instantiation. (The point of the accounts is that they are silent on the instantiation of the property.) This won't always be the case. If it is a fundamental fact that p, the explanation of what makes it 'right' or metaphysically acceptable to judge that p will be: that p. Such an explanation shows that we can only understand the area of discourse by adding such facts to our metaphysical picture of reality. So saying what makes it metaphysically acceptable to judge that p tells us whether p belongs to reality-as-it-is-in-itself.

Facts about knowledge aren't fundamental. Everyone is anti-realist about knowledge in that sense. My distinctive claim is that knowledge is partially subjective: whether it is metaphysically acceptable to attribute knowledge to Sally differs between judges.

The anti-realist need not explain the metaphysical acceptability of a judgement in terms of R-facts. The anti-realist needs to understand the acceptability of judgements that p in more fundamental terms, not in absolutely fundamental terms. For example, it is metaphysically acceptable for me to judge that the cake is delicious, because it is disposed to taste good to me. Judging something to be red is metaphysically acceptable because of the precise shade the thing is. The explanans are not facts-in-reality. The anti-realist explains the practice of judging whether p, and one does not understand a way of life in terms of fundamental physics. One should also be anti-realist about how metaphysically acceptable a given judgements is, for that won't be a fundamental fact. (There are different possible views about what makes judgements of metaphysical acceptability metaphysically acceptable.)

This metaphysics does not conflict with standard theories of the meaning of the word "know". My theory concerns when a judgement is metaphysically acceptable.

There's nothing linguistic about that. I have not given an (assertibility-conditions) theory of meaning. On the other hand, a compositional meaning theory should be about purely linguistic matters. Its job is to explain why "le gateau est delicieux" means in French that the cake is delicious, in terms of the way the sentence is made up of those words. It is not the job of linguistics to tell you whether the cake is delicious, or how to judge on that issue. On the truth-conditional approach, we want the theory to predict that "le gateau est delicieux" is true iff the cake is delicious. Again, that's independent of whether the cake is delicious, or how to judge on that issue. Anti-realism about such entities says that the cake exists, and that there is a property of deliciousness, though they are not 'metaphysically real' because they do not figure in any of the R-facts. So we can assign the cake as the meaning of "le gateau", and deliciousness as the meaning of "delicieux". Metaphysical antirealism is compatible with referential semantics.

On my view, our compositional semantics is neutral on metaphysical issues. Let's review two theories that build the metaphysics into the semantics. They suffer from essentially the same defect: they make it harder or impossible to explain compositionality. Michael Dummett (1993) suggests that the meaning of a sentence consists in its assertibility-conditions. He has an epistemic conception of assertibility in mind, but let's consider the corresponding view for metaphysical acceptability. Compositionality says that the meaning of a sentence is determined by the meanings of its parts. So if meanings are assertibility conditions, then the assertibility conditions of "p" and "q" must determine the assertibility conditions of "not-p", "p or q", and "p and q". But they don't.[19] For example, "Quentin's car is red" is assertible, and "Quentin's car is orange" is assertible, but "Quentin's car is red and Quentin's car is orange" is not assertible. I say that there is a top-down rule that red and orange are to be treated as incompatible. The assertibility conditions of the whole is not determined by the assertiblity conditions of the parts. But that's incompatible with meanings being assertibility-conditions (which would therefore have to compose). This kind of argument

[19] Dummett 1993 p. 473 admits that, after thirty years of trying, he can't see how to do negation.

against epistemically constrained conceptions of meaning is championed by Jerry Fodor and Ernest Lepore (Fodor 1998, chapters 4 and 5; Fodor and Lepore 2002).

Matthew Chrisman (2007) argues that the lottery puzzle motivates expressivism about knowledge attributions. Expressivism says that superficially similar sentences have very different semantics. The meaning of a descriptive sentence is the belief it is conventionally used to express. The meaning of a non-descriptive sentence is the non-cognitive attitude it is conventionally used to express. Expressivism about knowledge-attributions says that "Bob knows I've taken the spare key" is a non-descriptive sentence. Compositionality then requires that we explain, purely in terms of the mental states expressed by the disjuncts, the mental state expressed by: "Either Bob hasn't looked in the kitchen drawer, or he knows I've taken the spare key." Contemplating this challenge should give one a sinking feeling. I can't think what kind of mental state would be expressed, let alone how to compute it from the meanings of the disjuncts and of 'or'. Here's how Mark Schroeder puts the complaint against expressivism about moral language.

> The essence of the Frege-Geach Problem is that moral and descriptive terms play exactly the same kind of semantic role in every kind of complex linguistic construction in natural languages. Since non-cognitivist views consist centrally in the idea that moral terms like 'wrong' have a different kind of meaning from ordinary descriptive terms like 'green', that makes noncognitivism look like a very unpromising hypothesis about natural language semantics. (Schroeder 2008 p. 717)

For detailed argument that the challenge can't be met, see Schroeder 2009.[20]

Prima facie, neither expressivism nor Dummett's theory are sensible attempts to explain compositionality. By contrast, truth-conditional or propositional approaches do bear fruit. It seems to me a methodological mistake to hold the attribution of content in psychology and linguistics hostage to metaphysical concerns. Linguistics seeks to explain how linguistic communication works. That's a matter of speakers and hearers knowing

[20] This problem does not arise for truth-relativism, including Field 2009. But I argue in section 3c.iii and in my 2009 that relative truth won't do the desired metaphysical work.

the linguistic conventions of word meaning and sentence construction, and understanding each others' psychology. Compositional semantics is concerned with the former, or more precisely, what the conventions are. (How knowledge of the conventions is implemented in the brain is another matter.) For example, one convention might be that "knows" refers to knowledge. One cannot go to the linguistics department and tell them that they are pursuing a defective research programme because moral properties are metaphysically 'queer' and so we have to be expressivists. Metaphysicians don't get to say what counts as good cognitive science. Similarly, I think it is a methodological mistake to argue from the metaphysical thought that there is no 'objective fact' about whether curry is delicious, to the linguistic conclusion that we should be contextualists about "delicious". It is also a mistake to argue from invariantism about "delicious" to the metaphysical conclusion that there must be an 'objective fact' about what's delicious. My theory gets these things right by making the truth of p irrelevant to the metaphysics of p.

My metaphysics does show we shouldn't read too much into certain models of meaning we use to explain compositionality. A standard implementation of truth-conditional semantics assigns objects and functions as the 'semantic values' of expressions. Assertoric utterances get a semantic value (Kaplan's 'content') that is a function from possible worlds to {True, False}. That function models the truth-conditions of the utterance. An utterance of "the cake is delicious" maps the actual world to one of {True, False}. So there is no candidate semantic value it is metaphysically acceptable for all English speakers to conform to. This does not affect the explanatory power of positing semantic values as part of the scientific project of explaining compositionality. That task is to explain the meaning of complex sentences, *given* the meaning of the atomic expressions. In the current terms, we explain how *whichever* semantic value is assigned to "delicious" affects the semantic value of uttered sentences containing that word. This project does not assume that there is exactly one function that it is metaphysically acceptable for all English speakers to conform to. Nor does it require figuring out the semantic value of the atoms; we only need to know their type. (That's fortunate. We would get into a pickle trying to settle on the semantic value of "knows". Any candidate satisfying Multiple Item Closure would be either completely skeptical or encode objectionable dogmatism.)

Modeling meanings as semantic values is a good way to explain how compositionality works. Specific explanatory purposes warrant certain idealizations. We can ignore vagueness and the relativity of deliciousness and knowledge when explaining compositionality. When we are interested in the topics of this paper, we should not take seriously a model of meaning that abstracts away from those very phenomena.

6. Salience and Knowledge Attributions.

Facts about what is psychologically salient to the judge affect whether it is metaphysically acceptable for them to make a knowledge-attribution. Our initial judgement that Sally knows her car is round the corner is metaphysically acceptable. Once the question of car-theft is psychologically salient to the judge, the only acceptable judgement is that Sally doesn't know. The metaphysical theory is silent on whether Sally knows – that's irrelevant to understanding the nature of knowledge.

This theory gets right what contextualism gets wrong. You end up judging that Sally doesn't know her car is parked round the corner, and you remember you just judged that she does know that. Thus you judge that your initial verdict was false, and in that sense wrong. You changed your mind. (This is compatible with recognizing that your previous judgement was metaphysically acceptable.)

I understand the salience of a proposition psychologically. Either the judge takes q into consideration in judging whether S knows that p; or the judge ignores q in so doing. Judgement is occurrent belief; taking into consideration is also an occurrent mental phenomenon. It is initially metaphysically acceptable to judge that Sally knows that p partly because the judge ignores whether Sally knows that her car hasn't been stolen. This is so even if someone judges purely on the basis of testimony that Sally knows her car is round the corner. If the possibility of car-theft is salient to the judge, and Sally's situation is as stipulated, then it is metaphysically unacceptable for them to judge that she knows where her car is.

A knowledge-attribution is part of a representation of the subject's epistemic position. It is metaphysically unacceptable to represent Sally as in a position to know that her car hasn't been stolen and taken elsewhere. When that issue is taken into consideration as a way for Sally's belief to be false, judging that Sally knows her car is

round the corner represents her as able to know that it hasn't been stolen. When the issue is not salient, judging that Sally knows her car is round the corner does not represent that she is able to know it hasn't has been stolen. That's why taking car-theft into consideration stops it being metaphysically acceptable for you to judge that Sally knows her car is round the corner.

This explains why affirmation is the only metaphysically acceptable attitude to Single Item Closure. To judge that S knows that p, taking into consideration that q is a way for p to be false, is to thereby represent that S is in a position to know not-q. Not to judge that this is so would be incoherent, a metaphysically unacceptable combination of mental states. When one is not taking q into consideration, there may be nothing wrong with not judging that S knows that not-q.

The judge's taking into account a matter they had been ignoring is also the key to the other phenomena described in section two. One can judge metaphysically acceptably that Sally knows her car hasn't been stolen, while ignoring the issue of assigning it a likelihood given Sally's evidence. Once one considers it, the judge must deny that the epistemic probability for Sally that her car has been stolen is zero. That represents car-theft as epistemically possible for Sally; i.e. it represents that Sally doesn't know it hasn't been stolen. (That's why only affirming the Epistemic Possibility Constraint is metaphysically acceptable.) Ceasing to ignore the question of the epistemic probability changes the acceptability of judging that Sally knows.

The metaphysical acceptability of judgements about epistemic probability will vary between judges, as it does for judgements about knowledge. For S is in a position to know that p iff there is no chance given S's evidence that not-p. A judge-invariant metaphysically acceptable answer as to whether [Sally's evidence leaves no chance that her car is not round the corner] would determine a judge-invariant metaphysically acceptable answer about whether she is in a position to know it is still there.

Similarly, there can't be a judge-invariant metaphysically acceptable answer to whether Sally is epistemically entitled to assert that her car is parked round the corner. Being known is the necessary and sufficient epistemic condition for being asserted, and for being taken as a premise. So whether it is metaphysically acceptable for someone to judge that [Sally is epistemically entitled to assert her car is round the corner] depends on

whether car-theft is salient to them. Insofar as we intuitively affirm KNA, AKP etc in the cases at hand, this is the right result.

In the preface puzzle, the judge ignores whether John knows the conjunction, when acceptably judging of each person on the list that John knows that they will be at the conference. If they consider whether John knows that all 100 people will be there, they must judge that he can't – given John's evidence, the probability that someone will get sick is quite high. Now suppose the judge addresses whether John knows Ally will be at the conference, taking into consideration that it is likely given John's evidence that someone will get sick, and that John is in the same epistemic situation with respect to the attendance of each person on the list. The only acceptable way to do so is to affirm that there is a chance given John's evidence that Ally will get sick, and hence that John does not know she will be at the conference. Judging that John knows that Ally will be at the conference ceases to be metaphysically acceptable, as it would then represent John as in different epistemic positions with regards to Ally's attendance and that of some of the other people on the list.

That's how the 'epistemic probability' version of the preface puzzle works. The 'closure' version is very similar. Once we take into consideration that John can't know that all 100 people on the list will be at the conference, we can't judge that John knows Ally will be there, on pain of representing his epistemic situation with regards to Ally's attendance as special.

Salience also explains our different reactions to winning a lottery versus Jonathan Vogel's 'heartbreaker' case. Prima facie, you can't know you won't win the lottery. But (prima facie) you can know that not all sixty golfers in the competition will get a hole-in-one on the fiendish 'heartbreaker' hole. (Vogel 1999 p. 165.) A matchbox might contain a few duds, but (prima facie) you can know that not every match in the box is a dud. (Hawthorne 2004 pp. 12-20.) What explains the difference in our reaction to a lottery versus the heartbreaker and the matchbox? I think it is the psychological proximity of preface-style considerations. It is usually salient when considering a lottery that every ticket has a chance of winning. (When it is not salient, a judge can acceptably insist that they know they won't win. For example, a judge can partition outcomes only into their winning versus their losing. Hawthorne 2004 pp. 17-20 gets the psychology exactly

right.) When the matter is salient, you must respect that you are in similar epistemic positions with regards to each ticket's winning. As in the preface, that forces the judgement that you can't know of any particular ticket that it will lose. By contrast, we usually consider the heartbreaker and matchbox cases without there being such a set of propositions the subject is similarly positioned towards, that being salient, and for which it is salient that the subject does not know that they are all true. When such a conjunction becomes salient, we must change our judgement about the case. For example, if we think of the matchbox as one amongst thousands made by a company, we must judge that we can't know that no box contains only dud matches. The decent probability that there is a box that's all duds is the result of the small but non-zero probabilities that any particular box is all duds. So we can't know that this box is not all duds. Similarly, we can imagine the golf-competition being played every day for a million years. There's a decent chance that on one such day, every golfer will get a hole-in-one on the heartbreaker. The actual competition isn't special, so there's a small chance that every golfer will achieve the feat today.[21]

If this is right, then the preface puzzle is responsible for the lottery puzzle. The reason we usually judge that you can't know your ticket won't win the lottery, and that Sally can't know her car hasn't been stolen, is that we naturally consider a set of propositions that mandates that verdict on preface-style grounds. For example, it is natural to think that Sally knows that cars are sometimes stolen in situations epistemically just like this one. No such proposition is automatically suggested by considering whether Sally knows her car is parked round the corner, or whether you know you won't be able to afford a vacation in Barbados next year. We don't automatically think of those cases as one among many epistemically similar ones. We standardly evaluate those knowledge-attributions by considering the normal way for things to be.

[21] This framework solves the problem with objective chance (Hawthorne 2004 pp. 4-5, 93-4). A judge can ignore that according to quantum theory, there is a non-zero objective chance that the cricket ball will pass through the window without breaking it. It is then metaphysically acceptable for them to judge they know the ball will break the window.

Let's see how this theory avoids the objection to truth-relativism. Initially it is acceptable for Polly to judge that John knows Ally will be at the conference. It is acceptable for Polly to judge that John is in the same epistemic position with respect to everyone on his list. Indeed it is acceptable for Polly to: judge that John knows Bert will be there, judge that Carol will be there, …, and judge that Zena will be there. It is acceptable for Polly to affirm MIC. But it does not follow that it is acceptable for Polly to judge that John knows no-one will get sick and miss the conference. We can't derive that conclusion even if we interpret metaphysical acceptability as a kind of requirement, and apply standard deontic logic. For there to be a problem, the following principle would have to be true (aping the closure of relative truth).

> If {P1,…,Pn} entails Q, it is metaphysically forbidden to: judge that P1, …, judge that Pn, and not judge that Q.

But that's not right. When one ignores Q, one does not have a bad occurrent attitude towards Q; one has no occurrent attitude towards Q. In the case at hand, Polly is ignoring whether John knows no-one will get sick and miss the conference. It is exactly the kind of case we can't apply the above principle.

Here's another way to put the point. Polly is required to treat John as in the same epistemic position with respect to each person on his list, and required to treat knowledge as closed. But Polly is doing nothing wrong when she judges that John knows Ally will be at the conference. She is not violating either requirement. For they forbid judging that p while denying or suspending judgement on q. Polly judges that p while ignoring whether q, i.e. having no occurrent attitude towards it. So Polly is not required to judge that John knows that no-one on his list will get sick, even if the theory is framed in terms of what she is required to judge, and that notion satisfies standard deontic logic.

If fact, I suggested at the end of section 4 that it is best to make the basic term of our theory a normative notion analogous to the sense in which all options in a moral dilemma are bad. There is no prospect of anything like standard deontic logic for such a notion. So there is no prospect of the objection to truth-relativism also applying to my theory.

If you ask me, I will answer that SIC, MIC and EPC are true, and knowledge is factive.[22] In theorizing about the lottery puzzle, we bring to salience whether Sally's car has been stolen. We theorists end up judging that Sally does not know that her car is parked round the corner. But the theory itself is silent on all those questions. It is silent on whether skepticism is true, and silent on whether dogmatism is true. Analogously, our theory is silent on whether curry is delicious. Those questions are not metaphysically relevant. All that matters is that the judgement that Sally does know, made before car-theft has become salient, is just as metaphysically acceptable as the subsequent judgement that Sally doesn't know.

Further, I suggested that metaphysical acceptability is a matter of a degree of goodness. The initial judgement that Sally does know is metaphysically more acceptable than the subsequent denial of that proposition. That's why we often feel annoyed when someone uses lottery reasoning to make us retract an attribution of knowledge. That denying knowledge is now metaphysically best does not mean that it is highly acceptable. (Similarly, the best action in a moral dilemma will still be morally quite bad.) Though they are overridden, the original considerations still weigh in favour of saying Sally does know, making the retraction required but not happy. Paradigm-case considerations might play into our judgements of knowledge, and be overridden by preface-type reasoning. More interestingly, the latter can distort our understanding of the subject's normative position: what they are in an epistemic position to assert, or take as a premise in practical

[22] We can't capture the factivity of knowledge by saying that if p is false, then any judgment that someone knows that p is metaphysically unacceptable. Suppose I think Jennifer falsely but acceptably judges that Sally knows that her car is parked round the corner. Then I should say the same thing about Jennifer's judgement that she knows that Sally knows her car is parked round the corner. Again, it is metaphysically acceptable for my Dad to judge he knows that curry is not delicious, even though it is. So let's replace the bad first-pass at capturing factivity with the following. It is metaphysically unacceptable to: judge that S knows that p and not judge that p. It is metaphysically unacceptable to: judge that not-p and not judge that S does not know that p.

reasoning. The normative role of knowledge means we understand the subject's position better by attributing it, but that's trumped by the preface reasoning.

Why is it a good idea to employ a concept that answers to reality in this way? Here's a Just So Story. We are believing creatures. We cannot escape the normative question, What should I believe? We aim for our beliefs to reach a certain epistemic standard. Label that status 'knowledge'. If p is worthy of belief, then it is worthy of being taken as a premise. We reason from multiple premises, and as beliefs the results should also be worthy of being acted on and reasoned from. So knowledge must obey Multiple Item Closure. On the other hand, it's not a good general answer to the normative question that there's nothing we can take as a premise. As the preface puzzle shows, those demands on the concept of knowledge are in tension. We accommodate them both by ignoring preface-style considerations most of the time. That knowledge is the epistemic standard for our premises explains why it is subject to the preface puzzle, and so why it has the metaphysics I have argue for.[23]

7. Conclusion.

I judge that curry is delicious. But the deep philosophical question is not whether curry is delicious, but rather which judgements about that matter are metaphysically acceptable. The answer to the deep question is that it is metaphysically acceptable for J to judge that curry is delicious iff curry tastes good to J. When we consider our lottery puzzle, we judge that Sally doesn't know that her car is parked round the corner. But again, the deep question is not whether Sally knows, but rather which judgements on that issue are metaphysically acceptable. The answer to the deep question is that people

[23] Philosophers defending a 'have your cake and eat it' response to the lottery usually extend it to deal with Brain-In-a-Vat (BIV) skepticism. But that's not obvious. BIV skepticism turns on issues of epistemic priority, i.e. whether the Moorean solution is correct; epistemic priority is not at issue in the lottery puzzle. The lottery puzzle concerns whether the subject knows that p, versus p being merely very likely given their evidence; that's not the contrast at issue in BIV skepticism. Also, I suggested that the lottery puzzle derives from the preface puzzle; BIV skepticism doesn't.

usually ignore lottery- and preface-style considerations, and this helps make their generous attributions of knowledge metaphysically acceptable.

Knowledge works this way because it is the epistemic standard we aim for our premises and assertions to meet. That standard must be attainable, but also guarantee that we can reason from multiple premises to a conclusion that meets the standard too. Those competing demands on knowledge can only be accommodated by loosening its ties to Reality.

Appendix: Pragmatic Encroachment.

Our assessment of the practical cost of a subject's having a false belief that p affects our judgement as to whether they know that p. Call this 'pragmatic encroachment' into out knowledge attributions. This appendix extends my theory to deal with puzzles about this. Here's our working example (following Stanley 2005 pp. 3-5).

> It's Friday afternoon. There's a long line at the bank, so Bob decides to leave depositing his cheque till Saturday. It's not important that he deposit it immediately. Bob is stopped in the street by Hannah, who asks him whether the bank will be open on Saturday. He answers, "The bank will be open on Saturday. I know because I was there on a Saturday only two weeks ago." It is really important to Hannah and Sarah that they deposit a cheque by Saturday - they have a big direct debit due. They also remember that the bank was open on Saturday two weeks ago, but want to be sure about this Saturday. Hannah reports back to Sarah, "No, that guy doesn't know either that the bank will be open on Saturday."

Intuitively, Bob spoke truly, Hannah spoke truly, but they are disagreeing about whether Bob knows that the bank will be open on Saturday. Those three things can't all be true. As with the lottery paradox, there are three broad responses to this puzzle. The first is that Bob is right and Hannah is wrong – the end. (This is what subject-sensitive invariantism says, e.g. Hawthorne 2004 chapter 4 and Stanley 2005. For it holds that what matters is what's at stake for the subject of the knowledge ascription, and Bob does not have much riding on the bank's being open.) The second is that Bob is wrong and Hannah is right – the end. The third is that they are both right in the important sense. For example, contextualism says that Bob and Hannah both spoke truly, but did not actually contradict

one another. I will extend the framework I proposed to deal with the lottery and preface puzzles to deal with this problem.[24]

We can have different purposes in evaluating whether Bob knows. We can be interested in whether Bob should leave depositing his cheque till Saturday. We can be interested in whether Bob should plan a bank robbery for Saturday. We can be interested in whether Hannah and Sarah should leave depositing their cheques till Saturday. We can be focused on different practical questions for different agents, who need not be the subject of the knowledge ascription. That focus affects whether it is metaphysically acceptable for us to attribute knowledge to the subject. In the example, we are interested first with Bob and then with Hannah's normative situation.

It is worse for Hannah to falsely believe that the bank will be open on Saturday than it is for Bob. So when we are focused on Hannah's normative situation in deciding when to go to the bank, we acceptably judge the attribution of knowledge to any subject more skeptically. Hannah can't leave the trip to the bank till Saturday, because she does not know that the bank will be open then. The Action-Knowledge Principle captures this intuitively compelling explanation. We are focused on Hannah's situation, so whether other subjects know is only relevant insofar as it affects whether she knows. Other subject's knowledge is relevant because they can tell her things. We judge that Bob doesn't know that the bank will be open because we judge that Hannah can't come to know that on the basis of his testimony.[25]

Suppose we focus on Hannah's normative situation, and so judge that Bob does not know that the bank will be open on Saturday. An obvious way to shift the normative focus from Hannah to Bob is to consider whether he should act on the premise that the bank will be open on Saturday. Considering how the AKP applies to Bob shifts the focus to his normative situation. My theory correctly predicts that when we consider how the

[24] See Hawthorne 2004 pp. 179 and Stanley 2005 pp. 114-9 on the prospect that contextualism will deal adequately with cases of pragmatic encroachment.

[25] We can preserve the principle that: if A and B both believe that p on the same evidence, then A knows that p iff B knows that p. Whose practical interests are salient to us affects whether we acceptably judge that they both know that p, or that neither do.

AKP applies to Bob, we revise the verdict reached when focused on Hannah's situation. Now we judge that Bob can leave his trip to the bank till Saturday, because he knows it will be open. That judgement is metaphysically acceptable because the judge focuses on Bob, just as the contrary verdict was metaphysically acceptable when the judge focused on Hannah. Thus my theory endorses these intuitions and the AKP.[26]

In the lottery and preface puzzles, the acceptability of a knowledge-ascription depends partly on which matters are psychologically salient to the judge. In the present case, what makes a difference to the metaphysical acceptability of a knowledge-attribution is: which practical question for which agent the judge is psychologically focused on. This addition to my treatment of the lottery and the preface puzzles is very modest, and of a kind already accepted. The theory says that Bob and Hannah's judgements are both metaphysically acceptable, though they disagree. There will also be corresponding variation in the metaphysical acceptability of judging that Bob is epistemically entitled to assert that the bank will be open.

In our initial example we shift which agent we are focused on in assessing the knowledge-attribution. Here's a case where we hold the agent fixed, but shift the practical question we are interested in. Consider whether Agnes knows that her car won't get stolen in Greenpoint. She currently lives in Greenpoint, but is struggling financially, and rent in Bushwick is substantially lower. If we are interested in whether Agnes should move from Greenpoint to Bushwick, then we might judge that Agnes knows that her car won't get stolen in Greenpoint, she doesn't know it won't get stolen in Bushwick, and that this is an argument she can give for staying in Greenpoint. However, if we are interested in whether Agnes should renew her insurance against car-theft, the best judgement is that she doesn't know that her car won't get stolen in Greenpoint. Intuitively, whether she should take it as a premise shifts, along with whether she knows it. This intuition is in perfect harmony with the Action-Knowledge Principle. Subject-Sensitive Invariantism is not well placed to accommodate this kind of example, because the subject's situation does not change just because a judge starts attending to a different

[26] Stanley thinks that's impossible. He claims we should accept the AKP and thus reject our intuition that Hannah spoke truly (2005 pp. 114-5).

practical question. Given AKP, cases like Agnes' immediately motivates my view about how to use the concept of knowledge.

References

Author. 2009. "The Ideology For Understanding Vagueness."

Author. Forthcoming. "Appearances, Rationality, and Justified Belief."

Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm of practical Reasoning," *Nous* 42:2, pp. 167-189.

Cappelen, Herman, and Hawthorne, John. 2009. *Relativism and Monadic Truth*, OUP.

Chrisman, Matthew. 2007. "From Epistemic Contextualism to Epistemic Expressivism," *Philosophical Studies* 135(2), pp. 225-254.

Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*, OUP.

Cohen, Stewart. 1988. "How to be a Fallibilist," *Philosophical Perspectives* 2, pp. 91-123.

DeRose, Keith. 2006. "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments Against Contextualism", *Philosophy and Phenomenological Research* 73:2, pp. 316-338.

---- 1996. "Knowledge, Assertion, and Lotteries," *Australasian Journal of Philosophy* 74, pp. 568-580.

---- 1995. "Solving the Skeptical Problem," *Philosophical Review* 104, pp. 1-52.

---- 1991. "Epistemic Possibilities," *Philosophical Review* 100, pp. 581-605.

Dummett, Michael. 1993. *The Seas of Language*, OUP.

Egan, Andy, Hawthorne, John, and Weatherson, Brian. 2005. "Epistemic Modals in Context," in eds. G. Preyer and G. Peter, *Contextualism in Philosophy*, OUP.

Fantl, Jeremy, and McGrath, Matthew. 2007. "On Pragmatic Encroachment in Epistemology," *Philosophy and Phenomenological Research* 75:3, pp. 558-589.

Field, Hartry. 2009. "Epistemology Without Metaphysics," *Philosophical Studies* 143, pp. 249-290.

Fine, Kit. 2008. "The Question of Ontology," in eds. Chalmers, Manley, and Wasserman, *Metametaphysics*, OUP.

---- 2001. "The Question of Realism," *Philosophers' Imprint*, vol. 1:1.

Fodor, Jerry. 1998. *In Critical Condition*, MIT Press.

Fodor, Jerry, and Lepore, Ernest. 2002. *The Compositionality Papers*, OUP.

Hawthorne, John. 2004. *Knowledge and Lotteries*, OUP.

Hawthorne, John, and Stanley, Jason. 2008. "Knowledge and Action," *Journal of Philosophy* CV.10, pp. 571-590.

Kolbel, Max. 2003. "Faultless Disagreement," *Proceedings of the Aristotelian Society* 104, pp. 53-73.

Lewis, David. 1996. "Elusive Knowledge," *Australasian Journal of Philosophy* 74, pp. 549-567.

MacFarlane, John. 2005. "The Assessment-Sensitivity of Knowledge Attributions," in eds. T. Gendler-Szabo and J. Hawthorne, *Oxford Studies in Epistemology Volume One*, OUP.

Pryor, James. 2000. "The Skeptic and the Dogmatist," *Nous* 34, pp. 517-549.

Richard, Mark. 2004. "Contextualism and Relativism," *Philosophical Studies* 119.1-2, pp. 215-242.

Schroeder, Mark. 2009. *Being For: Evaluating the Semantic Program of Expressivism,* OUP.

---- 2008. "What is the Frege-Geach Problem?" *Philosophy Compass* 3/4, pp. 701-20.

Stanley, Jason. 2005. *Knowledge and Practical Interests*, OUP.

Unger, Peter. 1975. *Ignorance: The Case for Skepticism*, OUP.

Vogel, Jonathan. 1990. "Are There Counterexamples to the Closure Principle?" in eds. Ross and Roth, *Doubting*, Reidel.

Williamson, Timothy. 2005. "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge," *Philosophical Quarterly*, 55:219, pp. 213-235.

---- 2000. *Knowledge and its Limits*, OUP.

Wright, Crispin. Forthcoming. "On the Characterisation of Borderline cases," in ed. Gary Ostertag, *Meanings and Other Things: Essays on Stephen Schiffer*, MIT Press.

---- 2008. "Relativism about Truth Itself: Haphazard Thoughts about the Very Idea," in eds. Manuel Garcia-Carpintero and Max Kolbel, *Relativising Truth*, OUP.

Yalcin, Seth. 2007. "Epistemic Possibility," *Mind* 116, pp. 983-1026.