

## The New and Old Ignorance Puzzles: How Badly Do We Need Closure?

Brent G. Kyle

Skeptical puzzles and arguments often employ knowledge-closure principles. Consider

The Old Ignorance Puzzle:

[Know] Moore knows that he has hands.

[Old-Not-Know] Moore does not know that he's not a handless-brain-in-a-vat (BIV).

[Closure] If S knows that P, and knows that P entails Q, then S knows that Q.

The Old Ignorance Puzzle is genuinely puzzling, because its three claims are individually plausible yet cannot all be true (assuming Moore knows having hands entails not being a BIV). The puzzle turns on its third claim, a version of knowledge-closure. In general, knowledge-closure principles state that if a subject knows a proposition then she also knows, or is in a position to know, other propositions suitably connected to the first.<sup>1</sup> Must we accept a knowledge-closure principle?

Many epistemologists think so (Bonjour 1987, 311; Feldman 1999, 95). After all, these principles are highly intuitive, and are useful for explaining why abominable conjunctions (e.g. 'Moore knows he has hands but doesn't know he's not a BIV') are abominable. Moreover, epistemologists widely believe that an adequate reply to the skeptic should explain why his reasoning is appealing albeit misleading (Sosa 2000, 15; Cohen 1999, 63; DeRose 1995, 3); but it's unclear what would explain the appeal of the skeptic's closure principle, if not for its truth. And lastly, closure principles are useful for capturing an intuitive idea—that we can always gain knowledge by inferring from what we already know (Kvanvig 2006, 257).

I aim to challenge this widespread commitment to knowledge-closure. But I shall proceed by first examining a new puzzle about failing to know—what I call the New Ignorance Puzzle

---

<sup>1</sup> Unless otherwise specified, I use the phrase 'knowledge-closure' to refer to general closure principles, like [Closure], as well as particular instances of such principles (e.g. If Moore knows he has hands then he knows he's not a BIV). My claims in this paper apply to any version of knowledge-closure no matter how it's formulated, but I shall focus on simpler versions, like [Closure], for brevity's sake. Closure principles whose consequents do not involve knowledge attributions (e.g. ones stating that justified belief is closed) are left unaffected by my arguments.

(sections I-III). This puzzle resembles the Old Ignorance Puzzle, although it does not involve a closure principle. It instead centers on the standard view of ignorance, a highly intuitive principle stating that ignorance is merely a failure to know. In sections II and III, I argue that the best way to solve the New Ignorance Puzzle is to reject the standard view of ignorance and to explain away its appeal via conversational implicature. I then use this solution to the New Ignorance Puzzle as a way of explaining why knowledge-closure principles would seem true, and why abominable conjunctions would seem abominable, even if such principles were false (section IV). The upshot is a new way of explaining why the skeptic's reasoning is appealing albeit misleading.

What is the connection between knowledge-closure and the standard view of ignorance? Focus temporarily on the idea that we can always *gain knowledge* by inferring from what we already know. If ignorance is more than just a failure to know, contrary to the standard view, then it's possible to fail to know something without being ignorant of it. This possibility allows closure-deniers to accept a principle that is weaker than, but similar to, knowledge-closure—roughly, that we can always *lose ignorance* by inferring from what we know. As I show in section IV, this weaker principle helps closure-deniers resist two major reasons for accepting closure.

### **I. The New Ignorance Puzzle.**

The New Ignorance Puzzle is similar to the Old Ignorance Puzzle. But instead of employing a closure principle, it centers on the standard view of ignorance, according to which ignorance is merely a failure to know. Alvin Goldman and Erik Olsson explain the theory like this:

Knowledge seems to contrast with ignorance. Not only do knowledge and ignorance contrast with one another but they seem to exhaust the alternatives, at least for a specified person and fact. Given a true proposition *p*, Diane either knows *p* or is ignorant of it (2009, 19-20).

On the standard view, knowledge and ignorance are mutually exclusive and jointly exhaustive; they are contraries and contradictories. Goldman and Olsson explicitly assume this view rather than argue for it. But they are not alone. Others, like Michael Zimmerman (1988, 75), Peter Unger (1975, 93), and Stephen Levinson (2000, 208), also explicitly assume the standard view without defending it. The theory can be stated as necessary and sufficient conditions for ignorance:

Standard View:

For any truth P, S is ignorant of the fact that P if and only if S does not know that P.<sup>2</sup>

The above theorists assume the standard view explicitly. But it is more common to find people assuming it implicitly. Without stating the standard view, many philosophers make implicit transitions from ‘not know’ to ‘ignorant’ as if these expressions are synonymous. For example, Norman Gulley (1968, 63) and Gregory Vlastos (1985, 22) each describe Socrates’ disavowal of knowledge as an “avowal” or “confession” of ignorance. We also find Bernard Williams discussing whether one can know that one does not know that P; and the name he gives to this higher-order knowledge is “knowledge of ignorance” (1995/2006, 174). And lastly, Timothy Williamson holds that vague sentences are “either unknowably true or unknowably false.” Either way, says Williamson, “vagueness is a kind of ignorance” (1997, 265). These philosophers all make implicit moves from ‘not know’ to ‘ignorant’, but it’s hard to see how these inferences would be justified without the sufficiency condition of the standard view.<sup>3</sup>

---

<sup>2</sup> Unlike some theorists (e.g. Goldman and Olsson 2009), I find it more natural to speak of ignorance as a relation to facts rather than to propositions or states of affairs. But this makes no substantive difference in what follows. When I make claims like ‘S is ignorant of *the fact that snow is white*’ the reader is welcome to rephrase these claims as either ‘S is ignorant *that snow is white*’ or ‘S is ignorant of *snow’s being white*’.

<sup>3</sup> To be fair, Williamson could instead be relying on this principle—if S is *not in a position to know* that P, then S is ignorant of the fact that P. Even so, the issues raised in this paper arise similarly for this weaker principle. For other philosophers who implicitly assume the standard view, see Driver (1989, 373ff), Flanagan (1990, 420ff), and Engel (2004, 204 and 212).

Although the standard view is appealing, there are apparent counterexamples. Rik Peels points out that certain Gettier-style cases seem to be counterexamples (2011, 352). Consider Bertrand Russell's stopped-clock example:

There is the man who looks at a clock which is not going, though he thinks it is, and who happens to look at it at the moment when it is right; this man acquires a true belief as to the time of day... (1948, 154).

Most agree that this man does not know the time is such and such. But he also seems not to be ignorant of this fact (especially if he relies on it in his practical reasoning). Generally speaking, it's well-attested that there are cases about which we're inclined to deny that a subject knows that P and also inclined to deny that she's ignorant that P. For example, according to Goldman and Olsson

It is highly inaccurate, inappropriate and/or misleading to characterize somebody who unjustifiedly believes (the fact that) p as being *ignorant* of p. Similarly, it is highly inaccurate, inappropriate, and/or misleading to characterize somebody who justifiedly believes p but fails to satisfy [the Gettier condition] as being ignorant of p (2009, 21).

The cases Goldman and Olsson describe here are paradigmatic examples of failing to know that P (at least under one sense of 'know'). So, the putative possibility of failing to know without being ignorant is already recognized in the epistemology literature.

But the above sorts of cases do not exhaust the possibilities. There are also cases where a subject doesn't even believe that P, and yet we still want to deny she's ignorant of the fact that P:

*Doctor Visit:* Suffering from a throat ailment, you visit your doctor. She claims that you most likely have strep throat, based on your symptoms, but that the final test results will come back in two days. However, because strep is contagious and can create further complications, she decides not to wait for the test results but to prescribe a strep-throat antibiotic now. At the end of your visit, you ask your doctor "Do you really believe I have strep?" Your doctor replies, "Well no, I cannot make that judgment at this point. We need to see the test results to rule out unlikely possibilities, such as tonsillitis. But, given your symptoms, it's very *likely* that you have strep, and it's best to proceed as if that's so. That's why I'm prescribing the antibiotic." Two days later, the test comes back positive—you indeed have strep. Fortunately, your doctor suspected this all along, and has already begun treatment.

Intuitively, your doctor is not ignorant of the fact that you have strep (before seeing the test results). After all, she believes it's likely that you have strep, and proceeds as if you in fact have it. Your doctor is *onto it* in a way that seems to disqualify her from counting as ignorant. For that reason, I shall call her 'Ontoit'. But it also seems that Ontoit does not know you have strep (before seeing the results). After all, we can interpret this case as one where she withholds belief about whether you have strep.<sup>4</sup> And, if knowledge that P requires belief that P, she thereby does not know that you have strep. To be sure, Ontoit certainly believes it's *likely* that you have strep. But that's not the required belief. In short, if these intuitive judgments are correct,<sup>5</sup> the standard view is mistaken—failing to know is not sufficient for ignorance.<sup>6</sup>

So far, I have only *conditionally* rejected the standard view. This is because I believe the above cases generate a puzzle rather than obvious counterexamples. After all, there is a widespread intuition in favor of the standard view, and we've seen no reason to distrust that intuition. So, if we rejected the standard view by relying only on our intuitive judgments about the above cases, we would be privileging one set of intuitions over the other, for no good reason.<sup>7</sup> But we've also seen no reason to distrust our intuitions about the above cases. So, if we rejected an intuitive judgment

---

<sup>4</sup> The case description does not by itself *require* us to understand the doctor as withholding belief. Perhaps she only withholds diagnosis, but not belief. Nevertheless, the point is that the case *can* be interpreted in the way I have proposed without altering our initial intuition about her lacking ignorance.

<sup>5</sup> Some may not share these intuitions. But universal assent is unlikely for any intuitive claim (indeed, I do not have the intuition that the standard view is true). Still, lack of universal assent shouldn't deter us. Those who don't share the conflicting intuitions assumed in this paper can still consider whether the intuitions of *others* can be adequately explained away.

<sup>6</sup> An alternative to the standard view has been defended by Rik Peels (2011, 345): For any truth P, S is ignorant of the fact that P if and only if S does not believe that P. However, Doctor Visit is equally a problem for Peels' view, because Ontoit lacks belief in the fact that you have strep without being ignorant of that fact. I do not here focus on Peels' view because it lacks the intuitive support had by the standard view—i.e. it's not so initially appealing that people feel comfortable assuming it without argument. So, it cannot be featured within a puzzle like the one below.

<sup>7</sup> Peels (2010; 2011) rejects the standard view but does not explain away its appeal. It is unsurprising that opponents, like LeMorvan (2011; 2012) and Goldman and Olsson (2009), embrace the standard view in spite of the opposing intuitions that Peels cites.

about one of those cases by relying on our intuitive judgment about the standard view, we would again be privileging one set of intuitions over the other, for no good reason. This is puzzling.

Moreover, the problem for the standard view is fairly general—it can arise from consideration of at least three different kinds of cases. As noted, Gettier-style cases and cases of unjustified true belief often seem to be counterexamples. Moreover, we can easily describe cases that fit the general mold of Doctor Visit, involving four main features (for a subject S and truth P):

- (i) S believes that it is probable that P.
- (ii) S withholds belief about whether P.
- (iii) Because of (i) and because there's practical significance in whether P, S acts as if P.
- (iv) S seeks further evidence to rule out certain possibilities in which  $\neg P$ .

These features suggest a general threat to the standard view.<sup>8</sup> Features (i) and (iii) make it hard to see how S could be ignorant of the fact that P; (ii) and (iv) make it hard to see how S could know that P. In this paper, I focus primarily on Doctor Visit, but many of my claims about Doctor Visit could be made similarly for Gettier-style cases and cases of unjustified true belief.

In short, rather than taking Doctor Visit as an obvious counterexample, we should see it as generating a puzzle, one that could arise with many other examples. Assuming you in fact have strep, it appears the following cannot all be true.

The New Ignorance Puzzle:

[Not-Ignorant] Ontoit is not ignorant of the fact that you have strep.

[New-Not-Know] Ontoit does not know that you have strep.

[Sufficiency] For any truth P, if S does not know that P, then S is ignorant of the fact that P.

---

<sup>8</sup> Although it's possible to interpret Doctor Visit in other ways, I shall interpret it as an instance of (i)-(iv). It might be thought that (i) and (iii) jointly entail belief, and thereby exclude (ii). But this is surely false. If anything, (i) and (iii) only entail that S has *some degree of confidence* that P, but this need not be enough confidence for *outright belief*—thus (ii) is not excluded by (i) and (iii). An example may help: Suppose I believe it's probable that this lottery ticket is a loser—satisfying (i)—and for that reason I sell my ticket to the highest bidder—satisfying (iii). I may still fail to *believe outright* that this ticket is a loser—satisfying (ii)—even if I have *some confidence* that it's a loser.

The third claim states the sufficiency condition of the standard view. The first and second claims state the intuitions about Doctor Visit.

On the surface, the New Ignorance Puzzle resembles the Old Ignorance Puzzle. Just like the Old Ignorance Puzzle, one could use any two of the above claims as an initially plausible argument against the third. Moreover, both puzzles center on a subject's failure to know (i.e. [New-Not-Know], [Old-Not-Know]). And both puzzles involve an axiomatic principle about knowledge, one that is often assumed without argument (i.e. [Sufficiency], [Closure]).

But my reason for addressing these puzzles together runs deeper than these surface similarities. As we will see, many traditional approaches to the old puzzle fail when applied to the new puzzle. Moreover, the best approach to the new puzzle points towards a neglected solution to the old one, an approach that rejects [Closure]. So, it's worth our while to address both puzzles together, and it is best that we begin with the New Ignorance Puzzle.

Finally, it is widely agreed that an adequate solution to the Old Ignorance Puzzle would either explain away the appeal of the rejected claim or explain why the three claims could all be true contrary to appearances. I shall assume the same for the New Ignorance Puzzle.

In the next two sections, I reject several solutions to the New Ignorance Puzzle, and then propose my own solution, which is a denial of [Sufficiency] and an attempt to explain away its appeal via conversational implicature. Finally, I show that two major reasons in favor of [Closure] can be explained away by relying on my solution to the New Ignorance Puzzle.

## **II. Failed Solutions to the New Ignorance Puzzle.**

*II.1.* One attempt at resolving the New Ignorance Puzzle would reject [Not-Ignorant], holding that Ontoit *is* ignorant of your having strep. Why then are we inclined to think she's not?

One response would hold that ‘ignorant’ is a negatively evaluative term, and that we’re reluctant to call Ontoit ‘ignorant’ because we don’t want to evaluate her in such a negative light.<sup>9</sup> The trouble is that this only explains why we’re *reluctant to assert the affirmative*—that Ontoit is ignorant of the fact that you have strep. It doesn’t explain why we’re *inclined to assert its negation*—[Not-Ignorant]. And the New Ignorance Puzzle arises because we’re inclined towards this negation.

One might think we’re inclined to assert the negation because we’re reluctant to assert the affirmative. But this is not enough to explain our inclination to assert the negation, since the negation is supposed to be false, and asserting it is not the only way to avoid the affirmative. Instead of (i) asserting the false negation, we could (ii) refrain from asserting either the affirmative or the negation. And the present view fails to explain why we’re inclined towards (i) rather than (ii). Consider an analogy: suppose you correctly believe that Max is fat but are reluctant to assert that he’s fat because this would insult him; furthermore, suppose you go on to assert its negation—that Max is not fat—which is false. We cannot explain why you asserted this negation, rather than refrain, until we have *additional information* from the context—such as that asserting the negation was the only way to avoid insulting him (i.e. remaining silent would also have insulted him). But there is no such additional information available with regard to Doctor Visit. We can avoid the affirmative and its evaluative connotation by simply refraining from asserting either the affirmative or the negation—for example, we can say ‘I’m unsure.’ Unlike the example involving ‘fat’, there’s no reason why we should think refraining is not an option with regard to Doctor Visit. So, the present

---

<sup>9</sup> Instead of citing the negative evaluation, one may hold that ‘ignorant’ connotes (without entailing) that the subject significantly lacks information or evidence. This too could explain why we are reluctant to call Ontoit ‘ignorant’. My reply to this alternative would be parallel to what I say below—it doesn’t explain why we’re inclined to assert the negation. So I do not discuss this alternative directly.



explanation fails to explain why we're inclined to assert the false [Not-Ignorant] rather than to refrain.<sup>10</sup>

A different approach to our puzzle comes from Pierre LeMorvan, who draws a distinction between factual and propositional ignorance in order to explain away our inclination to deny ignorance in a Gettier-style case (2012, 388). According to LeMorvan, a subject has factual ignorance when she's ignorant of a fact (i.e. of a proposition's truth-conditions being satisfied). A subject has propositional ignorance when she's ignorant of a proposition's truth-conditions; this requires that the subject is unable to entertain the proposition in question. Applying this distinction to Doctor Visit, notice that [Not-Ignorant] is explicitly stated as a denial of factual ignorance; notice also that Ontoit is obviously not propositionally ignorant. So, LeMorvan may claim that we're mistakenly inclined to accept the proposition expressed by [Not-Ignorant] because we confuse it with the following truth:

- (a) that Ontoit is not ignorant of *the truth-conditions* of the proposition that you have strep.

There are other possible approaches similar to LeMorvan's. Perhaps we're inclined to accept the proposition expressed by [Not-Ignorant] because we confuse it with the following truth:

- (b) that Ontoit is not ignorant of *the possibility* that you have strep.

But neither approach is sufficient to explain away our inclination to accept [Not-Ignorant]. The reason is simple. We are typically able to distinguish propositions like (a) or (b) from their corresponding propositions that deny factual ignorance, such as the one expressed by [Not-Ignorant]. Even brief reflection can make one aware of the general sorts of distinctions drawn

---

<sup>10</sup> One might claim that we're inclined to assert [Not-Ignorant] because we mistakenly *believe* it's true. But this only shifts the problem: we now need an explanation for why we mistakenly believe [Not-Ignorant]. And our reluctance to assert its corresponding affirmative is not enough to explain this. By analogy, the mere fact that you're reluctant to call Max 'fat' doesn't predict that you'll mistakenly come to believe he's not fat.

above. So, we need some reason to think we would be unable to distinguish these propositions with regard to Doctor Visit, Gettier-style cases, and cases of unjustified true belief. But it's unclear what this reason would be.

It's worth noting that the external-world skeptic could take an analogous approach to the Old Ignorance Puzzle, but we would find it unsatisfactory for the same reason. Recall that we're inclined to accept [Know], that Moore knows he has hands. But the external-world skeptic, who rejects [Know], could try to explain away our inclination in ways analogous to the above. Such skeptics can normally accept the following claims:

- (c) that Moore knows *the truth-conditions* of the proposition that he has hands.
- (d) that Moore knows *it's possible* that he has hands.

So this skeptic could claim that we're mistakenly inclined to accept the proposition expressed by [Know] because we confuse it with either (c) or (d). But I take it as obvious that this explanation is not sufficient to explain away our intuition that Moore knows he has hands. The reason is that we typically can distinguish between the sorts of propositions that the skeptic says we confuse, and it's unclear why we would be unable to distinguish them when considering the Old Ignorance Puzzle.

I see no reason to treat the New Ignorance Puzzle differently. With regard to the new puzzle, the proposed explanation distinguishes between various contents that can appear within the scope of the non-ignorance-operator. With regard to the old puzzle, the skeptic draws the same kinds of distinctions between contents that can appear within the scope of the knowledge-operator. But in both cases, we lack an explanation for why we'd be apt to blur the distinctions with regard to the relevant puzzle, but not apt to blur them in normal cases. So, the proposals aren't sufficient to explain what they aim to explain; and it's unclear how to supplement them to make them sufficient.

II.2. If [New-Not-Know] is false, then Ontoit knows you have strep. This is counterintuitive, but there have been attempts at explaining away [Old-Not-Know], and we might think these attempts could be applied analogously to [New-Not-Know].

According to Ernest Sosa (1999, 146ff; 2000, 15ff) the intuitive appeal of [Old-Not-Know] roots in the fact that we confuse the correct requirement for knowledge, namely

Safety: S knows that P only if S believes that P via method M and, were S to believe that P via M in close possible worlds, it would be true that P,

for a mistaken requirement:

Sensitivity: S knows that P only if S believes that P via method M and, were it false that P, S would not believe that P via M.

The embedded counterfactuals are easy to confuse, since they are contrapositions. But they're not equivalent, since counterfactuals do not contrapose. The obvious problem with this approach is that Ontoit does not satisfy either safety or sensitivity, because she withholds belief on whether you have strep. Of course, Ontoit may have some *confidence* that you have strep, even if her confidence is not enough for *outright belief*. So, we could modify safety and sensitivity by replacing 'believe' with 'have confidence'. However, even if Ontoit satisfies this modified version of safety, the fix is only temporary. As we've seen, the New Ignorance Puzzle could be formulated with a Gettier-style case, like Russell's stopped-clock example. The man in that example fails to satisfy safety (even on the modified version). So, this strategy does not solve the New Ignorance Puzzle.

Can we instead explain away the appeal of [New-Not-Know] by distinguishing between low- and high-grade factual knowledge, and by claiming that Ontoit has low-grade knowledge but lacks high-grade knowledge?<sup>11</sup> There are two main ways of holding that there is both low- and high-grade

---

<sup>11</sup> Since Ontoit withholds belief, this view is feasible only if low-grade knowledge does not entail belief. I allow this for the sake of argument, although many proponents of the present view hold that low-grade knowledge entails belief (e.g. see Goldman and Olsson (2009, 20) and Mylan Engel Jr. (2004, 229 fn 24)).

knowledge: we can claim that ‘know’ is ambiguous or that it’s context-sensitive. Let’s begin with the ambiguity view.

*II.2.a. Ambiguity.* On this view, we solve the New Ignorance Puzzle by positing an extra sense of ‘know’. The weak sense of ‘know’ refers to low-grade knowledge while its strong sense refers to high-grade knowledge. Under this assumption there are three ways of interpreting the New Ignorance Puzzle—weak, strong, and mixed interpretations. On a weak interpretation, all occurrences of ‘know’ are used in the weak sense, but [New-Not-Know] is false, because Ontoit has low-grade knowledge of your having strep. On the strong interpretation, all occurrences of ‘know’ are used in the strong sense, but [Sufficiency] is false, because failing to have high-grade knowledge does not suffice for ignorance. For both the weak and strong interpretations, we can explain away the appeal of the false claim by holding that this claim is true under the other interpretation. Lastly, there is a mixed interpretation where each occurrence of ‘know’ is used in a different sense, which means the puzzle equivocates. Here, the three claims are not genuinely inconsistent; we only think they are, because there are two other interpretations—the weak and strong—where those claims are genuinely inconsistent. So, on all three interpretations, the New Ignorance Puzzle can be solved.

The basic problem with the ambiguity view is that there is no reason to posit an extra sense of ‘know’ except for the fact that doing so may help solve puzzles like the ones we’ve been considering. To see why, let’s observe that the putative ambiguity of ‘know’ cannot be detected by our linguistic tests for ambiguity.<sup>12</sup>

---

<sup>12</sup> In this section, I ignore what Cruse calls “indirect tests for ambiguity”—tests that find two occurrences of a word with different meaning relations (e.g. synonymy, oppositeness) to other words. As Cruse notes, these tests are not reliable in distinguishing ambiguity from lack of specificity (1986, 54-7). I also ignore the definitional test, which would presuppose that we can give necessary and sufficient conditions for knowledge (Lewandowska-Tomaszczyk 2007, 143).

First consider the conjunction-reduction test. On this test, we take two sentences that use the term in its supposedly different senses, and then we conjoin the sentences in a way that uses the term only once. If the term is ambiguous, it is expected that the conjoined sentence will be awkward—i.e. zeugmatic (Zwicky and Sadock 1975, 17-8; Cruse 1986, 61). For example, ‘call’ is ambiguous. It can mean ‘request’ as in ‘Sue called him a taxi’, but it can also mean ‘label’ as in ‘Sue called him a bigot.’ If we conjoin these sentences while using ‘call’ only once, the result is awkward:

(1) # Sue called him a bigot and a taxi.<sup>13</sup>

The awkwardness of (1) suggests that ‘call’ is ambiguous. But now consider how the test applies to ‘know’. Proponents of the ambiguity view hold that ‘Moore knows that he has hands’ is true in the weak sense. And they would likely hold that ‘Moore knows that he exists’ is true in the strong sense. But the conjoined sentence does not seem the least bit awkward:

(2) Moore knows that he exists and that he has hands.

So, the conjunction-reduction test provides no reason to posit an extra sense of ‘know’.

The same is true for the ellipsis test. This test combines the two sentences by using the anaphoric ‘so do’ construction (Cruse 1986, 62). For example, ‘expire’ can mean ‘finish or run out’ as in ‘John’s license expired last month’. But it can also mean ‘die’ as in ‘John expired last month’. When we combine these two sentences anaphorically, the result is awkward (i.e. zeugmatic):

(3) # John’s license expired last month, and so did John.

But this test does not detect any ambiguity in ‘know’. The sentence ‘I know that Moore has a mind’ can be true in the weak sense, whereas ‘Moore knows that he has a mind’ can be true in the strong sense. However, the anaphoric combination is not the least bit awkward:

---

<sup>13</sup> Henceforth, let ‘#’ indicate that the sentence in question is linguistically awkward. For conjunction-reduction and ellipsis tests, the awkwardness is supposed to be of a particular kind, known as zeugma, which occurs when a single occurrence of a word must be interpreted in two distinct ways.

(4) I know that Moore has a mind, and so does Moore.

Of course, ambiguity theorists can explain why (2) and (4) are felicitous, because both sentences could be interpreted as expressing *only* the weak sense of ‘know’. Ambiguity theorists typically hold that the strong sense of ‘know’ entails the weak sense, which means it will always be possible to interpret the combined sentences as expressing only the weak sense. However, if this explanation works in the present instance, it only exacerbates the problem. It predicts that these tests could *never* detect the putative ambiguity of ‘know’. Since the strong sense of ‘know’ entails the weak sense, there should always be a felicitous reading of the combined sentences. But then the conjunction-reduction and ellipsis tests could not detect the supposed ambiguity of ‘know’ even in principle.

The contradiction test also fails to detect the putative ambiguity of ‘know’. According to Quine, “an ambiguous term [...] may be at once clearly true of various objects [...] and clearly false of them” (Quine 1960, 129). For this test, we try to observe a lack of contradiction in sentences that would be contradictory if they involved no ambiguous terms. Consider

(5) The minister married his sister, and he didn’t marry his sister.

(5) can be interpreted as non-contradictory, because ‘marry’ can mean something different in each conjunct. The first conjunct can mean that the minister performed the marriage ceremony, whereas the second can mean that he didn’t get married to his sister. By contrast, there appears to be no established way of reading (6) that would be non-contradictory:

(6) # Moore knows he has hands, and he doesn’t know he has hands.

Thus, the contradiction test also fails to detect any ambiguity in ‘know’.<sup>14</sup>

---

<sup>14</sup> Although ‘know’ fails the contradiction test, ambiguity theorists may reply. Matthias Steup claims that “‘know’ is an instance of hidden ambiguity because ordinary speakers are not even aware of the distinction between [low-] and [high-grade knowledge]. It’s only after some exposure to epistemology that people grasp that there is a difference here” (2005, 7). This would explain why (6) *seems* contradictory to ordinary speakers, even if it’s not. However, if ordinary speakers don’t recognize the putative difference between weak and strong senses of ‘know’, this

The above tests provide no reason to posit an extra sense of ‘know’. Can we instead posit an extra sense *solely* to solve the New and Old Ignorance Puzzles? No, doing so would violate Grice’s Razor: “Senses are not to be multiplied beyond necessity” (1989, 47). The minimal version of Grice’s Razor that I assume is this: other things being equal, we should not posit an extra sense of a word if we can explain the same data by citing forms of conversational implicature that we already have within our linguistic theory.<sup>15</sup> In sections III and IV, I show that we can solve these puzzles by appealing to widely-accepted forms of conversational implicature. If I’m right then we should not posit an extra sense of ‘know’.

Thus far, I have only examined the view that ‘know’ is ambiguous, but it’s also possible to posit an extra sense of ‘ignorant’. This view faces the same challenges outlined above—it fails the linguistic tests for ambiguity (which I shall leave the reader to observe), and Grice’s Razor warns against positing an extra sense if we can solve our puzzles via conversational implicature.

*II.2.b. Context-Sensitivity.* The most prevalent way of allowing for low- and high-grade knowledge is to hold that ‘know’ is context-sensitive.<sup>16</sup> It’s also possible to posit the context-sensitivity of ‘ignorant’. Under these postulates, sentences containing such words have different truth-values in different contexts, depending on which propositions are expressed. We could claim that [Not-Ignorant], [New-Not-Know], and [Sufficiency] are each true relative to some context—which might explain their plausibility—but they’re not all true relative to the same context.

---

suggests that at least one sense is not an *established* sense. If there are two senses at all, then one of them is likely a non-established, technical usage, and thereby cannot help solve puzzles that arise in ordinary language. Moreover, there are many people, including me, who have had exposure to epistemology but do not recognize the putative ambiguity of ‘know’. Thus, neither competence in English nor exposure to epistemology is sufficient to grasp this putative difference between senses. Why then should we think there *is* a difference?

<sup>15</sup> See Hazlett (2007, 674) for a similar version.

<sup>16</sup> The contextualist theory discussed in this section is attributor-contextualism, where the proposition expressed by a knowledge-utterance is dependent upon features within the attributor’s psychology and conversational or practical situation. For a consideration of certain variants on attributor-contextualism, see footnote 20.

I shall not contest the view that ‘know’ or ‘ignorant’ is context-sensitive. Instead, I will argue that, if either term is context-sensitive, then [Sufficiency] is false *in at least some contexts*.<sup>17</sup> Why is this significant?—because an appeal to context-sensitivity would help solve the New Ignorance Puzzle only if Doctor Visit is associated with a context where [Sufficiency] is true. But once we grant that there are contexts where [Sufficiency] is false, we’ll have every reason to think that Doctor Visit is associated with one of the “false” contexts. In what follows, I argue that the supposed context-sensitivity of ‘know’ or ‘ignorant’ only creates more room for counterexamples to [Sufficiency], which means an appeal to context-sensitivity does not help solve the New Ignorance Puzzle.

First, let’s suppose that only one of ‘ignorant’ or ‘know’ is context-sensitive while the other is invariant. Under this assumption, there are simple counterexamples to [Sufficiency] just by virtue of context shifts wherein one clause changes truth-value while the other remains the same.<sup>18</sup> Consider an analogy. On the present view, [Sufficiency] is treated like the following tallness principle: If S is tall, then S is taller than six feet. The predicate in the antecedent is context-sensitive while the one in the consequent is invariant. This principle is true in some contexts, but counterexamples arise when we shift to contexts with lower standards for tallness, such as a third-grade classroom where one can be tall without being taller than six feet. Since the antecedent

---

<sup>17</sup> I assume the antecedent and consequent of [Sufficiency] are to be evaluated with respect to the same context.

<sup>18</sup> To see this in detail, let’s first suppose that ‘ignorant’ is context sensitive, but ‘know’ is invariant. Take a context C1 such that [Sufficiency] is true in C1. In C1, suppose that S does not know P and is ignorant of P. Now shift to some other context C2 such that it is false that S is ignorant of the fact that P in C2. Since ‘know’ is contextually-invariant, it will still be true that S does not know P in C2, just as it was in C1. Thus, in C2, S does not know P and S is not ignorant of P—here we have a counterexample to [Sufficiency]. Now suppose ‘know’ is context sensitive, but ‘ignorant’ is invariant. Take a context C3 such that [Sufficiency] is true in C3. In C3, suppose that S knows that P and S is not ignorant of P. But now go to some other context C4 such that it is false that S knows that P in C4. Since ‘ignorant’ is contextually invariant, it will still be true that S is not ignorant of P in C4, just as it was in C3. Thus, in C4, S does not know that P and S is not ignorant of P—here, we have another counterexample to [Sufficiency]. Therefore, if one term is context-sensitive, while the other invariant, there will be counterexamples to [Sufficiency] in some contexts.



changes truth-value, while the consequent remains the same, there are easy counterexamples. The same problem arises for [Sufficiency], if only one of ‘ignorant’ or ‘know’ is context sensitive.

Alternatively, let’s suppose that both ‘ignorant’ and ‘know’ are context-sensitive. There are counterexamples to [Sufficiency] even under this assumption. But this time the counterexamples arise because the context-sensitivity of each term does not dovetail with that of the other. For comparison, consider another tallness principle where both predicates are context-sensitive: If S is tall then S is not short. It’s hard to see how there could be counterexamples to this conditional, because the standards for tallness and non-shortness dovetail with one another—i.e. when the standards for tallness increase due to context shifts, so do the standards for non-shortness, and vice-versa. But this is not true for knowledge and ignorance.

To see why, we could consider nearly any example aimed at illustrating the context-sensitivity of ‘know’. But for present purposes, let’s focus on Stewart Cohen’s:

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, “Yes I know—it does stop in Chicago.” It turns out that Mary and John have a very important business contact they have to make at the Chicago airport, Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John agree that Smith doesn’t really *know* that the plane will stop in Chicago. They decide to check with the airline agent (1999, 58).

According to Cohen, both knowledge claims are true relative to different contexts. Smith correctly claims to know that the plane stops in Chicago, and John and Mary correctly deny that Smith knows this. The reason is that Smith’s context of utterance is different from Mary and John’s. There is a context shift, and the standards for knowledge increase in Mary and John’s context.

Interestingly the standards for ignorance do not change. Let’s suppose it’s true that the plane stops in Chicago. In Smith’s context it seems correct for him to say ‘I’m not ignorant of

whether the plane stops in Chicago.’ But nothing changes when we shift to Mary and John’s context. That is, it would seem false for Mary to say ‘Smith is ignorant of whether the plane stops in Chicago’. And it would seem true for her to say ‘Smith isn’t ignorant of whether the plane stops in Chicago but I’d still like to double-check.’<sup>19</sup> The former seems false and the latter seems true, even if Mary mentions counter-possibilities such as the schedule having a misprint. Thus, assuming the standard contextualist treatment of this example, the context-sensitivities of ‘know’ and ‘ignorant’ do not dovetail. There are context shifts with regard to one term without the other, which means that [Sufficiency] is quite unlike our second tallness principle: If S is tall then S is not short. And most importantly, this provides an opening for counterexamples to [Sufficiency]. In Mary and John’s context, Smith fails to know without being ignorant of the fact that the plane stops in Chicago.<sup>20</sup>

In sum, if either ‘know’ or ‘ignorant’ is context-sensitive then [Sufficiency] is false in at least some contexts.<sup>21</sup> [Sufficiency] is now starting to look dubious. Before we reject it, we must see what else is puzzling about this principle.

---

<sup>19</sup> I have intentionally phrased these utterances in terms of ignorance of *whether P*, rather than ignorance of *the fact that P*. The latter presupposes that P is true, and it would be odd for Mary to presuppose that the plane stops in Chicago given that this very fact is being questioned in her context. However, I assume the context-sensitivity of the former is not substantially different from that of the latter. I also assume that, if P is true and S is not ignorant of *whether P*, then S is not ignorant of *the fact that P*.

<sup>20</sup> Variations on attributor-contextualism can be dealt with similarly. For example, according to contrastivism (Schaffer 2004), knowledge ascriptions denote a three-place relation—S knows that P rather than Q—and the value for Q can vary depending on the attributor’s context. On this view, we must ask whether ignorance-ascriptions also denote a three-place relation and vary by context. If they don’t, then there will be counterexamples to [Sufficiency]. If they do, then their contextual variations do not dovetail with that of knowledge-ascriptions, which also allows for counterexamples to [Sufficiency].

<sup>21</sup> Although there is much debate over whether ‘know’ is context-sensitive (e.g. Stanley 2004; 2005 ch. 2; Ludlow 2005; DeRose 2009 ch. 5), some of the main arguments against the context-sensitivity of ‘know’ do not apply with respect to ‘ignorant’. For example, it has been observed that ‘know’ lacks the key features exhibited by context-sensitive gradable expressions—expressions that can be combined with modifiers (e.g. ‘very’, ‘rather’, ‘really’) and are conceptually related to comparative constructions (e.g. ‘taller than’, ‘flatter than’) (Stanley 2004, 124). But ‘ignorant’ clearly has these features, as evinced by sentences like “Jill is rather ignorant of the fact that P” and “Jack is more ignorant of that fact than Jill is.” So, those who reject the context-sensitivity of ‘know’ on these grounds will have to find other reasons for resisting the claim that [Sufficiency] is false in some contexts.

II.3. Without considering the logical relationship between the following two enthymemes, ask yourself which, if any, seems intuitively more plausible:

Enthymeme (A)

It's true that P and S does not know that P

---

Therefore, S is ignorant of the fact that P

Enthymeme (B)

It's true that P and S is not ignorant of the fact that P

---

Therefore, S knows that P<sup>22</sup>

Most people, who don't notice the equivalence, claim that (A) seems more plausible than (B). A natural reaction when considering (B) is something like this—"It's not *that* easy to show that S has knowledge!" No corresponding reaction comes to mind when considering (A). The relevant sort of asymmetry is that it seems more plausible that (A) is valid than that (B) is valid. But this asymmetry is strange, since (A) and (B) are logically equivalent—one of them is valid if and only if both are.

Interestingly, philosophers often advance (A)-type enthymemes but rarely (B)-type enthymemes. As cited earlier, Gulley, Vlastos, Williams, and Williamson make *implicit* transitions from 'not know' to 'ignorant' without explicitly stating the principle they rely on. But it's rare to find implicit transitions from 'not ignorant' to 'knows'. Indeed, I can find only two instances of that inference—Goldman and Olsson (2009, 21) and Levinson (2000, 208)—but neither is enthymematic because both rely *explicitly* on the standard view.<sup>23</sup> Why do we commonly base our ignorance-claims solely on a failure to know, but rarely base our knowledge-claims solely on a lack of ignorance?

One might think the asymmetry arises because (B)'s premise involves a hidden double-negation, whereas (A)'s premise does not. This assumes that 'not ignorant' is semantically equivalent to a double-negation (e.g. 'not not know'). But even if we accept this assumption, it's not enough to

---

<sup>22</sup> For Enthymeme (B), we must stipulate that the speaker is not identical to S, otherwise the premise would presuppose the conclusion *without entailing it*. For example, my utterance of 'I am not ignorant of the fact that P' presupposes P, and my presupposing P suggests that I know P is true.

<sup>23</sup> For example, Levinson says that "*not ignorant* logically implies *knows* (because ignorance and knowledge are contradictories)..." (2000, 208). His parenthetical remark precludes the inference from being a (B)-type enthymeme.

explain the asymmetry between (A) and (B). Compare two more enthymemes, where the second enthymeme's premise involves a hidden double-negation:

Enthymeme (C)

P is not possible

---

Therefore, P is impossible

Enthymeme (D)

P is not impossible

---

Therefore, P is possible

These are also equivalent. And, although (C) is *more natural* than (D), it seems *equally plausible* that (C) and (D) are valid. So, the relevant sort of asymmetry is absent here. Merely citing the putative double-negation of 'not ignorant' is not enough to explain the asymmetry between (A) and (B).

The asymmetry between (A) and (B) closely relates to the New Ignorance Puzzle, because this puzzle is framed in terms of the more plausible transition from 'not know' to 'ignorant', as found in [Sufficiency]. The puzzle could not be framed *as* persuasively with the contraposition of [Sufficiency], involving the less plausible transition from 'not ignorant' to 'knows':

[Contra-Sufficiency] For any truth P, if S is not ignorant of the fact that P, then S knows that P.

Strangely, [Sufficiency] seems intuitively more plausible than [Contra-Sufficiency].

Even if we reject [Sufficiency] and explain away its appeal, as I hope to do, we still may fail to explain the above asymmetry. For example, some philosophers hold that 'know' is an inconsistent term (e.g. Weiner 2009, 2; Schiffer 1996). On this view, our semantic competence with 'know' disposes us to accept false principles. In like manner, we may attempt to solve the New Ignorance Puzzle by claiming that we're inclined to accept [Sufficiency], even though it's false, because our competence with 'know' or 'ignorant' disposes us to accept it. However, this solution does not entail an explanation of the asymmetry. Perhaps our semantic competence disposes us to accept [Sufficiency], even though it's false. This does not entail that we have a *weaker* disposition to

accept [Contra-Sufficiency]. More generally, why would our semantic competence give us a weaker disposition to accept ‘If  $\neg P$  then Q’ than its contraposition ‘If  $\neg Q$  then P’? The view in question does not by itself entail an answer.

It may be possible to supplement the solutions discussed so far with additional resources so as to explain the asymmetry (e.g. by citing pragmatic factors at play in one enthymeme but not the other). But it should be clear that none of these solutions by itself *entails* an explanation of the asymmetry. And in the next section, I provide a solution to the New Ignorance Puzzle that does entail an explanation. In this way, my solution simultaneously solves the puzzle and explains why the puzzle is more compelling with [Sufficiency] than its contraposition. Although my solution is independently motivated, I assume its ability to explain the asymmetry provides additional support in its favor.

To be sure, natural language is rife with similar asymmetries between equivalent formulations (e.g. ‘if P then Q’ might seem more plausible than ‘P only if Q’ in certain contexts). But the mere fact that such asymmetries are abundant doesn’t mean hypotheses cannot gain further support by explaining an asymmetry. Moreover, these additional asymmetries may be explicable by citing certain pragmatic factors like the ones I shall go on to cite (e.g. see Grice 1989, 24ff).

### **III. Preferred Solution to the New Ignorance Puzzle.**

So far I’ve argued that [Sufficiency] has counterexamples that cannot be explained away. And I’ve argued that if ‘know’ or ‘ignorant’ is context-sensitive then [Sufficiency] is false in at least some contexts. These are good reasons to reject [Sufficiency], if we can explain why we’re inclined to accept it, and why [Sufficiency] is associated with the above asymmetry. Let’s begin with the former.

In order to explain why we're inclined to accept [Sufficiency], we must first consider a related principle that is also entailed by the standard view of ignorance:

[Necessity] For any truth P, if S knows that P, then S is not ignorant of the fact that P.

[Necessity] is obviously true, and nothing I've said in this paper calls it into doubt.<sup>24</sup> The question I shall focus on is this: Why are we inclined to accept [Sufficiency] *in addition to* [Necessity]? My explanation appeals to two forms of conversational implicature—conditional perfection and negative strengthening.

Conditional perfection is famously illustrated by Michael Geis and Arnold Zwicky, who draw our attention to conditionals like

(7) If you mow the lawn, I'll give you five dollars.

Geis and Zwicky point out that “students in elementary logic courses” often make the mistake of formalizing these sentences “with biconditionals rather than conditionals” (1971, 561-2). According to Geis and Zwicky, there is something right about their mistake, because sentences like (7) are often interpreted as if they implicate another conditional:

(8) If you don't mow the lawn, I won't give you five dollars.<sup>25</sup>

And (7) and (8) entail the biconditional ‘I'll give you five dollars if and only if you mow the lawn’.

More generally, Geis and Zwicky note that there is a mistaken tendency to “perfect” or strengthen conditionals of the form ‘If P then Q’ by inferring the perfected conditional ‘If  $\neg$ P then  $\neg$ Q’.<sup>26</sup> It's worth noting that this error is nothing peculiar to “students in elementary logic courses.”

---

<sup>24</sup> Potential counterexamples will arise if we give a *de dicto* interpretation of P in one instance, but a *de re* interpretation in the other (e.g. Lois knows that superman can fly, but she's ignorant of the fact that Clark Kent, who is in fact superman, can fly). To avoid this, we must insist that the same interpretation be given to each occurrence of P.

<sup>25</sup> For ease of formulation, I occasionally speak as if sentences can conversationally implicate other sentences. But, technically speaking, it is the *utterance* of the sentence that conversationally implicates the *proposition* expressed by the other sentence.

<sup>26</sup> There is an equivalent inference from ‘If P then Q’ to ‘If Q then P’, which Laurence Horn calls “Aristotle's conversion” (2000, 298). Although I only appeal to conditional perfection in this paper, the explanation I propose can

Professional philosophers make the mistake too. Consider Eva Feder Kittay's analysis of Antioch's sexual offense policy. The Antioch Policy contains the following conditional:

If sexual contact and/or conduct is not mutually and simultaneously initiated, then the person who initiates sexual contact/conduct is responsible for getting the verbal consent of the other individual(s) involved (1992, 140).

Kittay draws the following conclusion directly from the above conditional:

So, if my partner and I simultaneously are seized with the desire to kiss, we don't need to say a thing as our mouths spontaneously move toward each other [...]. Therefore, [...] we are not always obliged to obtain verbal consent, not forced to "mix the pleasures of language with the pleasures of the body" *if* the sex is mutually initiated. (1997, 154).

The policy's conditional takes the form of 'If  $\neg P$  then  $Q$ '. And Kittay infers a conditional of the form 'If  $P$  then  $\neg Q$ '. The latter is the perfected version of the former (eliminating double-negation).

So, Kittay has perfected the conditional. This is no amateur mistake.

What exactly is the relationship between a conditional statement and its perfected version? According to most linguists (e.g. Levinson 2000, 37-8; Horn, 2000), the relationship is one of conversational implicature.<sup>27</sup> People make the mistake of conditional perfection when they confuse a conversational implicature for an entailment. I will assume this in the present paper.

But linguists often note that there are many types of conditionals that we have no tendency to perfect. For example, we're not inclined to perfect conditionals whose consequents presuppose their antecedents—'If there's a King of France, he's bald' (Horn 2000, 318). I eventually want to claim that we're inclined to accept [Sufficiency] because it's the result of perfecting [Necessity]. So, for this explanation to work, we need some reason to think that there's a tendency to perfect

---

be easily framed in terms of Aristotle's conversion, rather than conditional perfection, provided that we start with the *contraposition* of [Necessity]—For any truth  $P$ , if  $S$  is ignorant of the fact that  $P$ , then  $S$  does not know that  $P$ .

<sup>27</sup> Geis and Zwicky (1971, 656) actually hold that the inference is an implicature, but not a *conversational* implicature. However, the pervasive trend, following their paper, has been to treat the inference as conversational.

conditionals *like* [Necessity], even if we're not inclined to perfect all conditionals. This is why my explanation appeals to *negative strengthening* in addition to conditional perfection.

As Laurence Horn notes, we tend to perfect conditionals when there is negative strengthening from the antecedent to the consequent of the perfected conditional (2000, 314-17). What is negative strengthening? Negative strengthening occurs when a negation is combined with certain words (e.g. 'happy', 'good') and, as a result, this combination implicates something stronger than what is literally said. For instance, from (9) below, one can normally infer (10):

- (9) Kim is not happy.
- (10) Kim is unhappy.

But (9) does not entail (10), since (9) allows for the possibility that Kim is so-so, whereas (10) does not. But linguists typically hold that (10) is a conversational implicature of (9) (Horn 1989, 331; Levinson 2000, 127; Blutner 2004). After all, the connection between (9) and (10) is both cancelable and reinforceable. It is cancelable because we can felicitously *negate* (10) after uttering (9), as in 'Kim is not happy, but she's not unhappy either'. The connection is reinforceable because we can felicitously *affirm* (10) after uttering (9), as in 'Kim is not happy; in fact she's quite unhappy'. In sum, the connection between (9) and (10) illustrates a form of conversational implicature known as negative strengthening.

Now consider an obviously true conditional employing the antonyms 'happy' and 'unhappy':

- (11) If S is happy, then S is not unhappy.

I take (11) to be analogous to [Necessity]. And Horn (2000, 317) correctly notes that we're inclined to strengthen (11) by also accepting its perfected conditional:



(12) If S is not happy, then S is unhappy.<sup>28</sup>

The claim here is only that we are initially inclined to accept (12) prior to seeing that it has counterexamples (e.g. S could be so-so). And we are initially so inclined. Even though (12) has counterexamples, it nonetheless seems trivial at first glance, unlike other false conditionals such as

(13) If S is not happy, then S is tall.

Both (12) and (13) have counterexamples. But, on the present view, we are initially inclined to accept (12) because it is the result of perfecting (11) and there is negative strengthening from its antecedent to its consequent. Nothing similar is true of (13).

Together, the two mechanisms of conditional perfection and negative strengthening can explain why we are inclined to accept [Sufficiency] in addition to [Necessity]. Notice that [Sufficiency] is what we would infer if we were to perfect [Necessity]. Moreover, the consequent of [Sufficiency] can be inferred from its antecedent via negative strengthening. For instance, (14) seems to implicate (15) via negative strengthening:

(14) Socrates does not know that P.

(15) Socrates is ignorant of the fact that P.

First notice that the inference from (14) to (15) is reinforceable:

(16) Socrates doesn't know that P—in fact, he's positively ignorant of that.

Is the inference cancelable? The inference from 'not know' to 'ignorant' is cancelable, but *not easily cancelable*. It can only be canceled in select contexts.<sup>29</sup> First consider that a straightforward attempt at cancellation seems to fail:

---

<sup>28</sup> Technically, Horn claims that we infer (12) from the contraposition of (11)—'If S is unhappy, S isn't happy'. This step is not an instance of conditional perfection, but is an equivalent inference known as Aristotle's conversion. See footnote 26 for explanation of Aristotle's conversion.

<sup>29</sup> An inference from 'A' to 'B' might be cancelable (and might be a conversational implicature), even if this inference cannot be canceled in *all* contexts where 'A, but not B' is uttered. It is enough if there are only *some* contexts where this inference is canceled. See Blome-Tillmann (2008) for discussion.

(17) # Socrates doesn't know that P, but he's not ignorant of the fact that P.

(17) seems odd. The implicature from 'not know' to 'ignorant' does not seem to be cancelled in this instance. But if we describe the right context of utterance, and replace P with the right propositional content, then the implicature can be cancelled. To illustrate this, let's suppose we're throwing a surprise party for Socrates, but I find out that Socrates has a hunch that we're doing this; you ask me whether Socrates knows about the surprise, and I reply in one of two ways:

(18) Socrates doesn't *know* we're throwing a surprise, but he isn't exactly ignorant of this either.

(19) Socrates doesn't *know* we're throwing a surprise, but I wouldn't say he's ignorant of this.

(18) and (19) seem normal in this context, which means the inference from 'not know' to 'ignorant' is cancelable. So, it's plausible that negative strengthening occurs from 'not know' to 'ignorant'.<sup>30</sup>

In sum, we correctly accept [Necessity] but mistakenly infer [Sufficiency] by way of conditional perfection. The reason we're inclined to perfect [Necessity], without being inclined to perfect *all* conditionals, is that the antecedent of [Sufficiency] conversationally implicates its consequent via negative strengthening.<sup>31</sup> Since conversational implicature is defeasible, this

---

<sup>30</sup> I also hold that equivalent sentences, like 'S does not have knowledge', 'S fails to know', and 'S lacks knowledge', will also have this implicature. This follows from the non-detachability of conversational implicature (i.e. truth-conditionally equivalent statements carry the same conversational implicatures).

<sup>31</sup> This explanation can be stated more precisely if we distinguish between sentences and propositions. Let <Sufficiency> and <Necessity> stand for the propositions expressed by [Sufficiency] and [Necessity] respectively. The explanation has two parts. **Part I:** We intuit <Sufficiency> because we accept the truth of [Necessity] and we mistakenly interpret [Necessity] as expressing the conjunction of <Necessity> and <Sufficiency>. **Part II:** We mistakenly interpret [Necessity] as expressing this conjunction because [Necessity] conversationally implicates <Sufficiency> via conditional perfection and negative strengthening. In Parts I and II, 'we' should technically refer only to English speakers, since [Necessity] is an English-language sentence. However, conversational implicatures are non-detachable, which means that 'we' could refer to non-English speakers whenever [Necessity] is replaced with its non-English equivalents. Moreover, the conversational implicatures in question are what Grice calls "generalized conversational implicature," which can be detected even when the sentence is not uttered in context (Grice 1989, 37). According to Levinson, these implicatures are "hard to distinguish from the *semantic* content of linguistic expressions, because such implicatures will be routinely associated with the relevant expressions in all ordinary contexts" (1983, 127). It's worth noting that Part II parallels the traditional Gricean explanation of why people interpret 'P or Q' as expressing the conjunction of <P or Q> and <It's not the case that both P and Q>.

explanation leaves room for counterexamples to [Sufficiency], like Doctor Visit. In a nutshell, this explanation holds that we have confused a conversational implicature for an entailment. Other instances of this type of mistake are revealed by Grice, who points out that certain inferences drawn from ‘either...or’ and ‘if...then’ statements are actually conversational implicatures, even though they have been mistakenly seen as entailments (1989, 8ff). Lastly, it’s worth noticing that the present solution only cites conversational implicatures that we already need within our general linguistic theory. So, we need not posit extra senses of ‘know’ or ‘ignorant’, and can thereby respect the modest form of Grice’s Razor mentioned earlier.

It might be objected that we are more strongly inclined to accept [Sufficiency] than we are analogous conditionals like (12). However, our stronger inclination to accept [Sufficiency] is due to at least two related factors. First, the inference from ‘not happy’ to ‘unhappy’ is easily cancellable, whereas the inference from ‘not know’ to ‘ignorant’ is not. And second, the counterexamples to [Sufficiency] are much less obvious than they are with respect to (12). It takes little reflection to see that (12) is false because of the obvious possibility that S is just so-so. But counterexamples to [Sufficiency], such as Doctor Visit and Gettier-style cases, are not the least bit obvious. These two factors—the difficulty of cancelation and the obscurity of counterexamples—are at least partly responsible for our stronger inclination to accept [Sufficiency].

Further support for my solution comes from the fact that it explains the asymmetry noted in II.3. Negative strengthening functions asymmetrically. Notice that (21) cannot normally be inferred from (20):

- (20) Kim is not unhappy
- (21) Kim is happy.

This illustrates a general pattern of asymmetries that holds for many other antonym pairs (Levinson 2000, 142; Blutner 2004, 499-500). Of course, there are unusual contexts where (21) can be inferred

from (20). For example, (20) can be interpreted as an understatement, especially when the speaker uses a rise-fall tone of voice. In this context one may infer that Kim is extremely happy. Still, there is no available inference from ‘not unhappy’ to ‘happy’ without the help of contextual factors indicating an understatement (e.g. rise-fall tone of voice).<sup>32</sup>

How then should we characterize the asymmetry associated with negative strengthening? We must first distinguish between marked and unmarked terms. Marked terms generate certain assumptions that their antonyms do not generate. For example, the question ‘How unhappy is she?’ assumes that she is unhappy, whereas ‘How happy is she?’ leaves open the possibility of her being either happy or unhappy. Of these two antonyms, ‘unhappy’ is marked whereas ‘happy’ is unmarked, because ‘unhappy’ generates certain assumptions that ‘happy’ does not generate. We can see that ‘know’ is unmarked by virtue of a yes/no question—‘Does she know that?’—which leaves open the possibility that she is ignorant as well as the possibility that she knows. But the opposite question—‘Is she ignorant of that?’—suggests some lack of awareness. Thus, ‘knows’ is unmarked whereas ‘ignorant’ is marked. In the following, let U be the unmarked term within an antonym pair, while M is the marked term; and let Ux and Mx be sentences that apply these antonyms to an item x. For our purposes, we can characterize the asymmetry as follows.

Asymmetry Principle:

If  $\neg Ux$  conversationally implicates Mx via negative strengthening, then  $\neg Mx$  typically does not conversationally implicate Ux, unless there are contextual factors indicating that the speaker is making an understatement.

We can now explain the asymmetry associated with knowledge and ignorance. My solution to the New Ignorance Puzzle integrally involves the claim that ‘not know’ conversationally implicates ‘ignorant’ via negative strengthening. Since ‘know’ is unmarked and ‘ignorant’ is marked, the

---

<sup>32</sup> See Horn (1989, 356-57) for discussion on of a similar pattern with regard to ‘not bad’.

Asymmetry Principle predicts that ‘not ignorant’ doesn’t conversationally implicate ‘knows’ (unless the context indicates an understatement).<sup>33</sup> Because there is typically no implicature from ‘not ignorant’ to ‘knows’, Enthymeme (B) seems less plausible than Enthymeme (A). Similarly, [Contra-Sufficiency] seems less plausible than [Sufficiency] because the antecedent of [Contra-Sufficiency] typically does not implicate its consequent.

#### IV. The Old Ignorance Puzzle.

I have argued that [Sufficiency] is false even though appealing. I now tell a related story about how [Closure] might be false even though appealing. To be clear, I do not argue that [Closure] *is* false. I instead use the resources of the previous section to argue that two ways of supporting [Closure]—from intuition and abominable conjunctions—can be explained away.<sup>34</sup>

*IV.1. Intuitive Appeal.* Possibly the most powerful consideration in favor of knowledge-closure is its intuitive appeal. Even closure-deniers, like Robert Nozick (1981, 206) and Fred Dretske (2005, 18), admit that closure principles are intuitively plausible. This alone does not put them at a disadvantage, because all solutions to the Old Ignorance Puzzle must reject at least one of our deeply-seated intuitions.<sup>35</sup> Still, most will agree that an important part of a successful resolution is that it can adequately explain away the appeal of the rejected claim. And a major problem for those who want to resolve the Old Ignorance Puzzle by rejecting [Closure] is that there are few hypotheses

---

<sup>33</sup> An example of an understatement involving ‘not ignorant’ might be ‘Socrates is not *ignorant* of the fact that P—he’s no idiot!’ In this case, we clearly are licensed to infer that Socrates knows that P.

<sup>34</sup> It is unclear how the ideas put forth in this paper would provide a new response to Hawthorne’s “Tortoise” argument for knowledge-closure (2005 32). So I omit discussion of that argument here. But see Kvanvig (2006, 262-64) for a reply to Hawthorne.

<sup>35</sup> Even contextualists are committed to something counterintuitive. Some of them abandon [Closure] (e.g. Heller 1999). Most others, like Cohen (1999, 67), are committed to the counterintuitive claim that [Old-Not-Know] is false in ordinary contexts.

for explaining away the pervasive intuition in favor of [Closure].<sup>36</sup> I shall now present a new hypothesis, and will argue that it has significant plausibility, barring any empirical tests that might confirm or disconfirm it.

Recall that the standard view of ignorance is equivalent to two claims—[Necessity] and [Sufficiency]. I have argued that [Sufficiency] is false, albeit intuitive. [Necessity], on the other hand, is obviously true and I will continue to assume its truth. Given that [Necessity] is true and [Sufficiency] is false, we can now use [Closure] to generate a weaker principle that is about *losing ignorance* rather than *gaining knowledge*. More precisely, [Closure] combined with [Necessity] entails

[NI] If S knows that P, and knows that P entails Q, then S is not ignorant of the fact that Q.

[NI] is what I call a no-ignorance principle. The basic idea is that we cannot be ignorant of facts that we know to follow from other facts we know. This is similar to the basic idea behind closure, but, as I'll show, no-ignorance principles are importantly different.

What is the relationship between [Closure] and [NI]? [Closure] and [NI] are equivalent if and only if the standard view of ignorance is true. But since the standard view is false, they are not equivalent. Instead, [NI] is logically weaker than [Closure]. [Closure] entails [NI], since [Necessity]

---

<sup>36</sup> One explanation might be extracted from Heller (1999, 205ff) who rejects closure but admits that closure is true in some contexts. It could be claimed that closure's truth in some contexts can make it seem as though closure is true in all contexts. However, the explanation is insufficient as it stands. There are some contexts in which 'If x is tall, then x is taller than 7 feet' is true, but its truth in some contexts doesn't make this principle seem true in general. Another attempt at explaining away closure comes from Cohen (on behalf of Dretske and Nozick): "I suppose Dretske and Nozick could say that it is our failure to appreciate the deep truth about the nature of knowledge as revealed by their theory. We mistakenly believe in the closure principle because we fail to see how the tracking condition for knowledge falsifies it" (1999, 64). This proposal fails to explain the appeal of closure for reasons not mentioned by Cohen. The supposed fact that "we fail to see how the tracking condition... falsifies" closure only predicts that we will *either* believe in closure *or* suspend judgment on whether closure is true. But since the second disjunct is a possible outcome, the proposal thereby does not explain why we're inclined to accept closure.

is true.<sup>37</sup> But [NI] does not entail [Closure], since [Sufficiency] is false. Thus, despite initial appearances, the two principles are not equivalent—[NI] is weaker than [Closure].

Many people have the intuition that the proposition expressed by [Closure] is true. If it turns out to be false, the following is a way of explaining their mistake. Henceforth, let all bracketed expressions (e.g. ‘<...>’) stand for the proposition expressed by the sentence in question:

Closure Hypothesis:

We intuit <Closure> because we intuit <NI> and that <NI> is equivalent to <Closure>.

A few points should be explained. First of all, it is indeed plausible that we intuit <NI>. After all, <NI> is obviously entailed by <Closure>, which means that anyone who is intuitively inclined towards <Closure> should also find <NI> intuitive. And second, we presumably also have the intuition that <Closure> and <NI> are equivalent—this is because we are inclined to accept the standard view of ignorance. But this equivalence intuition is where our mistake lies. Because the standard view is false, these two propositions are not really equivalent, contrary to appearances.

As stated, the Closure Hypothesis explains away the intuition about one version of knowledge-closure. But there are other versions, some of which are more plausible than [Closure]. Still, the way in which [NI] was generated from [Closure] can be applied to any knowledge-closure principle. We can combine any such principle with [Necessity] in order to yield a weaker no-ignorance principle.<sup>38</sup> Since every knowledge-closure principle entails a weaker no-ignorance principle, there is always a hypothesis available to those who wish to reject the knowledge-closure

---

<sup>37</sup> This inference assumes that knowledge is factive.

<sup>38</sup> The no-ignorance principle may take on a different form depending on the closure principle from which it's generated. For example, some closure principles have a weaker consequent. They replace 'knows' in the consequent with 'is in a position to know'. In this case, the corresponding no-ignorance principle will say something about being in a position to lose one's ignorance. But, just as before, this sort of no-ignorance principle would be entailed by its corresponding closure principle via [Necessity].

principle employed by the skeptic.<sup>39</sup> Sometimes the skeptic does not employ a *general* closure-principle, but instead employs a *particular instance* of such a principle (e.g. ‘If Moore knows he has hands then he knows he’s not a BIV’). But, given [Necessity], each particular instances of closure entails a particular instance of a no-ignorance principle (e.g. ‘If Moore knows he has hands then he’s not ignorant of the fact that he’s not a BIV’). So, we can adjust the Closure Hypothesis by replacing <Closure> and <NI> with particular instances of these principles. In sum, whether the skeptic employs a closure principle, or a particular instance of it, the Closure Hypothesis provides a new way of explaining why the skeptic’s reasoning is appealing albeit misleading. By accepting a knowledge-closure principle (or its particular instance), in addition to its corresponding no-ignorance principle (or its particular instance), the skeptic has confused two very similar things: he has mistaken our *loss of ignorance* for a *gain in knowledge*.

In proposing the Closure Hypothesis on behalf of closure-deniers, I am assuming that closure-deniers will happily accept no-ignorance principles. However, they often reject [Closure] because they see this as the best way out of the Old Ignorance Puzzle (e.g. Dretske 2005, 23).<sup>40</sup>

---

<sup>39</sup> These claims will apply equally to the closure principles that John Hawthorne relies on (2005, 31ff). Hawthorne provides an argument based on two closure principles, and aims to problematize the closure-denier’s view by showing that they are committed to denying one of these principles. The two principles are as follows:

*Equivalence:* If S knows a priori (with certainty) that P is equivalent to Q and knows that P, and competently deduces Q from P (retaining the knowledge that P), then S knows that Q.

*Distribution:* If S knows the conjunction of P and Q, then, as long as S is able to deduce Q, S is in a position to know that Q.

These two principles are intuitively plausible, and the closure-denier’s reasons for rejecting closure typically do not lead them to rejecting the first principle. Hawthorne then goes on to show that these two principles lead to the type of conclusion that the closure-denier wants to avoid—namely that one can know that one is not a BIV. The closure-denier is therefore committed to rejecting at least one of these principles, most likely Distribution. My response above applies equally to Hawthorne’s argument. Equivalence and Distribution each generate a no-ignorance principle that the closure-denier can accept. And this means that my general strategy for explaining away the intuitive appeal of [Closure] can be applied also to the intuitions upon which Distribution and Equivalence are based. If the closure-denier is committed to denying one of them, then there will be a suitable version of the Closure Hypothesis that can explain away its appeal.

<sup>40</sup> To be sure, there are reasons for rejecting *particular versions* of knowledge-closure, such as [Closure], and these reasons might also license the rejection of the corresponding no-ignorance principle, like [NI]. (See David and Warfield 2008 for different problems for specific versions of closure). My claim is only that the reasons traditionally cited against *the whole class* of knowledge-closure principles are not reasons for rejecting *the whole class* of no-ignorance principles. Most arguments against knowledge-closure, such as those in Dretske (2005), are obviously irrelevant to no-ignorance



Shouldn't they also reject [NI] for a similar reason? In other words, isn't there an equally difficult puzzle involving [NI] instead of [Closure]? To see why there is not, consider this modified puzzle:

The Modified Old Ignorance Puzzle:

[Know] Moore knows that he has hands.

[Ignorant] Moore is ignorant of the fact that he's not a BIV.

[NI] If S knows that P, and knows that P entails Q, then S is not ignorant of the fact that Q.

The trouble is that [Ignorant] does not seem intuitively plausible, unlike the original [Old-Not-Know]. To see this, suppose that Moore believes he's not a BIV on the basis of experiential evidence, and that this belief is true. Here, it seems strange to say that he's ignorant of the fact that he's not a BIV.<sup>41</sup>

Since [Ignorant] is not intuitively plausible, it must be supported by argument before the Modified Puzzle should pressure closure-deniers into rejecting [NI]. What argument could there be? One argument for [Old-Not-Know] might also apply to [Ignorant]. It is sometimes thought that knowledge of P requires infallibility—that the subject can eliminate all relevant possibilities in which  $\neg P$  (Lewis 1996, 549ff). And the possibility that Moore is a BIV is relevant to whether he knows he's not a BIV.<sup>42</sup> But Moore cannot eliminate this possibility, because all his experiences would be the same if the possibility were to obtain. This infallibility-based argument might support [Old-Not-

---

principles. So I omit discussion of them. The only one that might seem to apply also to no-ignorance principles is the challenge discussed in this paragraph.

<sup>41</sup> Some may have the intuition that Moore is ignorant of *whether* he's a BIV. I don't share this intuition, especially if we stipulate that Moore believes he's not a BIV. But even so, there's a distorting factor in the whether-formulation. Unlike [Ignorant], the whether-formulation doesn't presuppose that Moore is not a BIV. As such, those who contemplate the whether-formulation could be leaving it open as to whether Moore is a BIV; and if he is a BIV, then he would be ignorant of whether he's a BIV, since he would have the false belief that he's not one. This distorting factor is absent from the Modified Puzzle as I've stated it, because [Ignorant] presupposes that Moore is not a BIV. The distorting factor is also absent from the original Old Ignorance Puzzle, because [Old-Not-Know] presupposes that Moore is not a BIV, given that factive verbs (e.g. 'knows') presuppose their embedded contents when the verb is negated (e.g. 'S does not know that P' presupposes P).

<sup>42</sup> This is true according to Lewis' rule of attention—"if in this context we are not in fact ignoring but attending to it, then for us now it is a relevant alternative" (1996, 559 and 564). It's also true according to DeRose's variation, where the skeptical possibility becomes relevant when it occurs within the scope of an epistemic operator—"S does not know that..." (1995, 15 fn 22).

Know], but no similar argument could support [Ignorant] because it's possible to lack ignorance without meeting the infallibility requirement. Recall that Ontoit lacks ignorance of the fact that you have strep even though she's unable to eliminate all relevant possibilities in which you don't have strep, such as the possibility of you having tonsillitis.<sup>43</sup> Thus, infallibility is not a requirement for lacking ignorance, even if it's a requirement for knowledge. In short, I can see no reason—from argument or intuition—that should lead closure-deniers to accept [Ignorant], which means they shouldn't feel pressed to reject [NI] in escaping the Modified Puzzle.

Another objection to the Closure Hypothesis might go as follows: If the standard-view is false, then the concept of ignorance is distinct from the concept of not-knowing. So, it should be possible for someone to intuit <Closure> without even employing the concept of ignorance. The Closure Hypothesis couldn't explain this person's intuition.

It would be foolish to deny that the objector's scenario is *possible*. In general, an intuition could arise from any number of causes (e.g. a bump on the head, hallucinogens, etc.); and the Closure Hypothesis is not meant to exhaust all possible causes. It is only if the objector's scenario is *actual* that it presents a problem for the Closure Hypothesis. After all, the Closure Hypothesis is an empirical hypothesis about the genesis of our actual intuition that <Closure> is true. As such, the 'we' in the Closure Hypothesis refers only to actual persons who intuit <Closure>, not to all possible persons with this intuition. If there is an *actual* person who reports that she doesn't employ the concept of ignorance when intuiting <Closure>, this report would supply empirical evidence against the Closure Hypothesis, provided we accept the report as accurate. We've seen that there's reason to be skeptical of such reports, given that we're prone to confuse knowledge with non-

---

<sup>43</sup> There are differing views about what determines the relevance of a possibility, but none would make this possibility irrelevant. The possibility of you having tonsillitis was explicitly mentioned by Ontoit, has practical significance for all parties involved, and is likely enough that Ontoit seeks further evidence in order to exclude it.

ignorance. But such reports cannot be ruled out *a priori*. And if they actually occur, then they would require us either to reject the Closure Hypothesis or to modify it (e.g. by allowing other sources as giving rise to the closure-intuition).

But the same is true of any empirical hypothesis aimed at explaining away an intuition. If someone actually reports that she doesn't employ the concept that supposedly gives rise to her mistake, then this report would tell against the hypothesis, provided we grant the report's accuracy. For example, recall that Sosa holds that we mistakenly intuit <Old-Not-Know> because we employ our concept of sensitivity (which we confuse with the correct condition on knowledge—safety). If someone reports that she doesn't employ the concept of sensitivity when intuiting <Old-Not-Know>—but instead employs, say, the concept of infallibility—then this report would supply empirical evidence against Sosa's hypothesis, provided we grant its accuracy. The same is true of contextualism. Many of us believe that [Old-Not-Know] is true in ordinary contexts, and contextualists often claim we are mistaken (e.g. Cohen 1999, 67). On the typical view (e.g. DeRose 2009, 155ff), we make this mistake because we apply our high-standards concept of knowledge within what we mistakenly believe to be an ordinary context. But if a person reports that she's not employing her high-standards concept when intuiting the truth of [Old-Not-Know] then her report would be evidence against the contextualist hypothesis, provided we grant its accuracy.<sup>44</sup>

In general, the possibility of empirical counterevidence to an empirical hypothesis cannot be ruled out *a priori*. And I see no *a priori* reason to think we would grant the accuracy of reports

---

<sup>44</sup> One might wonder how a person could distinguish between the high- and low-standards concepts of knowledge. Contextualists typically hold that the high-standards concept entails a concept such as sensitivity (DeRose 1995) or unrestricted infallibility (Lewis 1996). So, for the same reason mentioned with regard to Sosa's hypothesis, if a subject reports that she's not employing the concept of sensitivity/infallibility, then this report would be evidence against the view in question.

weighing against the Closure Hypothesis, but not grant the accuracy of reports weighing against rival hypotheses, like Sosa's or the contextualist's.

One might argue that there's *a priori* reason to think empirical counterevidence *will* arise against the Closure Hypothesis. Surely there is an actual language L with a term equivalent to our 'know' but no term equivalent to our 'ignorant'. And, according to this objection, some unilingual speakers of L would lack the concept of ignorance and thereby fail to employ it when intuiting <Closure>. However, even if such speakers were to lack the concept of ignorance (which is by no means obvious<sup>45</sup>), we still have no *a priori* reason to think they'll have the intuition that <Closure> is true. For all we know, they might find <Closure> unintuitive or intuit that <Closure> is false. So, there's no *a priori* reason to think such counterevidence will arise against the Closure Hypothesis.

It should be clear that I am neutral on whether the Closure Hypothesis can be empirically verified. I have only argued that this hypothesis has significant plausibility barring any empirical tests that might confirm or disconfirm it. But even this modest conclusion is a big step forward for closure-deniers. And I take it as obvious that proponents of rival hypotheses (e.g. Sosa and the contextualists) must be equally modest, at least until their hypotheses are empirically tested.

*IV.2. Abominable Conjunctions.* Keith DeRose (1995, 28) has objected that closure-deniers, like Robert Nozick, must embrace "abominable conjunctions" such as

(22) # Moore knows he has hands, but he doesn't know he's not a handless-BIV.

---

<sup>45</sup> Joseph Greenberg has observed that certain African, Amerind, and Oceanic languages have an expression equivalent to our 'good' with no expression equivalent to our 'bad'. But it turns out that competent speakers of these languages still have the concept of badness. According to Greenberg, such speakers convey that something is bad by negating their term for 'good' (1966, 52). According to Horn, this information is conveyed by conversational implicature (Horn 1989, ch. 5, 336, for discussion). But for these speakers to convey that something is bad, and to understand this when it is conveyed, they must possess the concept of badness, even if there's no term for it in their language. These languages are closely analogous to language L. So, it's far from clear that unilingual speakers of L would not lack the concept of ignorance.

People who escape the Old Ignorance Puzzle by rejecting [Closure] typically hold that the other claims within the puzzle—[Know] and [Old-Not-Know]—could be true within the same context. These people must therefore grant that (22) could be true. But (22) seems *linguistically awkward*. So, it may be charged that there is no satisfactory way of explaining its linguistic impropriety without accepting a knowledge-closure principle.<sup>46</sup>

This objection brings up an important question: How would we explain the infelicity of (22) if we *did* accept closure? Initially, we might think that (22) is infelicitous because it's contradictory. The trouble is that one cannot derive a contradiction from (22) without some contingent background assumptions. For example, one might initially think (22) is contradictory because the following version of closure is necessarily true:

[Simple-Closure] If S knows that P, and P entails Q, then S knows that Q.

But most will agree that [Simple-Closure] is false (e.g. David and Warfield 2008, 142). Someone could know that P, but fail to realize that P entails Q, and thereby fail to know that Q. So, this explanation assumes a principle that's vulnerable to counterexample. Our original principle—[Closure]—can avoid this counterexample. The trouble is that [Closure] does not *by itself* allow us to claim that (22) is contradictory. We must also assume contingent background information, such as

(BG) that Moore knows that his having hands entails he's not a handless-BIV.<sup>47</sup>

Without assuming (BG), one cannot derive a contradiction from (22), even if [Closure] is true.<sup>48</sup>

---

<sup>46</sup> Peter Ludlow (2005, 35) and Mark Heller (1999, 198) also understand the problem as (at least partly) a call to explain the linguistic impropriety of abominable conjunctions.

<sup>47</sup> Alternatively, without accepting a general closure principle, like [Simple-Closure] or [Closure], one might instead assume a specific instance of such a principle as background information. For instance, one may assume the following: If Moore knows that he has hands, then Moore knows that he's not a handless-BIV. But the same problem arises for this approach—this specific instance of closure is only a contingent truth that must be assumed as background, before a contradiction can be derived from (22).

<sup>48</sup> Even [Closure] is in need of much improvement before it can avoid counterexamples (see David and Warfield 2008, 143ff). Nonetheless, my basic point in this paragraph applies for all plausible versions of knowledge-closure. We must assume certain contingent truths, like (BG), before a contradiction can be derived from (22).

How then does (BG) become part of the closure-accepter's explanation of the infelicity of (22)? Most likely, closure-accepters will appeal to the common ground. The common ground of an utterance is the mass of knowledge, beliefs, and suppositions that conversational participants believe they share with one another (Clark 1996, 93). So, closure-accepters can claim that (BG) is in the common ground of typical utterances of (22). And the first conjunct of (22), along with that common ground information, contradicts the second conjunct, provided [Closure] is true. More generally, the closure-accepter could provide the following explanation:

Closure-Accepter's Explanation:

Abominable conjunctions are odd because one of the conjuncts (together with the common ground) contradicts the other conjunct.

This explanation is initially plausible. Can closure-deniers match it? I shall argue that they can provide an explanation that is at least as good.

Instead of appealing to a closure principle, the closure-denier can appeal to the corresponding no-ignorance principle along with conversational implicature. As argued earlier, 'not know' conversationally implicates 'ignorant'. So, the second conjunct of (22) conversationally implicates that Moore is ignorant of the fact that he's not a handless-BIV. This implicature helps to explain the oddity of (22). Suppose that a no-ignorance principle like [NI] is true. In this case, the first conjunct of (22), together with certain common ground information like (BG), contradicts the implicature of the second conjunct. More generally, closure-deniers can hold

Closure-Denier's Explanation:

Abominable conjunctions are odd because one of the conjuncts (together with the common ground) contradicts a conversational implicature of the other conjunct.

Here, the closure-denier assumes a no-ignorance principle in place of closure.

There are other infelicitous conjunctions that would be explained by the closure-denier's explanation. For example, (23) below is equally as odd as (22):

(23) # The U.S. has won all wars, and they tried to defeat North Vietnam.

The two conjuncts in (23) are not contradictory. Instead, the oddity of (23) stems from a conversational implicature generated by the second conjunct. Saying that S tried to  $\varphi$  conversationally implicates that S did not succeed in  $\varphi$ -ing (Huang 2007, 44). So, the second conjunct of (23) conversationally implicates that the U.S. did not succeed against North Vietnam. But this implicature contradicts the first conjunct of (23), provided certain common ground information (e.g. that the conflict in question was a war). That's why (23) seems odd.<sup>49</sup> In short, the closure-denier's explanation of abominable conjunctions like (22) gains independent support from the fact that it also explains awkward conjunctions like (23). In both cases, the first conjunct contradicts an implicature generated by the second conjunct, given the common ground.<sup>50</sup>

One might object to the closure-denier's explanation by reminding us that conversational implicature is cancelable. This means their explanation must allow for the possibility of contexts where abominable conjunctions are felicitous. However, closure-accepters must also allow for such contexts. Recall that closure-accepters explain the infelicity of (22) by citing common ground information like (BG). But there are possible contexts where the common ground of (22) does not include (BG). And these are contexts where the awkwardness of (22) could disappear.<sup>51</sup> So, both explanations allow for the possibility of contexts where abominable conjunctions are felicitous.

It might be argued that closure-accepters have a better explanation, because they can predict the oddity of abominable conjunctions even when the implicature is canceled, such as in

---

<sup>49</sup> It cannot be claimed that (23) is awkward because the first conjunct entails the second conjunct (making the second one redundant). There is no such entailment—it's possible to win without trying (e.g. the opponent may forfeit).

<sup>50</sup> A further similarity is that the oddity of (22) and (23) persists even when their conjuncts are transposed.

<sup>51</sup> To be clear, it's possible that the awkwardness of (22) persists even when (BG) is excluded from the common ground, since there could be *something else* that causes the awkwardness of (22) whenever the common ground excludes (BG). But the same could also be said for the closure-denier's explanation. It's possible that all mechanisms that can cancel the implicature also create infelicity (e.g. see the discussion on redundancy in the next paragraph).

(24) # Moore knows he has hands, but doesn't know he's not a handless-BIV, although he's not ignorant of his not being a handless-BIV.

(24) seems awkward. But, according to this objection, it should sound fine, because the third clause cancels the implicature from 'not know' to 'ignorant'. To see why this objection fails, notice that (23) also remains infelicitous when we try to cancel its implicature, as in

(25) # The U.S. has won all wars, and they tried to defeat North Vietnam; in fact, they succeeded against North Vietnam.

Why are (24) and (25) awkward despite cancellation attempts?—because the third clause in each sentence is entailed by the first clause, given the common ground. Each attempt at cancellation creates redundancy, and redundancy is an additional source of linguistic oddity (Horn 1991). So, the closure-denier's explanation does not predict that the awkwardness of (22) should disappear in (24).

So far, I've argued that closure-deniers can explain the oddity of abominable conjunctions just as well as closure-accepters. If I'm right then these conjunctions do not support [Closure].

But John Hawthorne has objected that explanations of abominable conjunctions that rely on conversational implicature fail to explain why we *believe* the propositions expressed by the negations of abominable conjunctions (2005, 30-1). However, the negations of abominable conjunctions are equivalent to particular instances of closure. For example, the negation of (22) is equivalent to

(26) If Moore knows he has hands, then he knows he's not a handless-BIV.

And, as argued in section IV.1, the Closure Hypothesis can explain why we intuit the proposition expressed by (26), which can in turn explain why we believe the equivalent proposition expressed by the negation of (22). Moreover, the Closure Hypothesis makes no appeal to conversational



implicature.<sup>52</sup> So, conversational implicature is not invoked to explain why we *believe* the negations of abominable conjunctions. It's only invoked to explain the *linguistic oddity* of such conjunctions.

## V. Conclusion.

In examining the New and Old Ignorance Puzzles, four related conclusions have been advanced:

- (i) Contrary to intuition, ignorance is not merely a failure to know.
- (ii) 'Not know' conversationally implicates 'ignorant'.
- (iii) Because of (i), there are weaker no-ignorance principles that can be used to explain away the appeal of any knowledge-closure principle.
- (iv) Because of (ii) and (iii), we can explain why abominable conjunctions are linguistically awkward without assuming any knowledge-closure principles.

I have not argued that the skeptic's preferred closure principle is false; nor have I argued that its corresponding no-ignorance principle is true. But let's suppose we have reason to think both are the case. If so, we now have a new way of diagnosing why the skeptic's reasoning is appealing albeit misleading. The skeptic thinks we can always gain knowledge by inferring it from what we already know. But he has mistaken our *loss of ignorance* for a *gain in knowledge*. We might be guaranteed to lose ignorance by inferring from what we know. But that's an importantly different matter, with less troubling consequences.<sup>53</sup>

## Bibliography

Antioch College. (1992) 'The Antioch College Sexual Offense Policy,' *Date Rape: Feminism Philosophy, and the Law*, ed. L. Francis. University Park, PA The Pennsylvania University Press.

Blome-Tillmann, Michael. (2008) 'Conversational Implicature and the Cancellability Test,' *Analysis* 68:156-60.

---

<sup>52</sup> I have invoked conversational implicature to explain why we mistakenly believe or intuit <Sufficiency>, and *that mistake* is used to explain why we mistakenly think <NI> and <Closure> are equivalent. But, strictly speaking, the Closure Hypothesis could be true even if my way of explaining away <Sufficiency> is false. So, an appeal to conversational implicature isn't assumed by the Closure Hypothesis. Still, Hawthorne's objection could be refocused on the question of how conversational implicatures can help explain our belief or intuition that <Sufficiency> is true. In response, see my explanation in footnote 31 where I explain how such implicatures can play this role.

<sup>53</sup> I would like to thank Matti Eklund, Nico Silins, Adam Pelser, Pierre LeMorvan, Rik Peels, Adam Bendorf and three anonymous reviewers for their very helpful comments.

- Blutner, Reinhard. (2004) 'Pragmatics and the Lexicon,' *The Handbook of Pragmatics*, ed. L. Horn and G. Ward, 488-514. Oxford: Blackwell Publishing.
- BonJour, Laurence. (1987) 'Nozick, Externalism, and Skepticism,' *The Possibility of Knowledge*, ed. L. Luper-Foy. Lanham, MD: Rowman and Littlefield.
- Clark, Herbert. (1996) *Using Language*. Cambridge, UK: Cambridge University Press.
- Cohen, Stewart. (1999) 'Contextualism, Skepticism, and the Structure of Reasons,' *Nous Supplement: Philosophical Perspectives* 33, 57-89.
- Cruse, Alan. (1986) *Lexical Semantics*. Cambridge, UK: Cambridge University Press.
- David, Marian and Warfield, Ted. (2008) 'Knowledge-Closure and Skepticism,' *Epistemology: New Essays*, ed. Q. Smith, 137-87. Oxford: Oxford University Press.
- DeRose, Keith. (1995) 'Solving the Skeptical Puzzle,' *Philosophical Review*, 104(7), 1-52.
- DeRose, Keith. (2009) *The Case for Contextualism: Knowledge, Skepticism, and Context*, Vol. 1. Oxford: Oxford University Press.
- Dretske, Fred. (2005) 'The Case against Closure,' *Contemporary Debates in Epistemology*, ed. M. Steup and E. Sosa, 13-26. Oxford: Blackwell Publishing.
- Driver, Julia. (1989) 'Virtues of Ignorance,' *The Journal of Philosophy* 86, 373-84.
- Engel, Mylan. (2004) 'What's Wrong with Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox,' *Erkenntnis* 61, 203-31.
- Feldman, Richard. (1999) 'Contextualism and Skepticism,' *Philosophical Perspectives*, ed. J. Tomberlin. Atascadero, CA: Ridgeview.
- Flanagan, Owen. (1990) 'Virtue and Ignorance,' *The Journal of Philosophy* 87, 420-28.
- Geis, Michael and Zwicky, Arnold. (1971) 'On Invited Inferences,' *Linguistic Inquiry* 2, 561-566.
- Goldman, Alvin and Olsson, Erik. (2009) 'Reliabilism and the Value of Knowledge,' *Epistemic Value*, ed. A. Haddock, A. Millar and D. Pritchard, 19-41. Oxford: Oxford University Press.
- Greenberg, Joseph. (1966) *Language Universals*. The Hague: Mouton and Co.
- Grice, Herbert Paul. (1989) *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Gulley, Norman. (1968) *The Philosophy of Socrates*. New York: Macmillan.

- Hawthorne, John. (2005) 'The Case for Closure,' *Contemporary Debates in Epistemology*, ed. M. Steup and E. Sosa, 26-43. Oxford: Blackwell Publishing.
- Hazlett, Allan. (2007) 'Grice's Razor,' *Metaphilosophy* 38, 669-90.
- Heller, Mark. (1999) 'Relevant Alternatives and Closure,' *Australasian Journal of Philosophy* 77, 196-208.
- Horn, Laurence. (1989) *A Natural History of Negation*. Chicago: University of Chicago Press.
- Horn, Laurence. (1991) 'Given as New: When Redundant Affirmation Isn't,' *Journal of Pragmatics* 15, 313-36.
- Horn, Laurence. (2000) 'From *if* to *iff*: Conditional Perfection as Pragmatic Strengthening,' *Journal of Pragmatics* 32, 289-326.
- Huang, Yan. (2007) *Pragmatics*. Oxford: Oxford University Press.
- Kittay, Eva Feder. (1997) 'AH! My Foolish Heart: A Reply to Alan Soble's "Antioch's 'Sexual Offense Policy': A Philosophical Exploration",' *Journal of Social Philosophy* 28, 153-59.
- Kvanvig, Jonathan. (2006) 'Closure Principles,' *Philosophy Compass* 1: 256-267.
- LeMorvan, Pierre. (2011) 'On Ignorance: A Reply to Peels,' *Philosophia* 39: 335-344.
- LeMorvan, Pierre. (2012) 'On Ignorance: A Vindication of the Standard View,' *Philosophia* 40, 379-93.
- Levinson, Stephen. (1983) *Pragmatics*. Cambridge, U.K: Cambridge University Press
- Levinson, Stephen. (2000) *Presumptive Meanings*. Cambridge, MA: The MIT Press.
- Lewandowska-Tomaszczyk, Barbara. (2007) 'Polysemy, Prototypes, and Radial Categories,' *Oxford Handbook of Cognitive Linguistics*, eds., D. Geeraerts and H. Cuyckens, Oxford: Oxford University Press.
- Lewis, David. (1996) 'Elusive Knowledge,' *Australasian Journal of Philosophy* 74: 549-67.
- Ludlow, Peter. (2005) 'Contextualism and the New Linguistic Turn in Epistemology,' *Contextualism in Philosophy*, eds., G. Preyer and G. Peter, Oxford: Oxford University Press.
- Nozick, Robert. (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Peels, Rik. (2010) 'What Is Ignorance?' *Philosophia* 38: 57-67.

- Peels, Rik. (2011) 'Ignorance is Lack of True Belief: A Rejoinder to LeMorvan,' *Philosophia* 39: 345-55.
- Quine, Willard Van Orman. (1960) *Word and Object*. Cambridge, MA: The MIT Press.
- Russell, Bertrand. (1948) *Human Knowledge: Its Scope and Its Limits*. New York: Simon and Schuster.
- Schaffer, Jonathan. (2004) *Philosophical Studies* 119: 73–103.
- Schiffer, Stephen. (1996) "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society* 96, 317-333.
- Sosa, Ernest. (1999) 'How to Defeat Opposition to Moore,' *Nous* 33, *Supplement: Philosophical Perspectives*, 141-53.
- Sosa, Ernest. (2000) 'Skepticism and Contextualism,' *Philosophical Issues* 10, 1-18.
- Stanley, Jason. (2004) 'On the Linguistic Basis for Contextualism,' *Philosophical Studies* 119, 119-146.
- Stanley, Jason. (2005) *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Steup, Matthias. (2005) 'Contextualism and Conceptual Disambiguation,' *Acta Analytica* 20, 3-15.
- Unger, Peter. (1975) *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.
- Vlastos, Gregory. (1985) 'Socrates' Disavowal of Knowledge,' *The Philosophical Quarterly* 35, 1-31.
- Weiner, Matthew. (2009) "The (Mostly Harmless) Inconsistency of Knowledge Ascriptions," *Philosophers' Imprint* 9, 1-25.
- Williams, Bernard. (1995/2006) 'Philosophy and the Understanding of Ignorance,' *Diogene* 169, reprinted in *Philosophy as a Humanistic Discipline*, ed., A.W. Moore, Princeton: Princeton University Press, 169-79.
- Williamson, Timothy. (1997) 'Vagueness and Ignorance,' *Vagueness: A Reader*, eds. R. Keefe and P. Smith. Cambridge, MA: The MIT Press.
- Zimmerman, Michael. (1988) *An Essay on Moral Responsibility*. Totowa, NJ: Rowan and Littlefield.
- Zwicky, Arnold and Sadock, Jerrold. (1975) 'Ambiguity Tests and How to Fail Them,' *Syntax and Semantics* 4, ed. Kimball. New York: Academic Press, 1-36.