

Draft. Forthcoming in *Philosophical Studies*.

Strict Moderate Invariantism and Knowledge-Denials

Gregory Stoutenburg

Abstract:

Strict moderate invariantism is the ho-hum, ‘obvious’ view about knowledge attributions.¹ It says knowledge attributions are often true and that only traditional epistemic factors like belief, truth, and justification make them true. As commonsensical as strict moderate invariantism is, it is equally natural to withdraw a knowledge attribution when error possibilities are made salient. If strict moderate invariantism is true, these knowledge-denials are often false because the subject does in fact know the proposition. I argue that strict moderate invariantism needs an explanation of this phenomenon, but it does not have one. That is significant, for if strict moderate invariantism does not square with ordinary intuition, then it cannot rely on ordinary intuition for support.

Section 1 introduces the concept of *epistemic relevance blindness*, which says ordinary subjects are generally insensitive to whether or not error possibilities are relevant to knowledge attributions. Section 2 focuses on Patrick Rysiew’s influential strict moderate invariantist pragmatic explanation of knowledge-denials and argues that such pragmatic explanations of knowledge-denials depend on attributors being epistemic relevance blind. Section 3 targets psychological explanations of epistemic relevance blindness offered separately by Jennifer Nagel and Mikkel Gerken. I argue that strict moderate invariantists lack a plausible explanation of epistemic relevance blindness.

¹ Rysiew (2007, p. 629) offers a defense of what he calls the “ho-hum” view.

Keywords: semantics of knowledge attributions, invariantism, knowledge-denials

1. Introducing the Problem for Strict Moderate Invariantism

The practice of invoking the concept of knowledge for some purpose or other is as ordinary as any one might think of. Central to the interests of epistemologists and philosophers of language is how to give an account of the truth-values of our knowledge attributions—sentences of the form “S knows that p” —that fits with our intuitions about which knowledge attributions are natural to make in particular contexts. We are well aware of the fact that it is natural to make a positive knowledge attribution in some contexts, but not in others. Stewart Cohen’s (1999) now-canonical Airport Case illustrates this.

The Airport Case: Two passengers, Mary and John, are at the airport wondering whether the flight they plan to board has a layover in Chicago. They observe another passenger, Smith, look at his own itinerary and claim to himself that he knows the flight stops in Chicago. Mary and John have an important meeting in Chicago. Mary says to John, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John conclude that Smith does not know there is a layover in Chicago and decide to answer their own doubts by checking with a gate attendant for more information about the flight.²

According to Cohen, Smith is applying a lower standard for use of ‘knowledge’ than Mary and John because the standards applicable in Smith’s context are lower than in Mary and John’s context. The intuition that both Smith’s self-attribution of ‘knowledge’ that the flight stops in Chicago is true and that Mary’s denial that Smith ‘knows’ the flight stops in Chicago is also true

² It has been awhile since Cohen offered this example, but the pattern is general. Nowadays most people check internet-connected smartphones or airport television monitors for departure information. An updated airport case might have Smith rely on a terminal monitor for his knowledge claim. Mary’s skeptical worry would then be a little less intuitively forceful given that one would be reasonable to assume that an airport monitor is more likely to be accurate than an itinerary printed who-knows-when. Still, as anyone who has spent enough time at airports knows, those monitors do not always give accurate information.

is best explained, says Cohen, by different propositions being expressed by these knowledge attributions. This version of contextualism would have us understand Smith's claim, "I know that the flight stops in Chicago," as expressing the proposition *Smith knows_{low} that the flight stops in Chicago*, and Mary's claim, "Smith doesn't know that the flight stops in Chicago," as expressing the proposition *It is not the case that Smith knows_{high} that the flight stops in Chicago*. The operative standard of 'knowledge' in the context in which knowledge attributions are made partly determines the semantic content of the utterance.

Invariantism says "S knows that p" expresses the same proposition in every context of utterance (when S and p are fixed). *Strict* invariantism says the only factors that determine the truth-value of a knowledge attribution when the subject of the attribution believes a true proposition are factors that make the proposition likely to be true (Fantl and McGrath 2009, Gerken 2013). *Moderate* invariantism (MI) says sentences of the form "S knows that p" are often true (Hawthorne 2004). What makes MI 'moderate' is not its epistemic standards, but that it allows for lots of true knowledge attributions. One might, for instance, defend both MI and the view that when S knows that p, S knows that S knows that p, provided that one also holds that it is relatively easy to get knowledge of knowledge (and of knowledge of knowledge of knowledge...). Putting all of this together,

Strict Moderate Invariantism says knowledge attributions are often true, and they are made true by traditional epistemic factors like having a true, justified (or warranted or sensitive or whatever) belief that is not defectively accidental.

Strict moderate invariantism is the target of this paper. For convenience, I will often use the briefer 'moderate invariantism' and 'MI'.

Let us return to the airport case. Contextualism has an easy time accounting for the conversational appropriateness of Smith's 'knowledge'-affirmation and Mary's 'knowledge'-

denial, for according to contextualism, both sentences are true. According to MI, though, Smith, Mary, and John all know that the flight stops in Chicago.³ Smith's self-attribution of knowledge is appropriate because it is true. But Mary's knowledge-denial is false. Yet intuitively it seems that Mary's knowledge-denial was appropriate. Generalizing from the airport case, if MI is true, we are systematically mistaken when we deny knowledge under irrelevant error possibilities. Moderate invariantists must explain this. Concerning the airport case, MI must answer the question, "If Mary, Smith, and John all know that the flight stops in Chicago, why is Mary's knowledge-denial appropriate in this context?" In this paper I will argue that MI does not have a satisfactory answer to that question.

If there is no satisfactory explanation of the systematic mistake of taking merely salient error possibilities to be relevant to knowledge, that is a serious problem for MI. Surely one of the primary goals of a satisfactory account of knowledge attributions is to provide a semantics that reflects the standard for knowledge (thus 'knowledge') that speakers implicitly accept. Much of MI's initial plausibility is due to the fact that we so often attribute knowledge to ourselves and others, and we do not pre-reflectively think we are speaking or thinking falsely when we do so, nor that something different is said by "S knows that p" in one context than in another. Our epistemic standards somehow come out in our talk. But, according to MI, when speakers make knowledge-denials in ordinary contexts, these denials are false because subjects have irrelevant error possibilities in mind when they deny knowledge. So, for MI to retain its

³ A moderate invariantist could deny that in this situation, Mary, John, and Smith know that the flight stops in Chicago. But the worry generalizes: if one cannot get knowledge of where a flight is going by checking ordinary sources, then it is likely that knowledge attributions are very often false. That is incompatible with MI. So, even if one wants to reject the airport case, the problem for MI that I am pointing to by means of the airport case should be clear enough.

status as the ‘obvious’ view about knowledge attributions, moderate invariantists need to come up with an explanation of why we are systematically mistaken when we deny knowledge in these common cases. Moderate invariantists need to explain this systematic mistake without implying that the mistake accurately reflects speakers’ epistemic standards.

There are two basic, compatible strategies one might use to provide the needed explanation. The first is to provide a *pragmatic explanation*: to argue from conversational norms there is no mistake, just a certain kind of imprecision in ordinary discourse and thought. This moderate invariantist says that once we get clear about how knowledge attributions are *used* in ordinary contexts, we will have explained why knowledge-denials are so often used to express something true and/or useful in a particular context, even though they are false. The other strategy is to provide a *psychological explanation*: to argue that speakers are making a mistake that is partly responsible for why speakers take the knowledge-denials in question to be true. Moderate invariantists using this strategy say the psychological factors responsible for speakers’ confusion about the truth-values of knowledge-denials are general and apply to knowledge-denials as a special case.

If part of the MI psychological explanation of why speakers make false knowledge-denials is due to their inability to distinguish epistemically relevant error-possibilities from epistemically irrelevant error-possibilities, then speakers are *epistemic relevance blind*. I will argue that the MI psychological explanation does require this, and that moderate invariantists do not have a satisfactory explanation of it. The concept of epistemic relevance blindness appears throughout, so I will make it explicit.

Epistemic relevance blindness is the condition of being insensitive to whether or not an error-possibility (or class of error possibilities) is relevant to knowing some proposition (or class of propositions).⁴

It is plausible that epistemic relevance blindness, if there is any such thing, may be limited to particular classes of propositions, or may come in degrees. Maybe we are all epistemic relevance blind when it comes to error possibilities that support a kind of general skepticism, like brain-in-a-vat hypotheses. Or perhaps epistemic relevance blindness is more limited, and affects us only when we consider practically significant error possibilities. Some people may be less sensitive to which possibilities are relevant than others. The core idea is only that speakers are systematically mistaken to some degree about the relevance of (at least some) error possibilities to knowledge attributions.

If Mary is epistemic relevance blind, that partly explains why she says Smith doesn't know that the flight stops in Chicago. Epistemic relevance blindness partly explains that because it says Mary is unable to tell whether the possibility that there is a misprint on the itinerary is relevant to knowing that the flight stops in Chicago, or whether that possibility is merely an irrelevant possibility that she is thinking about. Of course, if the MI psychological explanation implies speaker epistemic relevance blindness, it needs to explain that, too.

In the next section, I evaluate what may be fairly called *the* MI pragmatic explanation of knowledge-denials. The explanation is due to Patrick Rysiew (2001, 2007), but has been cited approvingly by other moderate invariantists. I argue that the pragmatic explanation only works if speakers are epistemic relevance blind. So, MI needs a psychological explanation of

⁴ This basic notion plays a role in the availability heuristic explanation offered by (Hawthorne 2004) and (Williamson 2005b), as well as the raised standards-for-belief explanation offered by (Bach 2005, 2010), although they neither identify the condition nor attempt to explain it.

knowledge-denials regardless of the prospects of a pragmatic explanation. I then turn to two such MI psychological explanations—those offered by Jennifer Nagel (2008, 2010, 2012) and Mikkel Gerken (2012, 2013, forthcoming)—and raise problems for them.

One final note before moving forward. I have for convenience focused on the problem MI faces in explaining knowledge-denials and I will generally continue to do so. But the phenomenon that challenges MI is broader than straightforward knowledge-denial. When an error-possibility is mentioned after a knowledge attribution is made, speakers may also retract the previous attribution without denying it (“I guess I shouldn’t have said ‘I know’”), revise the attribution into a conditional knowledge attribution (“I should have said that I know that *if* there is no misprint, *then* the flight stops in Chicago”), or rebut the possibility (“Airlines don’t write just *anything* on those itineraries. We can trust it.”) Each of these responses assumes that the speaker takes the error-possibility in question to be relevant. So, each exhibits a mistake if the possibility in question is irrelevant.⁵

2. The Pragmatic Explanation

Patrick Rysiew claims an utterance of “S knows that p” pragmatically implicates “‘S’s epistemic position with respect to p is good enough...’ where the ellipsis is completed according to context” (2001).⁶ Context provides the threshold for how good S’s epistemic position must be in order to appropriately assert that S knows. As an invariantist, Rysiew argues that context

⁵ While I can only point for support to conversations with philosophers on this topic, a natural first move when confronted with the objection that MI needs to explain knowledge-denials is to claim that when given more information—say, about the reliability of itineraries—it is natural to go on to claim knowledge. That might be true (I am unsure), but it does not help MI. The problem is that if MI is true the possibility does not need to be eliminated in order to have knowledge. Such possibilities should simply be ignored. They are not ignored. MI needs to explain that.

⁶ Cf. Rysiew (2007) for similar ideas.

provides only that threshold and has no bearing on what is semantically expressed by a particular use of “S knows that p.” Rysiew endorses a relevant alternatives analysis of knowledge according to which S knows that p only if there are no *relevant* possibilities inconsistent with p which S cannot eliminate. For Rysiew, a possibility is relevant if and only if normal humans think it is likely. *Salient* possibilities are those “which the parties in a given situation ‘have in mind.’” (2001). A possibility may be salient without being relevant. This is especially clear if a possibility has been mentioned, and so is salient, but is one normal humans would think is unlikely, and so is irrelevant. On Rysiew’s moderate invariantism, the possibilities that are salient at a given moment in a conversation determine what proposition is pragmatically implicated by a knowledge attribution. Those possibilities do not, however, have any bearing on which possibilities are relevant.

To illustrate:

Speaker: “S knows that p.”

—(literally says)—→ “S truly, justifiably believes that p and can eliminate all not-p possibilities normal people would think are likely.”

—(pragmatically conveys)—→ “S can eliminate all salient not-p possibilities.”

Applying this to Cohen’s airport case,⁷ after the possibility of an itinerary misprint has been raised, Mary denies that Smith knows the flight stops in Chicago because at that point in the conversation the proposition expressed by an utterance of “Smith knows the flight stops in Chicago” would pragmatically implicate the false proposition that the salient possibility of an

⁷ Rysiew only applies the view to DeRose’s bank cases, but this is how the story would go for the airport case.

itinerary misprint can be eliminated.⁸ Rysiew supports this claim by appeal to Grice's maxim of Relation ("Be relevant!") (Grice 1975). That maxim requires conversational participants to contribute to the conversation in a way relevant to the topic at hand. Because the misprint possibility had just been raised, Rysiew holds that if Mary were to say "Smith knows that the flight stops in Chicago," that utterance would be taken by her audience to convey that Smith can eliminate the just-raised misprint possibility. Mary believes Smith cannot eliminate that possibility. And because she is prone to making the common confusion between what is semantically expressed by an utterance of "Smith knows that the flight stops in Chicago" and what would be pragmatically conveyed by that utterance, she asserts (falsely) "Smith doesn't know that the flight stops in Chicago" in order to avoid conveying (falsely) "Smith can eliminate the misprint possibility."

This is what Rysiew would have Mary's utterance expressing after the misprint possibility is made salient:

Mary: "Smith knows that the flight stops in Chicago."

—(literally says)→ "Smith truly, justifiably⁹ believes the flight stops in Chicago and can eliminate all not-(flight stops in Chicago) possibilities normal people would think are likely."

—(pragmatically conveys)→ "Smith can eliminate the misprint possibility."

⁸ Specifically, the utterance would, on Rysiew's view, pragmatically impart that the misprint possibility can be eliminated *by Smith*. In the airport case, no one can eliminate the misprint possibility at the time the conversations of interest are taking place.

⁹ It is clear in Rysiew (2007) that he accepts a JTB analysis of knowledge.

Now for the problem. If Mary did not regard the misprint possibility as relevant to knowing that the flight stops in Chicago, then she would not say “Smith doesn’t know...” The pragmatic story presupposes that speakers are epistemic relevance blind but does nothing to explain their epistemic relevance blindness. Mary’s epistemic relevance blindness is a necessary assumption of the appropriateness of her knowledge denial. If she were aware that the misprint possibility is irrelevant, her knowledge denial would be inappropriate, as I illustrate shortly. Rysiew does not address why speakers would think merely salient possibilities are relevant to knowledge when they are not. Nor do other moderate invariantists who endorse this style of pragmatic explanation, including Black (2005, 2008), Bach (2005, 2010),¹⁰ Williamson (2005a), and Brown (2006). Another story about someone saying something false in order to helpfully convey something true will illustrate the essential role played by epistemic relevance blindness in the MI pragmatic explanation of knowledge-denials.

One of the examples Grice (1975) used to illustrate the concept of conversational implicature concerns a driver who has just run out of fuel asking a pedestrian, “Is there a fuel station nearby?” Recognizing that the driver is asking the question for the purpose of finding fuel at that time, and knowing that there is a nearby station but that it is closed at that time, the pedestrian answers, “No.” In the example the appropriateness of the false “No” answer depends in part upon the pedestrian recognizing the driver’s intention to find fuel at that time. The

¹⁰ Bach’s explanation includes a psychological element, namely that attributors make false knowledge-denials under salient possibility of error because they mistakenly raise their standards for confidently believing the target proposition. Bach: “[W]hat is essential here is not the attributor’s lack of settled belief but her raised threshold for (confidently) believing. Before believing the proposition in question, at least with the confidence and freedom from doubt necessary for knowing, the attributor demands more evidence than knowledge requires” (2010, p. 114). For Bach’s explanation to work, attributors must be epistemic relevance blind because they fail to see that they do not in fact need more evidence in order to know.

pedestrian then knowingly says something false in order to convey something true, specifically that the driver will not be able to get fuel by going in the intended direction.

A modified version of that example illustrates the kind of mistake implied by the MI pragmatic story.

Giraffe Hunt

A is walking around an unfamiliar zoo looking for the tallest animal. B has been to this zoo several times and knows where nearly all of the exhibits are, which exhibits are open, and so on. While walking by, B overhears A's children asking to see the tallest animal in the zoo. A, knowing that if there is a giraffe at the zoo, it will be the tallest animal, turns to passerby B and asks, (pointing) "Excuse me, is there a giraffe down there?" There is a giraffe down there, and B knows it. B wants to be helpful. But B, falsely believing weight to be a determiner of height, wants A to find what B mistakenly believes is the tallest animal in the zoo: the rhinoceros. B, as a normal, cooperative speaker, recognizes that in this context, an utterance of "Yes, the giraffe is that way" would be taken by A to communicate that the tallest animal in the zoo is that way. B believes that is false. So, B says, "No": a claim B believes is false, but appropriate.

That story should strike you as a little bizarre. But it is exactly as bizarre as the MI take on the airport case: no more, no less. Here, B is sincerely trying to be helpful to A by saying something B believes is false. B is saying what is relevant both to the last thing said in the conversation and its topic. B's claim conforms to many of the Gricean maxims MI relies on to explain knowledge-denials in contexts where error-possibilities are salient.¹¹ The reason B denies there is a giraffe in the indicated direction is because he is badly mistaken about the conditions under which 'tall' is correctly predicated. But because B is also a cooperative speaker trying to lead A to the tallest animal, B says something false with the intention of thereby conveying something

¹¹ B's claim does not conform to the maxim of Quality—"Do not say what you believe to be false" (Grice 1975). But in this case, B chose to break one maxim in order to respect others, just as MI has speakers violate Quality by falsely denying knowledge in order to avoid falsely implicating that a certain possibility can be eliminated.

true. The appropriateness of B's claim depends on B's ignorance of what is and is not relevant to height. Change that fact, and B's claim is no longer appropriate.

The MI story is that speakers are aware that were they to say that they know some proposition, their utterance would convey that some possibility could be eliminated which in fact cannot be eliminated. Then, confusing what would be conveyed versus said, the speaker falsely says, "I don't know" or "Smith doesn't know" in order to avoid conveying the false proposition that some salient possibility of error can be eliminated. But that story only works if the speaker is also just plain wrong about what is relevant to possession of knowledge, like the story about B's denial that there are giraffes down there only works if B is just plain wrong about what is relevant to being tall. What makes Mary's false knowledge-denial appropriate is that *because* she is unaware that the misprint possibility is irrelevant, it is (allegedly) appropriate to assert falsely to avoid implicating falsely. But it is crucial to note that if she were aware that the misprint possibility is irrelevant, then the maxim of Relation ("be relevant!") would be violated by her knowledge-denial. That is because a knowledge-denial at that point in the conversation would be irrelevant to the topic of conversation, which is whether or not anyone knows the flight stops in Chicago. So, the MI pragmatic explanation of knowledge-denials assumes speaker epistemic relevance blindness.

If Mary were aware that the misprint possibility is irrelevant, she would not deny that Smith knows for the purpose of avoiding a false implicature, she would dismiss the relevance of the misprint possibility and get on the flight.¹² She would never be in the position of worrying

¹² Those who accept epistemic norms of action that require epistemic statuses 'stronger' than knowledge may deny that Mary is epistemically warranted in getting on the flight even though she knows the flight stops in Chicago. I think there is no epistemic status stronger than knowledge, but for discussion of such norms, see (Brown 2008) and (Gerken forthcoming, sect. 6.3).

about making the false implicature that the possibility can be eliminated, because her sensitivity to which possibilities are epistemically relevant and which are not would permit her to attribute knowledge even when the misprint possibility is salient.

Thus, MI cannot explain knowledge-denials using the pragmatic account alone because that account presupposes epistemic relevance blindness as a condition on the intelligibility of the account. If B were not badly mistaken about ‘tall’, the Gricean explanation of why B appropriately denies there is a giraffe down there would not work; if Mary were not epistemic relevance blind, the Gricean explanation of why she appropriately denies Smith knows would not work.

So, MI is committed to speaker epistemic relevance blindness. That is a psychological condition, so MI needs a plausible psychological explanation to account for it. I will now argue that the prospects of providing the needed explanation are dim.

3. Psychological Explanations

I will consider two psychological explanations that might account for epistemic relevance blindness, one from Jennifer Nagel (2010, 2012)¹³ and one from Mikkel Gerken (2012, 2013).¹⁴ These explanations are compatible with the sort of epistemic relevance blindness that is required to explain ordinary knowledge-denials. To anticipate my conclusions, however, Nagel’s account fails to explain epistemic relevance blindness because her explanation, like the pragmatic account,

¹³ In what follows I do not discuss Nagel (2008). In that paper, Nagel argued that the subject in a high-stakes case lacks knowledge. Her account in that paper does not address knowledge ascriptions under low-stakes error possibilities. Other than that, the arguments in her (2008) are very similar to the arguments in her (2010, 2012) that I do discuss.

¹⁴ Other psychological error theories include those of Hawthorne (2004) and Williamson (2005b). Nagel (2010) contains compelling empirical criticisms of their views. In addition, the Hawthorne-Williamson availability heuristic explanation, if it works, succeeds only in addressing knowledge-denials under high stakes. In this paper, I am also concerned with low-stakes knowledge ascriptions.

ultimately takes epistemic relevance blindness for granted. Gerken's account does not presuppose epistemic relevance blindness and is therefore *prima facie* better positioned to explain it. However, I will argue that his account does not explain knowledge-denials in a way that is compatible with MI.

The strategy Nagel and Gerken pursue is common in psychological studies on cognitive bias.¹⁵ First, the theorist identifies a supposed mistake. In our case, the mistake is regarding error possibilities that are merely salient as relevant to the truth-value of a knowledge ascription. Second, the theorist identifies some wider pattern of cognitive error or bias. Nagel and Gerken differ on which errors and biases are the culprit. Third, if successful, the theorist subsumes the particular error-type under the general pattern. Having done so, the theorist will have provided a psychological explanation of the data.

3.1 Jennifer Nagel's Epistemic Egocentrism Hypothesis

Jennifer Nagel (2010, 2012) argues that the knowledge-denials we naturally make under salient possibility of error is a result of *epistemic egocentrism*. Her epistemic egocentrism hypothesis says we attributors erroneously regard others as possessing background information we possess and we then hold them accountable to act on that information (2010, p. 301; 2012, p. 186).

When they fail to do so, we refrain from attributing to them knowledge that is properly theirs.

Our mistake, Nagel thinks, consists in judging the subject according to information we possess but which the subject does not. Insofar as the information we have that the subject lacks is

¹⁵ The strategy is transparent in Nagel and Gerken; their papers point the way to many other studies that apply the same general method.

irrelevant to whether or not the subject knows, we are mistaken in refusing to attribute knowledge to the subject.

Nagel applies the egocentrism hypothesis to intuitive judgments concerning a skeptical case:

John A./John B.

- (A) John A. is looking at a bright red table under normal lighting conditions and believes the table is red. Does he know the table is red?
- (B) John B. is looking at a bright red table under normal lighting conditions and believes the table is red. A white table under red lights would be indistinguishable to John B. Does he know the table is red?¹⁶

Subjects tend to answer the knowledge question affirmatively in (a) but not in (b). Nagel's hypothesis is that the (b) case prompts attributors to illegitimately evaluate John B. as though he shares our skeptical concerns and therefore we regard his belief as hastily and improperly formed, thus falling short of knowledge (2010, p. 303).

The egocentrism hypothesis may well be true. It may be that attributors do require that subjects of their knowledge attributions be able to rule out all error possibilities that are salient to attributors in order to be granted knowledge. However, the egocentrism hypothesis cannot explain epistemic relevance blindness. The egocentrism explanation of subjects' increased tendency to answer the knowledge question affirmatively in (a) but not in (b) is that *attributors* fail to realize that the unusual lighting error-possibility is *irrelevant* to whether or not John B. knows. That is, it assumes that attributors are epistemic relevance blind. The egocentrism hypothesis says attributors judge that John B. lacks knowledge because they regard as blameworthy his inability to rule out error-possibilities like unusual lighting. So, attributors think John B. *ought* to rule out various irrelevant possibilities of error in order to properly

¹⁶ The cases are shortened from Nagel's (2010, p. 287), which are adapted from Cohen's (2002) cases.

believe that the table is red. They think John B. *should not* believe that the table is red even though the error-possibility is not only irrelevant to whether or not he knows, it is also not salient to him.¹⁷ The only reason attributors would think John B. ought to rule out that error-possibility is that they think the error-possibility is relevant to John B.'s knowledge of the table's color. But to think that is just to conflate a salient possibility with an epistemically relevant possibility. What we wanted out of the egocentrism hypothesis was explanation of that conflation, and it does MI no good if the explanation of the conflation is that attributors really do regard all salient possibilities as epistemically relevant.

The egocentrism explanation assumes that attributors in the John A./John B. case are epistemic relevance blind because they fail to grant John B. knowledge of the table's color due to his inability to eliminate an error possibility that is irrelevant by MI's standard.¹⁸ Since the egocentrism hypothesis relies on the phenomenon that we hoped it would explain, it is clearly unsatisfactory as an explanation of epistemic relevance blindness.¹⁹

¹⁷ As a result, the standard for knowledge suggested by natural responses to the John A./John B. case is even more skeptical than the view that all salient error-possibilities are relevant: since attributors think John B. must be able to rule out the unusual lighting error-possibility even though is not salient to John B., that suggests the standard attributors hold requires that to know that p one must be able to rule out possibilities that are not salient but could be.

¹⁸ Nagel also uses the epistemic egocentrism hypothesis in her (2012) to attempt to explain why intuitions supporting a negative judgment as to whether a subject knows in skeptical case are misguided while intuitions supporting a negative judgment as to whether a subject knows in a Gettier case are accurate. I do not discuss that argument here for reasons of space. The replies to her (2010) which I have given here are relevant to her (2012) argument.

Some philosophers regard the ever-present epistemic possibility of *now* being subject to a Gettier-like accident as a skeptical concern (Reed 2009). If such an argument is well-motivated, then Nagel's attempt to separate intuitions supporting negative verdicts on Gettier cases from intuitions supporting negative verdicts in skeptical cases is bound to fail because the same intuitions are responsible for both verdicts.

¹⁹ Gerken (independently) draws the same conclusion about the egocentrism hypothesis (forthcoming, sect. 10.3).

3.2 Mikkel Gerken's Epistemic Focal Bias Hypothesis

Mikkel Gerken (2012; 2013; forthcoming) attempts to provide a general psychological account of our practice of knowledge attribution that is compatible with MI. The cornerstone of his account is his epistemic focal bias hypothesis. He describes it this way: “we tend to form epistemic judgments on the basis of the alternatives that are in focus—whether or not they are epistemically relevant alternatives” (2013, p. 48).²⁰ The basic idea is that our practice of knowledge attribution is driven by what possibilities we are focused on. If I am thinking about whether S knows that p and I am focusing on the possibility that q and I take q to undermine S's putative knowledge that p, I am as a result of epistemic focal bias likely to take q to be an epistemically relevant alternative to p and conclude that S does not know that p. When q is not epistemically relevant to S knowing that p, I will be mistaken, and my mistake will be a result of focal bias. The epistemic focal bias hypothesis has the potential to explain why speakers make false negative knowledge attributions under salient but epistemically irrelevant possibilities of error: knowledge attributions are driven by what is psychologically in focus, whether or not that which is in focus is epistemically relevant.

The epistemic focal bias hypothesis might explain the differing knowledge attributions in the case of John A./John B. this way: attributors in the John A. case have in focus that John A's belief is true and that it is formed on the basis of a truth-conducive belief-forming process, namely visual perception. Attributors in the John B. case, however, also have in focus an epistemically irrelevant alternative to John B.'s knowledge: the red lighting possibility. Because

²⁰ I focus here on Gerken (2013); Gerken (2012) has much the same arguments. Incidentally, Gerken says in (2013), “...[M]any pragmatic accounts are compatible with the epistemic focal bias account and may even require it” (57). He is surely right that a pragmatic MI account ultimately requires some psychological support of the sort that, if successful, the focal bias account would provide.

knowledge attributions are formed on the basis of the alternatives in focus, attributors in the John A. case get the knowledge attribution right while attributors in the John B. case get it wrong. The difference is due to alternatives to ‘John knows the table is red’ that attributors have in focus.

The problem for the egocentrism hypothesis was that it moved the problem of explaining epistemic relevance blindness up a level, assuming that attributors are epistemic relevance blind in order to explain the difference in attributions in the John A./John B. cases. But notice that epistemic focal bias does not assume epistemic relevance blindness. Epistemic focal bias just describes an important psychological mechanism responsible for making knowledge attributions. So, it is in principle possible that epistemic focal bias could help explain epistemic relevance blindness. It could be that attributors suffer from epistemic relevance blindness *because* they suffer from epistemic focal bias. If that were shown to be so, the epistemic focal bias hypothesis will have provided a successful psychological explanation of knowledge-denials by showing that focal bias explains epistemic relevance blindness. The outstanding question would then be why attributors think the red light possibility *undermines* knowledge rather than simply having some other effect on knowledge attribution when it is in focus.

The connection between focal bias and epistemic relevance blindness suggests a test for the epistemic focal bias explanation of knowledge-denials. To see if epistemic focal bias can explain epistemic relevance blindness in a way compatible with MI, construct a scenario where a knowledge attribution comes under salient possibility of error. Assume the epistemic focal bias hypothesis is true. Then get attributors to focus on the claim that the salient error possibility is unlikely. That last part is needed to ensure compatibility with MI because MI says knowledge is

compatible with unlikely error possibilities.²¹ If our knowledge attributions are driven by what is in focus, and what is in focus includes both a possibility of error and that the possibility of error is unlikely, then speakers should—if MI is true—be more willing to affirm knowledge in that case. If *q* being in *S*'s focus prompts *S* to judge that $\sim K_{sp}$, add to *S*'s focus that *q* is unlikely. If the focal bias hypothesis and MI are both true, making salient to *S* both *q* and that *q* is unlikely should persuade *S* that K_{sp} .

To illustrate, it should be possible on the focal bias account to tell attributors considering Nagel's case (b) that it is rare for white tables to be illuminated by red lights because stores selling tables have legal incentive to not disguise their products, that businesses hoping to stay afloat are unlikely to dupe consumers into buying things they do not actually want, etc. Once this is all in focus, the focal bias explanation predicts that speakers would be more likely to attribute knowledge because what is now in the speaker's focus plus MI supports a positive knowledge attribution. But speakers may still resist that the new information allows them to claim knowledge. Consider this exchange between our familiar John B.—who now shares our skeptical worry—and a friend, John C.:

Table Shopping

John B: "I'm not sure that's a red table. A white table under red lights would look just like that."

John C: "That's ridiculous! I mean, you're right, it *would*. But don't you think furniture stores probably go out of their way to make sure things look in the store the way they'll look in peoples' homes? If they didn't, people would *obviously* notice they didn't buy what they thought they did the second they got the table home. That would turn into a huge hassle for the store."

John B: "Well, *that one* might be white..."

John C: "Did you hear what I just said?"

²¹ Unlikely *relative to* something, that is. Both how unlikely a possibility is and how important it is that the possibility fail to obtain affect judgments of relevance.

John B: “Look, you’re right that making white tables look red would be a dumb move, but still, *that one* right there *could be* a white table, even if it *probably* isn’t. Don’t say you *know* that one *right there* is a red table. We’ll need to find a way to check.”²²

John B.’s resistance is natural. If, like John B., attributors continue to resist ascribing knowledge when they have in focus both some possibilities of error and the unlikelihood of those error-possibilities obtaining, epistemic focal bias could not explain that fact in a way consistent with MI. That is because it would show that speakers want the error possibility eliminated, rather than safely ignored. The success or failure of epistemic focal bias as an explanation of epistemic relevance blindness hinges on whether or not speakers focusing both on error-possibilities and their improbability continue to think that those error possibilities undermine knowledge.²³

There is significant precedent for resisting attributing knowledge in such a scenario. Only a committed fallibilist says one can know that a lottery ticket is a loser prior to the draw, even though everyone knows the dismal odds in advance. Banks do not regularly change their hours unannounced, but if you need your paycheck to hit your account before midnight tomorrow, you will not say that you know it will be open later when the line is shorter. I am physically healthy, I do not take a lot of risks, and I am not genetically predisposed to any

²² Note that this conversation does not require thinking of John B. as being unable to gather further evidence that would rule out the red-light possibility. That is why I have John B. suggesting they find a way to check the table’s color. But in the circumstances of the conversation, recognizing the improbability of *that table* being a white table under red lights does not allow John B. to conclude that they *know* the table is red.

²³ Gerken and Beebe (2016) supplement the epistemic focal bias hypothesis presented in Gerken’s other work with further accounts of what may be responsible for particular responses to contrast cases. One account they consider is an inadequate processing account. Gerken and Beebe hypothesize that inadequate processing may be responsible for attributors’ failing to take seriously (“adequately process”) alternatives that are out of focus (cf. p. 154). That may be. However, inadequate processing is not a viable candidate explanation of the reluctance to attribute knowledge in the Table Shopping scenario because there the red table alternative *is* in focus. I thank an anonymous referee for prompting me to discuss this.

significant health conditions down the road, but I am keeping my life insurance policy anyway—as people say, “You never know, today could be your last.” As far as explaining particular cases goes, the table shopping case is as easy as things get for MI. The intuitive pull to resist attributing knowledge only gets stronger if one modifies that story so that a red table is needed for some important practical purpose. Nor do intuitions change to favor MI if we shift to a third-person case, like this.

Third-Person Table Shopping

Mary and John are at the furniture store looking for a red table for their Chicago business meeting at Red, Inc. They overhear John B. ask another customer, John C., if the table in front of them both is red. John C. says it is: after all, it looks red. John B. says, “I’m not sure that’s a red table. A white table under red lights would look just like that.” John C. replies that the store would make sure that things look in the store the way they look outside the store, so it is very unlikely that the table is actually white rather than red. John B. is not convinced. Neither is Mary. Mary and John decide to start looking around for anything that might make the table appear red when it is not.

If all that is correct, the psychological explanation of knowledge-denials afforded by the focal bias hypothesis does not explain epistemic relevance blindness in a way compatible with MI. The focal bias hypothesis may be true in that our knowledge attributions are indeed driven by what is in focus at a time. But in the table shopping scenario and similar cases, the MI epistemic focal bias explanation of knowledge-denials predicts an increased tendency to make a positive knowledge attribution while it is natural to refuse to attribute knowledge. The best explanation of this is not that the epistemic focal bias explanation fails because the epistemic focal bias hypothesis is false, but that the explanation fails because MI is false.

My argument assumes that attributors should be able to appreciate the irrelevance of an error-possibility once its improbability comes into focus. Perhaps that is not so. Perhaps, for some reason, once an error-possibility is in focus, it becomes stuck there, undermining

knowledge attributions until it is no longer in focus. I doubt that is the case. If it were, then stressing the improbability of an error-possibility before mentioning the error-possibility should have the effect of undermining the relevance of an error-possibility. E.g. “What I am about to say is pretty unlikely, but if that were a white table under red lights, it would look red, just like that.” To my ear, the order in which an error-possibility and the improbability of the error-possibility become salient is irrelevant to the intuitive truth-value of a knowledge attribution made in a context where both are salient.

Those sympathetic to MI psychological explanation may point to ‘mindware gaps’ to undermine my response above (cf. Gerken forthcoming, sect. 10.4, Gerken and Beebe 2016). Mindware gaps are limitations in mental processing owing to a lack of particular concepts or exact sensitivity to particular distinctions, even if the subject with the mindware gap is somewhat sensitive to the distinction. The objector may claim that if there is a mindware gap that is partly responsible for the subject’s epistemic relevance blindness, then it should be expected that the mistake will persist across similar cases, like those presented above.

This objection may prove too much. MI is a folk epistemological thesis allegedly supported by reflection on ordinary thought and talk. If ordinary thought and talk is infected by mindware gaps, then ordinary thought and talk is an unreliable source of support for MI. That is not an argument against the existence of a mindware gap underlying epistemic relevance blindness. It is rather an argument that positing a mindware gap to protect one’s preferred epistemological theory from counterexamples puts one on thin methodological ice. That is because part of the basis for claiming that a mindware gap is responsible for a pattern of behavior is that there is a standard of correctness from which subjects deviate, where the deviation is due to a mindware gap. But in the case of MI, or any other folk epistemological

theory, the standard of correctness is derived from critical reflection on our ordinary thought and talk. The challenge, then, is to identify a standard of correctness from which ‘deviations’ are to be counted as effects of a mindware gap without taking one’s epistemological theory (e.g. MI) for granted.

The issue can be made clearer by contrasting the epistemologist’s conundrum with a case where the standard of correctness is easily identified. Everyone’s favorite example of cognitive bias is the Linda case from Tversky and Kahneman (1982). It is well-known that subjects regularly mess up by claiming that the probability of a conjunction of probabilistically independent claims (‘Linda is a feminist bank teller’) is higher than the probability of one of the conjuncts (‘Linda is a bank teller’). The argument for calling that behavior a “fallacy” crucially depends on there being a right answer to the question, “Which is more probable—that Linda is a bank teller, or that Linda is a feminist bank teller?” the correctness of which is independent of subjects’ answers to the question. The probability calculus tells us that the probability of any conjunction of probabilistically independent claims is never higher than the probability of one of the conjuncts alone. The probability calculus is the standard of correctness for probabilistic inferences, so when subjects conclude that it is more likely that Linda is a feminist bank teller than that Linda is a bank teller, we can point to the probability calculus to show that they are wrong. And when the mistake is common and systematic enough, we may be able to identify a cognitive bias that is responsible for the mistake.

But when we argue about which account of knowledge is true, there is no prior, independent standard of correctness that we can appeal to, and from which we can conclude that judgments about knowledge are correct if they fit that standard and incorrect if they do not. Tversky and Kahneman have the probability calculus, but epistemologists only have reflection

on our judgments. As a result, there is no clear basis for taking judgments that support MI to be reliable while the ubiquity of knowledge-denials under salient possibility of error is instead regarded as some kind of cognitive mistake. This may not only be a problem for MI. But the fact that others share one's problem does not imply that one does not have a problem.²⁴

4. Conclusion

If MI is true, speakers are epistemic relevance blind. That is, they are insensitive to the fact that salient error possibilities are often irrelevant to whether subjects of knowledge attributions know certain propositions. MI has speakers committing a battery of mistakes, some of which are rather surprising if they are competent with 'knows.' I have argued that MI has no good explanation of this fact. That seriously undermines the tenability of MI.²⁵

²⁴ I think this problem points to the superiority of pragmatic accounts to psychological ones, provided that those pragmatic accounts do not ultimately commit their defenders to positing systematic psychological errors. That goes for all the views in the debate about knowledge attributions. My view is that the assumption that speakers are generally sensitive to the semantics and pragmatics of language best supports the view that knowledge attributions are usually false.

²⁵ I thank Mikkel Gerken, Richard Fumerton, Ting Lau, Heather Stoutenburg, Landon Elkind, Emily Waddle, John Komdat, Ryan Cobb, and an anonymous referee for helpful comments on a draft of this paper. I also thank the audience of the Northwestern/Notre Dame Epistemology Conference for fruitful discussion.

References

- Bach, K. (2005). The Emperor's New 'Knows'. Contextualism in philosophy: Knowledge, meaning, and truth. G. Preyer and G. Peter. New York, Oxford University Press: 51-89.
- Bach, K. (2010). Knowledge In and Out of Context. Knowledge and Skepticism. J. K. Campbell, M. O.'Rourke and H. S. Silverstein. Cambridge, MIT Press: 105-136.
- Black, T. (2005). "Classic Invariantism, Relevance and Warranted Assertability Manœuvres." The Philosophical Quarterly **55**(219): 328-336.
- Black, T. (2008). "A Warranted-Assertability Defense of a Moorean Response to Skepticism." Acta Analytica **23**(3): 187-205.
- Brown, J. (2006). "Contextualism and Warranted Assertability Manœuvres." Philosophical Studies **130**(3): 407-435.
- Brown, J. (2008). "Knowledge and Practical Reason." Philosophy Compass **3**(6): 1135-1152.
- Cohen, S. (1999). "Contextualism, Skepticism, and the Structure of Reasons." Noûs **33**: 57-89.
- Cohen, S. (2002). "Basic Knowledge and the Problem of Easy Knowledge." Philosophy and Phenomenological Research **65**(2): 309-329.
- Fantl, J. and M. McGrath (2009). Knowledge in an Uncertain World. New York, Oxford University Press.
- Gerken, M. (2012). On the Cognitive Bases of Knowledge Ascriptions. Knowledge Ascriptions. J. Brown and M. Gerken. New York, Oxford University Press: 140-170.
- Gerken, M. (2013). "Epistemic Focal Bias." Australasian Journal of Philosophy **91**(1): 41-61.
- Gerken, M. (forthcoming). On Folk Epistemology: How we Think and Talk about Knowledge. New York, Oxford University Press.
- Gerken, M. and J. R. Beebe (2016). "Knowledge In and Out of Contrast." Noûs **50**(1): 133-164.
- Grice, H. P. (1975). Logic and Conversation. Syntax and Semantics, vol. 3: Speech Acts. P. Cole and J. Morgan. New York, Academic Press.
- Hawthorne, J. (2004). Knowledge and Lotteries. New York, Oxford University Press.
- Nagel, J. (2008). "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes." Australasian Journal of Philosophy **86**(2): 279-294.
- Nagel, J. (2010). "Knowledge Ascriptions and the Psychological Consequences of Thinking about Error." The Philosophical Quarterly **60**(239): 286-306.
- Nagel, J. (2012). Mindreading in Gettier Cases and Skeptical Pressure Cases. Knowledge Ascriptions. J. Brown and M. Gerken. New York, Oxford University Press.
- Reed, B. (2009). "A New Argument for Skepticism." Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition **142**(1): 91-104.
- Rysiew, P. (2001). "The Context-Sensitivity of Knowledge Attributions." Noûs **35**(4): 477-514.
- Rysiew, P. (2007). "Speaking of Knowing." Noûs **41**(4): 627-662.
- Tversky, A. and D. Kahneman (1982). Judgments of and by Representativeness. Judgment under Uncertainty: Heuristics and Biases. D. Kahneman, P. Slovic and A. Tversky. Cambridge, Cambridge University Press.
- Williamson, T. (2005a). "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." The Philosophical Quarterly **55**(219): 213-235.
- Williamson, T. (2005b). Knowledge, Context, and the Agent's Point of View. Contextualism in philosophy: Knowledge, meaning, and truth. G. Preyer and G. Peter. New York, Oxford University Press: 91-114.