Syracuse University
SURFACE at Syracuse University

Dissertations - ALL | SURFACE at Syracuse University

2-4-2015

Questions And Inquiry: Some Epistemological Applications

Jason Rourke
Syracuse University

Follow this and additional works at: https://surface.syr.edu/etd

 Part of the Epistemology Commons

Recommended Citation

Rourke, Jason, "Questions And Inquiry: Some Epistemological Applications" (2015). *Dissertations - ALL*. 1799.
https://surface.syr.edu/etd/1799

This Dissertation is brought to you for free and open access by the SURFACE at Syracuse University at SURFACE at Syracuse University. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE at Syracuse University. For more information, please contact surface@syr.edu.

Abstract: A recent and innovative research program in epistemology aims to connect the related phenomena of questions and inquiry with epistemological concerns. This dissertation project contributes to that program, taking as its inspiration the contrastive theory of knowledge developed by Jonathan Schaffer in a series of recent papers. The dissertation is comprised of three main parts. I begin by articulating a positive account of the evaluation of knowledge attributions, an account that aims to respect the basic insights of Schaffer's contrastivism while situating them in a modal framework that makes manifest the utility of questions for epistemological theorizing. Then I offer and discuss a counterexample that shows that any theory that fits the general structure that our respective accounts share cannot adequately account for certain contexts of knowledge attribution. I close by applying certain basic claims about the nature of inquiry to the problem of skepticism. I attempt to show that inquiry so understood is incompatible with a very strong version of skepticism, namely global skepticism about justification, with the result that non-skeptics can permissibly disregard it even if it is true.

Questions and Inquiry: Some Epistemological Applications

by

Jason Rourke

M.A., Syracuse University, 2010

Dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor in Philosophy in Philosophy

Syracuse University

December 2014

Copyright © Jason Rourke 2014

All Rights Reserved

Acknowledgments

The work that follows was developed with indispensable help and guidance from several people. Tom McKay, Andre Gallois, and Kris McDaniel offered insightful comments on early drafts of Chapters 2, 3, and 4. Chapter 3 benefitted a great deal from comments given by an anonymous referee for *Philosophical Studies*. I am also grateful to Paul DuBois, Kenneth Levin-Epstein, Larry Horn, Paul Hagstrom and Christopher Potts for helpful discussion. The work of Jonathan Schaffer is the inspiration for this project, and his ideas have been a constant spur to my thinking on the matters contained herein. I am grateful to him for being so generous with his time in helping me develop the ideas I argue for in what follows. Finally, my deepest thanks go to Mark Heller, who has contributed to this project in so many ways. Without his patience and unflagging support, the results would be much poorer indeed.

Table of Contents

Chapter 1: Background for the project

1 Skepticism

1.1 An overview

One way to distinguish the varieties of skepticism and situate them in relation to one another is along the dimensions of scope and strength, with the latter distinction cutting across the former. The scope of a skeptical view is determined by the range of propositions the view is thought to affect. Views with the widest scope are the various global skepticisms, which hold that for some epistemic property (such as knowledge or justification) agents in a certain class (such as beings with cognitive abilities like ours) fail to instantiate that property with respect to *any* proposition whatsoever. The strength of a particular brand of skepticism is determined by which epistemic property or properties the view targets. Among skeptical views, the view that is strongest and has the widest scope is global skepticism about justification- the thesis that for any subject S and any proposition p, S is not justified to any degree in believing that p.[1] For the global skeptic about justification, your current belief that you are reading a dissertation chapter enjoys no better epistemic credentials for you than, say, a spontaneously formed belief (were you to somehow form it now) that the highest temperature recorded in Syracuse on June 16th 2026 will be 81.6 degrees. If we assume that justification (or some cognate concept) for p is a necessary condition for knowing that p then global skepticism about justification entails but is not entailed by global skepticism about knowledge. The global skeptic about knowledge, unlike the global justification skeptic, need not claim that every belief is on an epistemic par with every

[1] Global skepticism about justification is less often defended in the literature than other versions of skepticism, but contemporary defenses can be found in Unger (1975), who presents it as a consequence of the connections he draws between knowledge and certainty, and Oakley (1976), who offers it as the outcome of the justification regress problem.

other. The global skeptic about knowledge can agree that your current belief that you are reading a dissertation chapter is more justified for you than the belief about the temperature in Syracuse on a day twelve years in the future would be. The global knowledge skeptic is committed to the weaker claim that there is no true belief for which agents like us can acquire justification sufficient unto knowledge- for any subject S and any proposition p, it is not the case that S knows that p.

Compared with global versions of skepticism, each of the local versions of skepticism is narrower in scope, challenging some restricted set of propositions. An external world skeptic, for example, will typically deny that we have any knowledge of the world around us, although she may concede that we can have knowledge of such things like our occurrent mental states. A skeptic about other minds is concerned with arguing that no agent knows any proposition describing the mental status of others, although she may happily concede that we may have knowledge of our own minds and much of the world beyond us.

No matter whether a skeptic endorses some version of global or local skepticism, she is interested in a defending a particular thesis concerning our relation to some epistemic property or properties. The skeptic is typically not interested in defending the claim that there are some propositions that every agent is in practice ignorant with respect to. The skeptic rather proposes the stronger thesis that for some class of propositions there is no way even in principle to come to know or to be justified in believing any member of that class. The skeptic may explain our failure by identifying the kind of evidence that is relevant to propositions of the specified class, and then showing that it is not possible for us to obtain enough of this sort of evidence to render our beliefs in propositions from this class known (or even justified).

Explanations of this kind are often made vivid by the skeptic's description of a world that is radically different from the world we take ourselves to inhabit, yet is utterly indistinguishable from it, even under the supposition that our evidence is optimal. Thus we are variously invited to suppose that each of us is merely a brain in a vat whose entire sensory experience is experience of a world that is wholly fictitious; that we are subject to a systematic and unrelenting delusion produced by some malicious occult power; that our fellows behave one and all just as they do now despite the fact that they are automatons bereft of consciousness; and so forth.

Such experiments, the skeptic thinks, force us to recognize that *no* amount of evidence that it is possible for us to obtain is sufficient to distinguish, say, the state of affairs in which you are in fact reading a dissertation chapter from a state of affairs in which sophisticated software is producing an utterly persuasive but delusory reading experience for your envatted brain.[2] Any sensory evidence that you might appeal to reject that latter possibility, no matter how vivid and clear, is of course open to the challenge that it is possible that such evidence is itself a product of the apparatus to which you are connected.

The skeptic's conclusion that you thus fail to know that you are reading a dissertation chapter right now is consistent with a number of theories concerning evidence and its connection to knowledge. A skeptic might assent to the conclusion against the background of broadly Cartesian approach to knowledge and evidence, according to which knowledge of a proposition requires that one's evidence for the proposition be such as to remove *all* grounds for doubting it. Put another way, evidence that is sufficient for S's knowledge that p must be evidence that allows S to be *certain* that p. Since the hypothesis that you are merely an envatted brain does

[2] The move from this claim to the claim that we cannot have knowledge of the external world is often underwritten, either explicitly or implicitly, by some thesis to the effect that one's knowing that p requires that one knows that every proposition that one knows to be incompatible with knowing that p is false. This thesis (and some variants on it) and the prospect of challenging it is nicely surveyed in Chapter 1 of Stroud (1984).

seem to constitute a genuine ground for doubting the proposition that you are reading a dissertation chapter right now, and since it is claimed that there is no way, even in principle, for you to obtain sensory evidence for this proposition that would also be sufficient for denying the hypothesis, the skeptic looks to be in a strong position to deny that you have knowledge of this proposition. Generalizing, the skeptic concludes that we fail to know any proposition that depends upon the evidence of the senses for its justification.

But a stringent Cartesian conception of evidence is not a prerequisite for endorsing skepticism. Skepticism, about both knowledge and justification, global and local, might instead be underwritten by appeal to a very plausible principle that is amenable to strong as well as weak conceptions of the strength of evidence required for knowledge or justified belief. That principle is the closure principle, and many contemporary arguments for skepticism exploit it, either explicitly or implicitly. In a slogan, the closure principle holds that knowledge is closed under known entailment. Although there is some dispute about how to best render this slogan more precise, a popular regimentation of the idea runs as follows: if S knows that p, and S knows that p entails q, then S knows that q.[3] The closure principle as stated is plausible enough to be treated as a near platitude by some epistemologists, and it is fair to say that nearly all believe that respecting the closure principle is an important desideratum for epistemological theories.[4]

The skeptic can appeal to this principle to ground a very seductive argument for the claim that we fail to have knowledge of even the kind of propositions that no reasonable person would ordinarily call into question, such as your putative knowledge that you are reading a dissertation

[3] As noted, a number of alternative formulations of the closure principle appear in the literature. Both Williamson (2002) and Hawthorne (2005) contain recent discussion of the issue.
[4] There are exceptions. Dretske (1970), Nozick (1981), and Heller (1999) reject closure. Schaffer (2007) rejects standard formulations of the principle and offers principles that capture the spirit of closure in a way that is compatible with his contrastivist view.

chapter right now. The argument is often presented with a nod to Moore's notorious anti-skeptical defense- the first premise of the argument states that if I know that I have hands, and I know that my having hands entails that I am not merely a handless envatted brain, then I know that I am not merely a handless envatted brain. As an instance of the closure principle, the premise is difficult to challenge. By appealing to the sort of evidential underdetermination lately discussed, the skeptic grounds the argument's second premise that I do *not* know that I am not merely a handless envatted brain. A consequence of the conjunction of these two premises is that at least one of the conjuncts in the antecedent of the first premise is false- either I do not know that I have hands, or I do not know that my having hands entails that I am not a handless envatted brain. But since the entailment is obvious, the skeptic thinks that we must accept that I do not, after all, know that I have hands.

It is on account of such implications that skepticism is typically viewed as a central problem in epistemology.

1.2 Responding to skepticism

This brief of skepticism can serve as a background against which much of contemporary epistemology can be surveyed; many epistemological theories are motivated, at least in part, by a desire to defend the kind of knowledge we ordinarily take ourselves to possess from skeptical threat. Some theories aim to provide a response that is ideal from a dialectical standpoint, namely, the refutation of skepticism. If we construe what counts as a refutation broadly, a theory that aims for this kind of response might attempt to establish either that the skeptic's thesis is false, or that the skeptic's view is internally inconsistent, or that it is incoherent. Of course the

history of epistemology is rich with arguments employing a number of general strategies that attempt to refute skepticism of various forms. In what follows we shall focus on only two of those strategies, as they are most relevant to the chapters to follow.

1.2.1 Transcendental arguments

A refutation of, say, external world skepticism might try to prove the existence of the material objects the existence of which the skeptic claims to doubt. The lesson of much epistemological investigation is that the prospects for a direct refutation of this sort look to be quite dim. As the skeptic urges, it appears that any justification for our belief in the existence of objects external to us would have to come from within experience; given the skeptic's arguments, it appears that there is no way to justify such beliefs.

But its worth noting that while the skeptic challenges the claim that our experiences carry justificatory force for our external world beliefs, she acknowledges that we *have* such experiences, or that we have thoughts about the external world that we can express in language. Some philosophers have thought that this acknowledgement might itself ground a kind of argument with anti-skeptical force. Arguments of this kind aim to show the impossibility of the skeptical challenge by proving that the truth of certain claims that the skeptic purports to doubt is a necessary condition for the very existence of the phenomena whose existence the skeptic must concede. These *transcendental arguments*, as they have come to be known, are sometimes taken to successfully refute the skeptic in that they conclusively establish the falsity of the skeptic's thesis. Sometimes they are thought rather to show that the skeptic's doubts are unintelligible or self-defeating in some way. Either way, the transcendental strategy promises to deliver a

powerful anti-skeptical result that does not depend on the kind of empirical evidence that has proved time and again to be ineffective against skepticism. For this reason, many have found the transcendental strategy to be an attractive one, and there are a number of examples of such arguments to be found in the literature.

The strategy was inaugurated by Kant (1787), who purports to demonstrate that the obvious temporal ordering of our experience (a fact that the skeptic cannot deny) would not be possible without genuine veridical experiences of an external world. The last century witnessed a revival of interest in Kant's approach, consequent to exposition suggesting that Kant's argument fails. The revival began with Peter Strawson (1959, 1966), who in the latter work argued against Berkeleyan idealism on the grounds that it would not be possible for us to experience the world as consisting of a manifold of particulars distributed in space and time (that we do so experience the world is taken to be undeniable) unless it were true that objects continue to exist when not perceived.

While Strawson and Kant used features of our inner experience to undermine the skeptic, other writers attempted to exploit claims about the nature of thought and language to similar effect. Sydney Shoemaker (1963) employs considerations about the meaningfulness of language to undermine a version of skepticism about other minds. For example, a skeptic about other minds may say that even if it is true of some subject S that S is in pain, it is impossible for me to know that S is in pain. But, argues Shoemaker, such knowledge is an essential precondition for an expression like 'pain' to have any meaning at all. Consequently, since the skeptic's claim features the term 'pain', her claim is either meaningless or it is meaningful but guaranteed to be false.

Hilary Putnam (1981) also employs considerations concerning meaning to argue against external world skepticism of the brain-in-a-vat variety. Putnam's key resource is semantic externalism, on which the meanings of terms used by a subject essentially depend on facts about the subject's environment. Putnam thinks that the BIV skeptic must endorse certain claims, e.g. the claim that it is possible that I am a BIV, and in so endorsing is committed to the claim that the constituent terms successfully refer. On Putnam's view one of the preconditions for successful reference is the existence an appropriate kind of causal connection between the user of the term and the entity to which the term is supposed to refer. But suppose that there is a subject who in fact really *is* a BIV. When such a subject utters 'I am a brain-in-a-vat', his use of the expression 'vat' is not causally connected with the vat that he is actually in- the causal connections trace back to the fragment of the program responsible for producing convincing experiences *as of* vats. So given semantic externalism, the BIV's use of 'vat' refers to some segment of program. But of course the subject in question is not located in a vat hallucination or in a segment of program, but is rather located in a genuine vat, so his utterance is false. Generalizing, no BIV can truly assert 'I am a brain in a vat'; and of course since no non-BIV can truly assert 'I am a brain in a vat', it is not possible for anyone to truly assert 'I am a brain in a vat'. But then the claim whose truth the skeptic is committed to, e.g. 'It is possible that I am a brain in a vat' is false, and BIV-style skepticism is thus refuted.

Finally, Donald Davidson (1991) presents an argument against skepticism about other minds that is transcendental in character. Such a skeptic, presumably, is committed to the claim that she, at least, is a thinking being with contentful mental states. For Davidson, as for Putnam, the content of one's thoughts are at least partially fixed by facts about one's external environment. In Davidson's case, part of what fixes the content of an individual's thoughts is the

way that others are disposed to interpret his behavior in response to external stimuli. Put another way, the content of an individual's thought essentially depends on facts concerning the mental states of others. Consequently, an individual who existed in a world with no other minds would, in effect, fail to be having thoughts at all. But since the other minds skeptic is committed at least to the belief that *she* is having thoughts with a more or less determinate content, her doubting the existence of other minds is incoherent.

Contemporary philosophers' interest in transcendental arguments was diminished by Barry Stroud (1968), who showed to the satisfaction of many that such arguments suffer a fatal weakness.[5] Stroud specifically considers the arguments of Kant, Strawson, and Shoemaker, and presents the advocate of transcendental arguments with an unattractive dilemma: either the argument succeeds against the skeptic only if the elements distinctive of a transcendental argument are superfluous, or the argument has no force against the skeptic. Simplifying somewhat, Stroud's argument runs as follows: all transcendental arguments feature a premise to the effect that the very possibility of our thinking or experiencing in the ways we obviously do essentially depends upon it being the case that the world *is* a certain way. Now since the way the argument claims the world must be is precisely what the skeptic calls into question, such a premise would indeed carry a powerful anti-skeptical force. But, Stroud thinks, the skeptic can respond that the possibility of our thinking or experiencing in certain way requires only that we *believe* that the world is a certain way, e.g. requires only that we believe that there is a world of material objects external to us. But then the second horn of Stroud's dilemma looms; the claim that we must believe the world to be a certain way is wholly ineffective for dealing with

[5] Other criticisms of transcendental arguments appeared at roughly the same time as Stroud's; among them are Korner (1966, 1967), and Gram (1971, 1977). Brueckner (1983, 1984, 1993, 1996) also presents a number of objections to the transcendental strategy. But I focus on Stroud's objection here since space precludes a full survey of objections, and since Stroud's objection is certainly the most well known and is generally regarded as *the* objection that defenders of transcendental arguments must grapple with.

skepticism. So if the transcendental argument is going to be effective, the argument must feature a resource that can block the retreat to the claim that mere belief in what the skeptic wishes to deny is adequate. The only way to do this, claims Stroud, forces the advocate of the transcendental strategy onto the first horn of the dilemma; for any adequate solution must incorporate into the argument some version of what he calls a verification principle, which will be some claim to the effect that it is possible for us to *establish* that the world is a certain way. But if such a principle were true, it would be sufficient on its own for refuting skepticism, and thus the distinctively transcendental structure in which the principle is embedded does nothing to contribute to the refutation. Stroud presents this problem as follows:

> The skeptic can always very plausibly insist that it is enough to make language possible if we believe that S is true [where S is the claim that e.g. there is an external world], or if it looks for all the world as if it is, but that S needn't actually be true… Any opposition to the skeptic on this point would have to rely on the principle that it is not possible for anything to make sense unless it is possible for us to establish whether S is true, or, alternatively, that it isn't possible for us to understand anything at all if we know only what conditions make it look for all the world as if S is true, but which are still compatible with S's falsity. The conditions for anything's making sense would have to be strong enough to include not only our beliefs about what is the case, but also the possibility of our knowing whether those beliefs are true; hence the meaning of a statement would have to be determined by what we can *know*. But to prove this would be to prove some version of the verification principle, and then the skeptic will have been directly and conclusively refuted. Therefore, even when we deal in general with the necessary conditions of there being any language at all, it looks as if the use of a so-called "transcendental argument' to demonstrate the self-defeating character of skepticism would amount to nothing more and nothing less than the application of some version of the verification principle, and if this is what a transcendental argument is then there is nothing special or unique, and certainly nothing new, about this way of attacking skepticism. (Stroud 1968, 255-56)

While Stroud directs his critique at the arguments of Kant, Strawson, and Shoemaker, it appears that a similar line can be pressed against the later arguments of Putnam and Davidson. For both arguments also rely on a principle, namely content externalism, that if true would be sufficient by itself to refute the kind of skepticisms their respective arguments target.

Subsequent to Stroud's critique few writers have attempted the kind of full-blooded transcendental arguments surveyed. But some writers, while agreeing that the kind of transcendental strategy challenged by Stroud isn't viable, have nevertheless tried to develop transcendental strategies that are more modest in their ambitions.[6] In particular, they often try to exploit the claim that Stroud concedes might be correct, namely the claim that we must have certain beliefs in order for certain ways of thinking to be possible. This sort of transcendental claim can be employed to show not that skepticism is false or unintelligible, but rather that the skeptical challenge is *idle*- the beliefs the skeptic tries to call into doubt are such that we are simply unable to give them up.[7]

In chapter 4, I'll develop something akin to a modest transcendental argument, one that exhibits an important kind of anti-skeptical force while avoiding Stroud's criticisms of transcendental arguments. In particular, I'll employ a claim that is broadly transcendental in spirit in such a way as to *block* the skeptic's attempt to engage with the non-skeptic. The anti-skeptical force this provides survives Stroud's criticisms because the transcendental claim I'll endorse, even if true, does not attempt to refute skepticism; indeed the account I'll propose is entirely consistent with skepticism's truth. Furthermore, unlike the kind of claims Stroud targeted, the transcendental claim I'll support does not identify or presuppose any sort of necessary connection between a psychological fact on one hand and a non-psychological fact on the other.

---

[6] Brueckner (1996) is perhaps the most visible critic of even this less ambitious role for transcendental arguments, although his target is restricted to arguments that have a broadly Kantian emphasis.

[7] Strawson (1985) is a good example of this kind of strategy. Other more recent examples include Stern (2000), whose purpose is to identify the forms of skepticism a more modest argument of this kind can be effective against, and even Stroud (1994, 1999).

1.2.2. Contextualism

While a conclusive refutation of skepticism would be welcome, no extant strategy for refutation has enjoyed anything like consensus that it has been successful. Given the difficulties that beset the refutation project, some epistemologists have tried to develop weaker strategies for dealing with skepticism. The contextualist project is among them. Rather than attempting to refute skepticism, the outcome of the contextualist approach, if successful, will be to *mitigate* the threat skepticism poses to our claims to knowledge. It is worthwhile to emphasize the importance of the expression 'claims to knowledge' just used. Contextualism is best understood not as an epistemological theory that tells us something about the property of knowledge, but rather as a semantic theory that tells us something about the *term* 'knowledge'. The guiding idea is that by developing the proper treatment of the semantics of the verb 'knows', and thus a partial semantics for constructions in which the verb paradigmatically features, we might preserve the truth of many of our commonsense knowledge claims while conceding that the skeptic does sometimes get it right in asserting that we do not in fact know certain propositions that seemed epistemically secure.

We can work toward an understanding of the contextualist project by way of an analogy that many contextualists employ to explain the view. What proposition is expressed by the sentence 'I am in Spain now'? A popular answer to this question is that which proposition gets expressed depends on the facts that prevail in the context in which the sentence is uttered, namely, the identity and location of the speaker as well as the time at which she is speaking. My token utterance of this sentence as I write this is false, since I am writing this in New York. My utterance of this sentence in May 2006, however, was true, as is a current utterance of the sentence by a resident of Madrid. The proposed explanation for this is that expressions such as

'I' and 'here' and 'now' are *indexical* expressions, and indexical expressions are a species of context sensitive expressions. On David Kaplan's (1989) influential account of such expressions, they are analyzed along two dimensions. One dimension is the *character* of the expression, which is roughly the rule for the expression's use grasped by competent speakers of the language of which the expression is a part. The *content* of the expression is fixed by the interaction of its character with features of the context in which the expression is uttered. To illustrate, the character of the expression 'I' may be expressed by the following simple rule: 'I' refers to the speaker. The content of the expression on a particular context of use is fixed by the identity of the speaker in that context.

For our purposes, the important lesson to draw is that sentences that feature such expressions do *not* express a proposition acontextually; such sentences require enrichment from context in order to express a proposition, and distinct utterances of the same sentence can express different propositions when there are differences in the respective contexts in which the utterances occur. Consequently, while it is true that the truth-value of distinct utterances of the very same sentence may vary, as in the case described above, the crucial point to note for our purposes is the explanation for variation in truth-value: the truth-*value* varies because of variations in truth *conditions*; since the truth-condition of a sentence is associated with the proposition it expresses, and sentences featuring indexicals express a proposition only when supplemented by context, the truth conditions of such sentences themselves vary with context.

The contextualist presents a semantic treatment of the verb 'knows' that is analogous to the treatment of 'I' just discussed. The constructions in which the verb 'knows' paradigmatically features are knowledge attributions and denials, sentences of the form 'S knows that p' and 'S does not know that p' respectively. For the contextualist, the character of a knowledge

attribution may be expressed as follows: S's belief that p is true and S is in a *good enough* epistemic position with respect to p. This is the dimension of the locution's meaning that remains fixed across every context of utterance. This character interacts with features particular to a given context of utterance to yield a content, the proposition the sentence expresses in that context on that occasion of utterance. The contextual feature that yields a content in concert with the utterance's character is the *epistemic standard* that is operative in that context.[8] Intuitively, an epistemic standard may be thought of as the threshold that a subject's evidence for p must meet or exceed for the subject to count as knowing that p. One central claim of contextualism is that there are many such epistemic standards, each more demanding or less demanding than each of the others. Some of these standards set a low bar for subjects- for example there is some standard that a subject may satisfy even though she possesses very little evidence for some proposition p. Others are quite demanding- for example, there is some standard that is satisfied only by a subject whose evidence for p entails that p.

The other central claim of contextualism is the claim that none of these many standards are privileged as the referent for 'knows'- a given knowledge attribution may be governed by a very demanding standard in one context, and, perfectly appropriately, governed by a quite relaxed standard in another context. Consequently, for the contextualist 'knows' is an expression

[8] Keith DeRose is the contextualist who is perhaps most responsible for pressing the analogy between 'knows' and indexicals, and most contextualists have adopted the model. But one can be a contextualist without endorsing any quasi-indexical thesis for knows; Ludlow (2005) and Schaffer (2004, 2005a, 2007) among others articulate broadly contextualist approaches that reject the indexicality model. But I'll adopt DeRose's account as the template for exposition, since my purpose here is merely to outline the contextualist position and because DeRose is surely the epistemologist who has done the most to articulate and defend the contextualist project. DeRose has brought together his wide body of work on contextualism in his 2009. Since he takes this to be the authoritative presentation of his views, I have used it as my main source for the discussion here.

that can pick out different properties in distinct contexts of utterance, in much the same way that 'I' picks out different individuals in distinct contexts of utterance.[9]

Given the contextualist's quasi-indexical treatment of 'knows', contextualism emerges as a semantic theory concerning the truth-conditions of knowledge attributing (and knowledge denying) sentences. For the contextualist, just as the sentence 'I am in Spain now' does not possess a truth condition outside of a particular context, sentences of the form 'S knows that p' have a truth condition only when supplemented by features of the context in which the sentence is uttered, viz,. when supplemented by the epistemic standard operating in that context.

Contextualists think that an important virtue of the account just sketched is its potential for defending many of our claims to mundane knowledge from skeptical threat. The defense they propose is novel in that its anti-skeptical force is more modest than the kind of defense anti-skeptics have historically sought- contextualism, if correct, offers a way to *contain* skepticism rather than *refute* it. The contextualist concedes that the skeptic can, and sometimes does, succeed in showing that our claims to know even the most ordinary sorts of propositions (e.g. the proposition that you are reading a dissertation chapter right now) are false. But, the contextualist claims, not nearly as often as the skeptic thinks, and only when certain conditions that favor her are in effect. When those conditions are not in effect, our everyday claims to knowledge can be true in just the way we typically take them to be.[10]

---

[9] See Heller (1999), who offers an illuminating discussion of how one should understand the contextualist's thesis that there are many standards for knowledge and the way in which uses of 'knows' are connected with them. I borrow his picture wholesale for the framework I articulate in Chapter 2.

[10] This is one part of the view that contextualists have historically had difficulty explaining in a wholly satisfactory way. What are the conditions that favor the skeptic and what are the conditions that do not? What mechanism or mechanisms operate to fix a very demanding standard in one context and a lax one in others? Lewis (1996) has been one of the few contextualists to take this challenge on by providing a detailed account of the various factors that can push the standard around. DeRose (1995) and Cohen (1998) also propose single rules that purport to explain

To illustrate, let's suppose we want to judge whether an utterance of the attribution "Jones knows that the cat is on the mat" is true. The contextualist tells us that the answer depends on features of the context in which the utterance occurs. If the context is such that the participants in that context are in some sense aware of the possibility that Jones is not in fact looking at an enmatted cat because he is an envatted brain having a deceptive experience *as of* an enmatted cat, then it might be true to say that "Jones does not know there is a cat on the mat" even when Jones really is looking at an enmatted cat in good conditions. Yet if the context is one in which exotic possibilities like this are not, so to speak, in the air, then it can be true to say that "Jones knows that the cat is on the mat", even though Jones's evidence is in all respects the same in each context.

Simplifying things considerably, contextualists by and large endorse some variation on the following thought: the standard that is operating in a given context of utterance is correlated with the range of alternative possibilities to the putatively known proposition p that are salient in the context of an utterance that "S knows that p". In the case of Jones above, the salience of the possibility that Jones is an envatted brain means that the standard that Jones's evidence must meet in order for an attribution to be true in that context is remarkably (perhaps even unattainably) high, and it is contexts like this in which the skeptic prevails.

Evidence for contextualism, and a survey of its competitors

If contextualism is to be more than simply an ad hoc response to skepticism, there must be some independent evidence for its semantic view. Contextualists think that such evidence is

---

shifting standards, and Blome-Tillman (2009) offers a revision of Lewis's model that aims to salvage it from what are commonly thought to be defects.

readily available in the knowledge-attributing and knowledge-denying practices of ordinary competent speakers. Contextualists typically present this evidence by certain test cases. Chief among these cases are the Bank Cases presented by Keith DeRose,[11] which are worth reproducing here in full:

> *Bank Case A*: My wife and I are driving home on a Friday afternoon. We plan to stop by the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.'
>
> *Bank Case B*: My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago an discovered that it was open until noon. But in this case we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, 'Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no, I don't know. I'd better go in and make sure.' (DeRose 2009)

Contextualists think that ordinary speakers will judge that the epistemic appraisal is true in *both* cases, despite the fact that the two knowledge claims appear to contradict one another. Generalizing, the contextualist thinks that the cases reveal the truth of the certain principle. For the sake of a label I'll call it the Compatibility Thesis, and formulate it as follows:

> (CT): It is possible that speaker A1's utterance of 'S knows that p' in one context and speaker A2's utterance of 'S does not know that p' in another context are *both* true, where A1 and A2 are

[11] There are other prominent cases that contextualists employ to underwrite their view. The cases that appear in Cohen (1999) and Fantl and McGrath (2002) are historically important. But DeRose's Bank Cases have without question been the cases most discussed and challenged in the contextualist literature.

simultaneously speaking of the same subject S who possesses the same justification for the same true proposition p.

The contextualist, armed with her semantics for 'knows', claims that the explanation for the superficially conflicting knowledge claims is that the epistemic standard has shifted from one context to the other. CT is made true by those cases in which A1's context is governed by an epistemic standard that is met by S's epistemic condition with respect to p, while A2 is speaking in a context governed by a more demanding standard that S fails to meet. The contextualist will then insist that in such a situation the appearance of contradiction can be explained away: since the epistemic standard operating in a context is part of the content of a use of 'knows' in that context, and different standards are picked out by A1 and A2's respective uses of the verb, the proposition which A2 denies is *not* the proposition that A1 affirms.[12]

The Bank Cases (and other well known cases) have generated a debate in contemporary epistemology that has produced a number of competitors to contextualism, and each can be usefully explicated in reference to them. The first view, which I'll call *classical invariantism*, deserves to be thought of as the historically dominant approach. The classical invariantist holds that 'knows' picks out the very same epistemic standard in every context of use. Accordingly, the classical invariantist denies the contextualist's interpretation of the cases and the generalization CT that they indicate.

For such an invariantist, the knowledge claims in the bank cases (and by extension the utterances of A1 and A2 in CT) are straightforwardly contradictory- in the situation described,

[12] Critics do not allow the contextualist to get away with this claim so lightly. They contend, plausibly, that in certain contexts the intuition that two speakers really are contradicting each other remains very powerful even after the contextualist explanation is given- see (Richard 2004). DeRose attempts to meet this problem with an account of conversational dynamics very much inspired by Lewis's (1979) influential discussion. McFarlane (2007) counters that DeRose's apparatus works for some types of conversations but fails for others.

either the affirmation of knowledge in Case A is false or the denial of knowledge in Case B is false. But the classical invariantist need not pay the apparent cost of outright denying the intuitive pull of the cases. The invariantist could agree that we are in fact pulled toward the contextualist's interpretation of the cases, but claim that our judgments are misled by confusing standards for *warranted assertibility* with standards for truth. What shifts between the cases is the standard for permissibly asserting that one knows. For example the classical invariantist might claim that, given the importance of depositing the check, it is permissible for Keith to assert that he does not know the bank is open in Case B despite the fact that his claim is false.[13]

It is important to note that the classical invariantist is distinguished from the other views to be considered by denying the intuition concerning the truth *values* of the respective knowledge claims. Other views join the contextualist in affirming the intuition, but diverge from the contextualist over the correct explanation for the intuition. The contextualist explains the intuition concerning truth value by appealing to variation in the truth *conditions* of the respective sentences, as described earlier. One of contextualism's primary rivals, subject-sensitive invariantism, explains the cases with a view that is something of a hybrid of the contextualist and classically invariantist positions.[14] The subject-sensitive invariantist agrees with the classical

[13] While this objection to contextualism is not prominent in the literature, DeRose (2009) spends a good deal of time developing such an objection and responding to it. I mention it here because, like DeRose, I find that an objection of this kind appears quite often when I'm involved in discussions of contextualism.

[14] SSI is one of a number of views that are commonly referred to as "pragmatic encroachment" accounts. To my knowledge, the earliest such account is that of Fantl and McGrath (2002). Other accounts that have generated considerable discussion are those of Hawthorne (2004) and Stanley (2005). Pragmatic encroachment views will play little role in the substantive chapters to follow, so I give them short shrift here. Painting with a broad brush, such views unite in the claim that whether a subject counts as knowing depends on practical facts concerning the *subject* of the knowledge attribution. But these views deny that these affect the content of a knowledge attribution or denial, and also deny the contextualist's thesis that it is facts about the context of the *speaker* of a knowledge attribution that make a difference to its evaluation. Yet another player in the debate is MacFarlane's (2005) account, which holds that it is facts about the context of the *evaluator* of a knowledge attribution that are the relevant facts.

invariantist that the semantics for 'knows' is invariant, but agrees with the contextualist that a principle very like CT is true.[15]

The subject-sensitive invariantist accomplishes this by denying an assumption that she will claim is shared by both classical invariantism and contextualism. This assumption is the thesis that Jason Stanley calls *intellectualism*: the thesis that the factors that convert true belief to knowledge are exclusively factors that are truth-conducive, i.e. factors that affect the degree to which the belief is likely to be true. The novelty of SSI lies in its proposal that whether a true belief counts as knowledge can be affected by facts about the *practical situation* of the subject of a given knowledge attribution or denial. While the contextualist explains the truth value intuitions in the Back Cases by identifying shifting referents for the context-sensitive 'knows', the invariantist appeals to the practical differences between the cases: it is because the costs of being wrong in Case B are quite significant that the subject of the attribution must possess stronger evidence for the bank's hours than in the former case.

Some objections to contextualism

Some epistemologists challenge contextualism by arguing that one of its key theses is false or has unacceptable consequences. Others argue that even if the theory is workable it fails to deliver the benefits that contextualists claim for it. We'll start with objections of the first kind, and concentrate only on those objections that seem to be serious ones.

[15] Such theorists will reject CT as I have stated it here, because they will deny that a given attribution and its negation can be true at the same time. Since the facts that fix what it takes to know are solely facts about the subject of the attribution, those facts constrain any simultaneous attributions made by multiple speakers, so that one utterance in a simultaneously occurring attribution-denial pair must be true and the other false. These theorists, however, will presumably accept a variant of CT that specifies that the relevant utterances occur at different times.

A part of the theory that is frequently challenged is the thesis that 'knows' is a context sensitive expression. Critics have sometimes directly challenged this thesis by appealing to linguistic evidence. Jason Stanley (2005) has been the most prominent exponent of this anti-contextualist strategy. Using a number of generally accepted tests for context-sensitivity Stanley purports to show that 'knows' fits into none of the classes of context-sensitive expressions known to semanticists.[16] Stanley claims that the thesis that 'knows' is context sensitive runs afoul of principles and tests for context sensitivity that are generally accepted by linguists.

Stanley begins by challenging the claims of some contextualists that 'knows' can be assimilated to gradable adjectives like 'tall' and 'flat'. Stanley considers two tests for gradability, and concludes that 'knows' either fails the test, or at least performs much less naturally on the test when compared with uncontroversial cases of gradable adjectives. For example, one test is that gradable expressions can accept modifiers like 'really' and 'very'. As Stanley notes, the negation of degree-modifying uses of such expressions can be conjoined with unmodified forms without inconsistency. 'Knows' seems to fail this test, as the following pairs illustrate

1) John is tall, but he's not *really* tall.
1a) #Keith knows that the bank is open, but he doesn't *really* know that it's open.

Stanley also thinks that 'knows', unlike standard cases of gradable adjectives, do not easily pair with comparative constructions. Gradable adjectives like 'tall' have natural comparative forms like 'taller than'. 'Knows' has no such natural form, as the following case pairs illustrate:

[16] It's important to note that Stanley's claim here is controversial. In speaking of 'the class of context-sensitive expressions known to linguists", one might suppose that there is such a generally agreed upon class. There isn't. For example, Cappellen and Lepore (2005) argue for a very restricted list of context-sensitive expressions that doesn't include the expressions Stanley uses to attack contextualism. Other writers, such as Recanati (1986), argue that the class of context-sensitive expressions is extremely broad. Since this debate isn't settled, it is not clear that Stanley's attack is successful.

2) John isn't tall, but he's taller than Sally.

2a) #Keith doesn't know that the bank is open, but he knows it better than Sally.

While Stanley opts to challenge the context-sensitivity of 'knows' directly, other critics have done so indirectly by identifying what appears to be a false consequence of the thesis. One such objection claims that the truth of the contextualist thesis attributes systematic semantic incompetence to ordinary speakers of the language. This has come to be known as the 'semantic blindness' objection to contextualism.[17] We can frame the objection by returning to the closure argument for skepticism and the contextualist response. Most will take it as uncontroversial that once presented with the closure argument, many ordinary speakers will feel a strong intuitive pull toward the skeptic's reasoning. But, the objection goes, why should this be so if the contextualist is correct that 'knows' is context-sensitive? If it were, we would certainly expect most competent speakers to be aware of its context sensitivity, and to immediately respond to the closure argument in a way that mirrors the sort of response that ordinary speakers would make to parallel cases involving clearly context sensitive expressions. DeRose formulates the critic's parallel case as follows:

> It's snowing in Chicago, but not Los Angeles. Cher is in Chicago, talking on the phone with Lonnie, who is in Los Angeles. Cher says, 'It's snowing here.' Lonnie replies, 'It's not snowing here.' Obviously, these two assertions contradict one another, so they can't both be true. So I ask you: Which speaker is making a true claim, Cher or Lonnie? (DeRose 2009).

The critic notes that no ordinary speaker will be taken in by this dialogue, as their competence with 'here' will permit them to immediately see what's wrong with it. But, the critic continues,

[17] Objections of this type can be found in Hawthorne (2004) and Schiffer (1996). Hawthorne's discussion of the objection relies on principles that he thinks underlie our practice of belief reporting, and shows that if ordinary speakers were even dimly aware of the context-sensitivity of 'knows' no such principles would govern our belief reporting practices. Since Schiffer's argument is much easier to summarize compactly, the discussion to follow relies on Schiffer's formulation.

the context-sensitivity of 'knows' in the contextualist's handling of the closure argument is supposed to explain what's wrong with the skeptic's reasoning in just the way that the context-sensitivity of 'here' is supposed to explain what's wrong with the reasoning above. Since ordinary speakers not only fail to see what's immediately wrong with the skeptic's reasoning, but to the contrary often find the reasoning compelling, it seems to follow that if contextualism were correct, then ordinary speakers are in fact blind to the workings of a term that they use frequently.

DeRose (2009) responds to this invariantist objection by bringing it to bear on the invariantist herself. If the results from the Bank Cases are robust, as the contextualist thinks they are, then if invariantism is true it will also be the case that ordinary speakers are blind to the invariantist semantics of 'knows'. For precisely what the Bank Cases reveal is that ordinary speakers strongly incline to the view that whether a use of 'S knows that p' is true will depend on the context in which it is uttered. Consequently, DeRose thinks, the objection from semantic blindness is an objection that cuts both ways, and so carries no special force against contextualism.

Cohen (1999, 2001, 2005) offers a different defense against this objection. Cohen notes that there are precedents for the kind of semantic blindness about 'knows' that the contextualist appears to be committed to. Cohen thinks that ordinary speakers are disposed to view a term such as 'flat' as an absolute term that admits of no degree. But Cohen thinks that we can reveal what we might think of as 'flatness contextualism' to such speakers by raising the standards for flatness, say by drawing their attention to the possibility that there are microscopic bumps on any surface they had been inclined to apply 'flat' to. But such speakers will nevertheless continue to apply the term 'flat' in ordinary contexts. The lesson Cohen draws is that there are different

degrees of context sensitivity in our language, and we can be more or less blind to it. For cases of pure indexicals such as 'I' and 'here' the sensitivity is obvious to everyone, while it is less so for terms such as 'flat', and may be quite unobvious in other cases, such as 'knows'.

One final important criticism of contextualism is relatively recent. The advent of experimental philosophy has generated a number of results that bear on the claims that invariantists and contextualists alike make on behalf of ordinary competent speakers. Philosophers in both camps, in presenting or discussing cases like the Bank Cases, predict that a certain judgment about the cases is one that such speakers will make. Recent epistemological literature features a number of empirical studies of the actual judgments ordinary speakers will make about these cases. While the research is ongoing, it is fair to say that no clear consensus has emerged; for each of the major views in the dispute, there is a study that favors that view over the others.[18]

How I address the contextualist-invariantist debate

The debate between contextualists and invariantists can, I think, be construed in large part as a debate about the truth of CT. Historically the principle has featured, sometimes implicitly, as the key point of contention; many contextualist arguments seem directed toward supporting a judgment that CT is true, while invariantists have worked to impugn the principle. It is reasonable to suppose that their motivation for doing so is a shared belief that the truth of CT is strong evidence in favor of the contextualist's interpretation of the notion of epistemic

[18] Buckwalter (2010) challenges the contextualist's interpretation of the Bank Cases. May, Sinnott-Armstrong, Hull, and Zimmerman (2010) challenge Stanley's interpretation of the Bank Cases, as do Feltz and Zarpentine (2010). The research of Beebe and Buckwalter (2010) suggests that moral factors can influence subjects' verdicts about cases. Schaffer and Knobe (2012) obtain results that suggest that subjects incline to the contrastive view. DeRose (2011) mounts a defense of contextualism from these studies.

standard, while its falsity points toward the correctness of the invariantist's conception. For if an attribution and its negation can both be true even when we hold all the standard factors fixed, the only explanation for the variance in truth-value must be a corresponding variance in the epistemic standard (or so contextualists would have us suppose).

In Chapter 2 I develop and defend a framework from within which I aim to show that although the debate between contextualists and invariantists concerning the connection between CT and epistemic standard is genuine, the disputants have tended to overestimate the significance of that debate. I argue that CT is true and that its truth is independent of the question of which conception of epistemic standard is correct. By employing cases that seem to reveal underappreciated facts about justification, I show how CT can be true even when we assume that the invariantist is correct in insisting that there is a single epistemic standard. On the account that I provide, in some situations what can shift in such a way as to make CT is true is not the epistemic standard, but rather the *question under discussion* in the attributor's context. If the view I'll elaborate is correct, one consequence is that settling the debate about the correct conception of epistemic standard is *not* sufficient for settling the question of whether knowledge attributions and denials can behave in the way CT describes. While the view that emerges exploits basic insights of some recent theories that will be explained in Section 3, the view provides a new foundation for those insights that does a better job of rendering them plausible.

To this point I have touched upon the main epistemological issues that form a part of the background against which the chapters of the dissertation are to be understood. To recapitulate, I have provided an outline of skepticism and some of the dimensions along which a given skeptical theory may be evaluated, and I have explained and illustrated two historically prominent strategies for responding to the skeptic- the transcendental strategy, which aims to

refute the skeptic, and contextualism, which aims to constrain the skeptic by showing how the range of claims affected by skepticism is narrower than the skeptic has led us to believe.

I have provided a glimpse of how I intend to take each of these strategies as a starting point for two of this dissertation's projects: (i) to develop a new kind of transcendental argument that is immune from standard objections while retaining a minimal but nevertheless significant kind of anti-skeptical force; and (ii) to articulate a broadly contextualist account that, by providing a new foundation for explaining and justifying the idea that the semantics of knowledge attributions exhibit important connections to the question under discussion in the attributor's context, holds some promise for offering a new way to view the debate between contextualists and invariantists. In the next section, I will sketch some of the semantic and pragmatic issues that complete the background against which the other chapters are to be understood.

2. A linguistic perspective on questions and inquiry

A given natural language construction may be investigated along any or all of three dimensions: the syntactic, the semantic, and the pragmatic. For an example with which readers are likely to be familiar, consider those constructions which we typically use to communicate facts, to express beliefs, and to transmit knowledge. Syntactically, the kind of sentence by which we typically accomplish these aims is the *declarative* sentence. From the point of view of semantics, it is widely (though not universally) accepted that the meaning of a declarative sentence is to be explained by appeal to the notion of a proposition: a declarative sentence expresses a proposition (or in cases of lexical ambiguity or amphiboly, a declarative may be used

to express any of several propositions), and to understand a sentence is to grasp the proposition it expresses. We may also be interested in finding out what speakers can *do* by uttering declarative sentences, and the various ways in which hearers might interpret what the speaker intends by such an utterance. To undertake these sorts of explorations is to engage in the pragmatics of declarative sentences.

The natural language phenomena that concern us here are colloquially referred to as *questions*. As above, these phenomena can be examined from a syntactic or semantic or pragmatic perspective. The fundamental syntactic notion associated with these phenomena is the notion of an interrogative sentence. Just as a declarative sentence semantically expresses a proposition, an interrogative sentence semantically expresses a question, and to understand a given interrogative is to grasp the question it expresses. The pragmatics of these phenomena involve, among other things, the conditions that must be fulfilled to successfully perform the speech act of questioning. In a more extended sense, the pragmatics centrally involve the practice of inquiry, which is typically inaugurated by this speech act and constrained by the presuppositions of the question expressed.

In what follows I will set to one side consideration of the various theories that purport to account for the syntax of interrogatives; for my purposes readers may rely on their pretheoretical understanding of which sentences are syntactically interrogative (e.g. 'Did John go to the store?') and which sentences are not (e.g. 'Go to the store' and 'John went to the store'). My primary focus shall be to identify the major semantic accounts of interrogatives and situate them in relation to one another, with a secondary focus on some relevant pragmatic issues.[19]

[19] It is important to note that the separation between the syntax and the semantics of questions that I have implied here is not accepted by all question semanticists. A number of contemporary theorists argue that the semantics of

2.1 The semantics of interrogatives

The analogy drawn between declaratives on the one hand and interrogatives on the other seems to fail in at least one respect. The semantics of declaratives are typically held to be truth-conditional; a declarative expresses a content that is individuated by its truth condition, and to understand the sentence is to understand the conditions under which the proposition it expresses would be true. A problem immediately encountered by semanticists interested in other syntactic natural language forms, such as interrogative constructions, is the items expressed by such forms seem not to be the sorts of things capable of being true or false.

Each of the most historically influential semantic accounts of questions attempt to close the apparent gap between declarative and interrogative semantics by assimilating the meanings of interrogatives to the propositions we already knew how to handle. Such accounts are often referred to as *answer set* or *proposition set* accounts, and although they agree that an interrogative expresses a question and not a proposition, they unite in construing the notion of a question in terms of sets of propositions.[20] These accounts typically diverge along two dimensions that cut across one another: there is divergence concerning (i) the linguistic data a theory should be in the business of accounting for, and (ii) which class of propositions the notion of question is to be analyzed in terms of

Many summaries in the literature treat the proto-theory presented in Hamblin (1958) as the first formulation of an answer-set theory. Hamblin himself rejected the theory shortly after

---

questions should be tightly constrained by the syntax of the interrogatives that express them. See Reinhart (1998), Rullmann and Beck (1998a, b), von Stechow (2000), and Sauerland and Heck (2003) for discussion.

[20] In recent years a number of alternative approaches have been motivated by the worry that proposition set theories are not fine-grained enough to account for certain bodies of linguistic data. Ginzburg (1992, 2005) is probably the leading exponent of such alternatives, but other prominent proposals include Krifka (2001) and Aloni (2007).

proposing it, and eventually replaced it with the view developed in Hamblin (1973). The very influential 1973 account proposed that an interrogative denotes every proposition that counts as a possible answer. The criterion for this set is found in the syntactic structure of the interrogative that expresses the question. Consider the interrogative 'Who walks?': since 'who' is construed as a placeholder for people (but not for other kinds of object or for properties) the possible answers will be a list of propositions, each of which is derived by substituting for 'who' a term that refers to one of the people in our unrestricted domain, e.g., 'John walks', 'Sarah walks', 'Socrates walks' etc. When the list is exhausted, the resulting list is what the question expressed by the interrogative 'Who walks?' denotes. The guiding idea of most answer set approaches is that the meaning of an interrogative is a function that that maps worlds to sets of answers. In Hamblin's system, obviously, a given interrogative will map each world to the same set of propositions, since the question it expresses yields a proposition for each and every individual in an unrestricted domain.

Hamblin's analysis serves as a touchstone for many of the accounts that followed. Karttunen's 1977 account revised Hamblin's proposal in response to worries that it did not account for certain empirical data. Chief among these data, for Karttunen, is that Hamblin's proposal does not easily accommodate *indirect questions*, constructions that embed an interrogative within the scope of an operator. Karttunen focused especially on interrogatives embedded within the scope of the verb 'knows' as in the declarative 'John knows who walks'. Karttunen argued that such constructions demand that we restrict the denotation of an interrogative to a subset of its possible answers, namely the set of propositions such that each member is a *true* answer, rather than a merely possible one. Thus on Karttunen's proposal the meaning of an interrogative is a function that maps a given question to a subset (typically, but

not always, a proper subset) of the propositions that count as an answer on a Hamblin-style semantics. Further, unlike Hamblin's theory Karttunen's proposal entails that the set of propositions mapped by the question function can differ from world to world- at some worlds the proposition 'Sally walks' will be among the true answers to the question 'Who walks?' and at other worlds it will be among the false (but still possible) answers.

The narrowing of Hamblin's 1973 proposal continued with the advent of work by Jeroen Groenendijk and Martin Stokhof (1984, 1997), who in a series of jointly authored works developed what has come to be known as partition semantics. On this approach, an interrogative functions to induce a partition of logical space. Each 'cell' of the partition contains one and only one of its answers. Simplifying somewhat, partition semanticists characterize the partition via two properties: (i) exclusivity- each answer (i.e. each cell of the partition) is inconsistent with each of the others, and (ii) exhaustivity- the collection of cells is jointly exhaustive of logical space. For example, in a world whose domain we'll restrict to Tom and Sally, the interrogative 'Who walks?' functions to induce a partition with four cells: 1) both Tom and Sally walk; 2) neither Tom nor Sally walk; 3) only Tom walks; and 4) only Sally walks. In general, where *n* is the number of individuals in the domain in the context in which the interrogative occurs, the partition it induces will have cells numbering 2 to the *n*th power.

Groenendijk and Stokhof first argued for this approach by setting their semantics against Karttunen's and showing that the partition approach does a better job handling the data and respecting ordinary intuitions. I cannot here catalog the corpus of data or the mechanics of the theory in any kind of compact way, but it is worth pausing to identify at a broad level the kind of data that served as their primary motivation for developing partition semantics. As with Karttunen, Groenendijk and Stokhof take interrogatives embedded within the verb 'knows' and

other factive operators as revelatory of the foundations of interrogative semantics. Karttunen claimed that whenever there is a true utterance of a knowledge-attributing sentence that embeds an interrogative it will be the case that the subject of the attribution can supply the true answer to the question expressed by the embedded interrogative (or each of the true answers if there are more than one). For example, if 'John knows who walks' is spoken truly in a situation in which Tom and Sally did the walking, then John knows the propositions *Tom walks* and *Sally walks*.

Groenendijk and Stokhof insisted that our intuitions in cases like this are in fact much stronger. They claim that in such situations, our intuition is that the subject is able to evaluate the embedded interrogative with respect to every individual in the domain, not just the individuals that ground the true answer. In other words, Groenendijk and Stokhof insist that when such a knowledge attribution is true, for any individual *x* in the domain the subject knows that *x* walks if *x* does, and the subject knows that *x* doesn't walk if *x* doesn't. Thus, while Karttunen's approach implies that our intuitions regarding knowledge are weakly exhaustive (a knowing subject is able to provide all of the true answers), Groenendijk and Stokhof think that our intuitions are strongly exhaustive and partition semantics is expressly formulated to respect that fact. When a subject really does have knowledge her information state or her evidential position is such to eliminate every cell of the partition save one.

It is fair to say that although there is not yet consensus concerning which semantics of interrogatives (among the accounts considered here and several others that have not been discussed) is correct, partition semantics has for the better part of three decades underwritten a good deal of work in the semantics of interrogatives, and it probably deserves to be thought of as the leading approach. But it is worth noting that a Hamblinian approach to interrogative semantics is enjoying something of a revival on two grounds. First, there has been a surge of

theoretical interest in the phenomenon of focus (which, speaking very broadly, is the highlighting or special emphasis of a sentential constituent, accomplished in paradigm cases by distinctive vocal intonations on the part of a speaker that distinguishes that constituent from others, although focus may not depend on speaker intentions and may be fixed by contextual parameters). As it happens, the most influential and widely accepted semantics of focus, the *alternative semantics* developed by Mats Rooth (1985), is almost entirely parallel to the interrogative semantics of Hamblin. This parallel, combined with the widely acknowledged success of Rooth's account, looks to provide a strong sort of indirect support for the correctness of Hamblin's interrogative semantics.[21] Furthermore, although I haven't space to review the voluminous data here, recent work by a number of semanticists suggests that many of the results of partition semantics can in fact be duplicated by a broadly Hamblinian approach. [22]

2.2 The role questions semantics plays in this project

In Chapter 3 I will present some additional considerations that favor the choice of Hamblin-style semantics over the partition approach. In effect, I will be denying the claim that motivated the accounts of both Karttunen and Groenendijk and Stokhof. Recall that on Karttunen's view, the truth of 'S knows who V'd' (where 'V' marks a verb clause) requires weak exhaustivity, while partition semantics requires strong exhaustivity. Since the latter entails the former, both can be challenged by showing only that weak exhaustivity is false. By appealing to a special class of cases, I hope to show that weak exhaustivity is in fact false; in some cases, the

[21] See Hagstrom (2003) for discussion that supports this conclusion. Contrary to what I've said above, Hagstrom thinks that a broadly Hamblinian approach, rather than partition semantics, is the dominant view among linguists. But of course that would just strengthen the case for my choice to adopt Hamblin semantics below, so I'm happy to defer!

[22] Hamblinians reach this conclusion by drawing on Heim (1994) to show how strong and weak exhaustivity can be derived in a Hamblin-style framework.

truth of 'S knows who V'd' does not imply that S knows of each person who V'd that they V'd. Such cases, however, are nicely accommodated with the framework of a Hamblin-style semantics.

Given these reasons and others, I have assumed a Hamblin-style semantics for the positive project that I undertake in Chapter 2. This distinguishes my approach from other recent views that make questions and interrogatives relevant in some way to the evaluation of knowledge attributions and denials, most notably the contrastive account of knowledge developed by Jonathan Schaffer in a series of papers (although it remains unclear to what extent this change alone would help my account avoid certain difficulties with contrastivism as presented by Schaffer; it seems likely that this change in semantics would have to be supplemented with other revisions in order to yield a theory that holds the promise of full generality). In order frame and situate the accounts I will present in Chapters 2 and 3, it is to the consideration of these epistemological theories that we now turn.

3. Question-based epistemology

3.1 Dretske and relevant alternatives

The view that best deserves to be considered the progenitor of the group of theories that are the primary focus of this dissertation is the view put forward by Fred Dretske in a series of articles. A central work in this series is "Epistemic Operators" (1970). That article's reputation is due largely to its association with Dretske's notorious rejection of the closure principle. Here I'd like to focus not on that rejection itself, but rather on the on the deeper reasons that lead Dretske to that rejection. As we shall see, these reasons can be separated from a rejection of

closure, and also serve as the springboard for a number of recent theories that have been prominent in contemporary epistemology. These underlying reasons lead to the view that knowledge is essentially a *contrastive* state- it is never that case that 'S knows that p' *simpliciter*; rather when S knows, S knows that p *rather than q*, where q is some state of affairs that conflicts with p . While a full analysis of Dretske's arguments in "Epistemic Operators" is beyond the scope of our concerns here, a brief discussion of some salient points will prove instructive.

Dretske's purpose is to assimilate the propositional operator 'S knows that…' to other operators that express propositional attitudes. Dretske begins by showing that many of these latter operators do not respect the closure principle discussed at length earlier. For example, consider the propositional operator 'S regrets that…'. Suppose that John ate sushi for dinner last night, and suppose further that the meal made him very sick today. The proposition 'John ate sushi last night' presumably entails the proposition 'John existed last night'. It is true that John regrets that he ate sushi last night, and suppose it is true that John is aware of the entailment just noted. Does it follow that John regrets that he existed last night? Surely not, thinks Dretske, and similar examples may be advanced for a host of other propositional attitude operators such as 'S is happy that…', 'S is surprised that..', 'S is angry that…' and others. Using a series of examples that have since become very well known, Dretske claims that the operator 'S knows that…' displays similar behavior, and concludes that a closure principle for knowledge is incorrect. Dretske then uses this rejection to answer skepticism with what has come to be known as the *relevant alternatives* account of knowledge. On such accounts, one need not be in a position to rule out *every* proposition incompatible with p in order to know that p. Rather, what

is required is that one be in a position to rule out every proposition that is incompatible with p that is relevant in the context in which one's knowledge is at issue.[23]

Though it is not often appreciated, Dretske can be interpreted as offering a further related reason for thinking that the view that knowledge is contrastive is plausible. Note that knowledge is commonly thought to bear close ties with explanation. Dretske's analysis above can be extended to the phenomenon of explanation, and one might reasonably expect that deep features of explanation might be paralleled by knowledge. Here's how it works in the case of explanation: suppose that you were having dinner last night with John. You suffer from a severe sushi allergy, so you choose to order a noodle dish instead of one of the many sushi offerings at the restaurant. The explanation for your ordering as you did is that you are allergic to sushi. But note that the proposition that you ordered noodles entails the proposition that you ordered *something*. But the explanation isn't closed under the entailment; the explanation for the fact that you ordered something is *not* that you are allergic to sushi, but rather some other fact such as the fact that you were hungry. Dretske thinks that the reason for the failure of closure in the case of explanation is due to the nature of the phenomenon. As the example above illustrates, explanation is essentially sensitive to contrast; in the situation described, the contrast to the proposition that you ordered noodles is the proposition that you ordered sushi, and your allergy explains why *you ordered noodles rather than sushi*. But the contrast to the proposition that you ordered something is the proposition that you ordered nothing, and your allergy does nothing to explain why you ordered something rather than nothing.

[23] Classic statements of the relevant alternatives approach to knowledge can be found in Goldman (1976), Stine (1976) and Dretske (1970, 1981). It is important to note that embracing a relevant alternatives account does not require the rejection of closure, as Dretske himself (2005) is prepared to concede. The consistency of some relevant alternatives accounts with standard versions of closure is argued in Luper (2006).

There is much more that can be said about Dretske's view, but this much is hopefully clear- a virtue of Dretske's view is that it claims to identify deep symmetries between knowledge and explanation, and given the intuitive connection between these two phenomena, that is some reason for thinking that Dretske's view concerning knowledge is plausible.

3.2 Contrastivism

Dretske's initial proposal was subsequently refined and modified by a number of epistemologists. One approach that is inspired by the relevant alternatives approach has come to be known as *contrastivism*. Although several theories that embrace a contrastive perspective predate it, in what follows I will focus only on the most fully developed contrastive view, the approach proposed by Jonathan Schaffer in a series of prominent papers.[24] I will provide a broad outline of the theory and close by sketching the broadly contrastive view that I will develop in subsequent chapters.

The relevant alternatives approach is probably best understood as a view about what it takes for a subject to have knowledge of some proposition p; it is a theory about the *property* of knowledge. Unlike relevant alternatives, Schaffer's contrastive theory is best understood as a theory concerning the *structure* of the knowledge relation rather than a theory concerning the analysis of the property itself. Schaffer's key claim is that traditional epistemology has been misled by an inappropriate semantics *and* syntax for the verb 'knows'. Traditional epistemology treats the verb as a lexically binary term that expresses a semantically binary relation, canonically represented as *Ksp*, glossed as 'S knows that p'. Schaffer proposes that the verb is

[24] Other contrastive approaches include Morton and Karjalainen (2003), Sinnott-Armstrong (2004), Johnsen (2001), and Blaauw (2004).

actually lexically *ternary* and expresses a three place relation *Kspq*, glossed as 'S knows that p rather than q' *a la* Dretske.

Schaffer takes as his model for 'knows' the verb 'prefers'. The verb 'prefers', just like 'knows', often occurs in constructions that appear to be binary in form, e.g. 'John prefers noodles'. But of course, people don't prefer something *simpliciter*, they prefer it *to* some other thing. To evaluate the truth of a sentence like this, it must be enriched contextually. If the range of preferences considered includes only sushi, at the level of logical form this superficially binary sentence is represented as *John prefers noodles to sushi*, in which case the sentence comes out true. If the range of preferences considered is restricted only to John's favorite dish, at the level of logical form that sentence is represented as *John prefers noodles to Beef Wellington*, in which case it comes out false. In short, although 'prefers' might sometimes appear to be a lexically binary term, analysis reveals it to be lexically ternary.

By appealing to a range of linguistic tests for the presence of constituents that appear only in the logical form but not the surface form of a sentence, Schaffer concludes that there is good reason for thinking that 'knows', like 'prefers', is lexically ternary despite its appearance of binarity. And since lexical structure projects onto syntactic structure, sentences in which the verb features, paradigmatically knowledge attributions and denials, will themselves exhibit a syntactically ternary structure. Consequently, although knowledge attributions are often binary at the level of surface form, e.g. 'I know that John ate noodles', at the level of logical form such sentences are always ternary, e.g. 'I know that John ate noodles rather than sushi'. Generalizing, the deep structure of knowledge attributions is to be understood, contrary to epistemological tradition, as 'S knows that p rather than q'.

With this syntactic/semantic background in place, Schaffer must explicate the third argument place *q*. Taking his cue from Dretske, Schaffer conceives of *q* as recording contrasts to *p*; *q* is to be understood as a proposition that is inconsistent with *p*, or a disjunction of propositions each of which is inconsistent with *p*. What mechanism is responsible for fixing the proposition or propositions that feature in *q*? Although he has vacillated somewhat on the details, the value of *q* is always closely tied to the notion of a *question under discussion* in the context in which the knowledge attribution occurs. How it works is perhaps best revealed by example. Consider the knowledge attribution 'Ann knows that there is a goldfinch in the garden'. On the contrastive account, the superficially binary form of the attribution must have a ternary structure in logical form, 'Ann knows that there is a goldfinch in the garden rather than q'. The third relatum is fixed by the question under discussion in the context of the utterance. If the question under discussion is *What kind of bird is in the garden?*, then the third argument place will be occupied by propositions concerning non-goldfinch birds, e.g. *There is an ostrich in the garden*. And whether the attribution is true will depend on exactly what those contrasts are: if the relevant contrast is ostriches, then it may be that the attribution is true, while it may well be false if the contrast is canaries- Ann's evidence will have to be much stronger to rule out the canary contrast than the ostrich contrast. On the other hand, the set of contrasts that occupy q will be entirely different in kind if the question under discussion is *Where is the goldfinch?*. There, the relevant contrasts will be non-garden locations, e.g., *The goldfinch is at the neighbor's*.

It is in this way that Schaffer's theory takes the evaluation of knowledge attributions to be essentially *question-relative*. Reminiscent of our discussion of contextualism above, contrastivism holds that knowledge attributions like 'Ann knows that there is a goldfinch in the

garden' simply cannot be evaluated outside of context; they are evaluable only relative to a context and the question under discussion within it.[25] To support the claim that declarative knowledge attributions (like the one above and like the kind of attribution that has been almost the sole focus of epistemologists) are question relative in this way, Schaffer notes that it is uncontroversial that other kinds of knowledge attribution such as interrogative attributions (Ann knows whether there is a goldfinch in the garden) and noun attributions (Ann knows the time) are question relative. Since other forms of knowledge attribution are question-relative, a theory that treats declarative knowledge attributions differently looks suspect.

Schaffer (2007) also offers another line of argument to support the claim that the semantics of knowledge attributions and denials are question relative. It's worth outlining this argument as several criticisms of contrastivism in the literature have targeted it (criticisms we will turn to later). Schaffer begins with a survey of the approaches taken to the problem of relating knowledge-wh (e.g. knowing what, knowing when, knowing who, etc.) to the propositional knowledge epistemologists have standardly focused on (knowledge that). The dominant approach in the literature is *reductionism*. According to reductionism knowing-wh reduces to knowledge-that: knowing *what kind of bird is in the garden* is just to know *that* there is a goldfinch in the garden, and to know that there is a goldfinch in the garden is just to know what kind of bird is in the garden. Generalizing, reductionism is the thesis that S knows-wh iff S knows that p. Schaffer argues for the question-relativity of 'knows' by arguing that reductionism is false. In brief, Schaffer argues that knowledge-wh does not reduce to knowledge-that, then

[25] In his early work on contrastivism (2004), Schaffer argues that contrastivism is not merely a version of contextualism. Since then he appears to have backed off this position somewhat, and now appears to accept the view of many critics that contrastivism can be viewed as a species of contextualism.

shows how the semantics of knowledge-wh are clearly question relative, and closes by arguing that knowledge-that must also be question relative.

Schaffer initiates this strategy by identifying an entailment of reductionism and arguing that the entailed thesis is false. Some terminology is useful in understanding Schaffer's argument. *Convergent questions* are distinct questions that have the same true answer. Knowledge-wh claims that embed convergent questions are *convergent knowledge claims*. Reductionism entails that convergent knowledge claims are equivalent. Schaffer argues that convergent knowledge claims are not equivalent, and thus that reductionism is false. Suppose that the bird in the garden is a goldfinch. Consider the following convergent knowledge claims:

(K-wh1) Ann knows whether there is a goldfinch or a raven in the garden
(K-wh2) Ann knows whether there is a goldfinch or a canary in the garden.

According to reductionism, each of these claims is equivalent to the following knowledge-that attribution:

(KT) Ann knows that there is a goldfinch in the garden.

Since each of K-wh1 and K-wh2 are equivalent to KT according to reductionism, reductionism is committed to the further claim that K-wh1 and K-wh2 are also equivalent. But, says Schaffer, they are clearly inequivalent; if Ann were only a novice bird-watcher, K-wh1 could be true while K-wh2 could be false. Thus reductionism is false. Schaffer thinks that the inequivalence is nicely explained if we treat the knowledge relation as including a question: the relation is best expressed via *KspQ*, glossed as "S knows that p as the true answer to Q". Thus, in our envisioned circumstance, Ann knows that there is a goldfinch in the garden as the true answer to the question *Is there a goldfinch or a raven in the garden?*, but she does not know that there is a

goldfinch in the garden as the true answer to the question *Is there a goldfinch or a canary in the garden?*.

It should be noted that although Schaffer appears to have two distinct contrastive accounts that feature different knowledge relations (*Kspq* v. *KspQ*), he treats them as equivalent, reflecting differences in emphasis and presentation rather than differences in content. For Schaffer, the disjunction of propositions that saturate the third argument place *q* in the former relation are just the answers to the question Q that saturates the third argument place in the latter relation.

Given this background, what does it take for a knowledge attribution to be true according to contrastivism? The answer is given by the theory of question semantics that Schaffer adopts for contrastivism. Schaffer (2007) explicitly endorses the partition semantics approach described earlier. Given this assumption, *q* in essence becomes the partition of logical space induced by the question expressed by the interrogative operating in the context, and knowing that p amounts to being in a position to eliminate every cell of the partition except the cell that p occupies. To return to our example above, suppose that the question under discussion is *What kind of bird is in the garden?* And suppose that the domain of candidates is given by {goldfinch, raven, canary}. The partition will be constituted by eight cells representing every possible permutation of the three elements, and it will be true that Ann knows that there is a goldfinch in the garden if her evidential state eliminates every cell but the cell containing *only a goldfinch is in the garden.*

With this outline of the theory in place, what advantages accrue to the theorist adopting a contrastive account of knowledge? Schaffer identifies several. First, he notes that the theory offers the advantages of contextualism without suffering some of contextualism's drawbacks. As noted earlier, contextualism has long been troubled by objections to the effect that it proposes

the wrong understanding of the verb 'knows'. According to these critics, 'knows' simply does not behave linguistically as a context-sensitive expression. Contrastivism avoids this worry because the theory, unlike contextualism, treats the verb as insensitive to context. In a related vein, it is widely held that it is an error to appeal to a context-dependence that does not have a representation in logical form. Contextualism is open to this objection while contrastivism is not, as the latter theory records the context-sensitive feature of the relation explicitly in its third argument place *q* (or equivalently, *Q*).

But contrastivism preserves contextualism's nice anti-skeptical implications, and in much the same way. Given the combination of contrastivism's syntactic thesis and its question-relative semantics, contrastivism makes room for the claim that the knowledge claims we would ordinarily take as true go false only when the skeptic's recherché possibilities are represented in *q*. It can be correct for someone to say that Moore knows that he has hands, just as the skeptic can be correct to say that Moore does not know that he has hands, so long as the questions under discussion in their respective contexts fix the appropriate contrasts. If the question under discussion fixes q in such a way that the former attribution's logical form is represented as *Moore knows that he has hands rather than stumps or latex prostheses* then the attribution is true; if the question under discussion fixes *q* in such a way that the skeptic's knowledge denial is represented as *Moore doesn't know that he has hands rather than vat-images of hands*, then the denial is true. As on contextualism, the appearance of contradiction vanishes once we appreciate that the proposition the former is affirming via the sentence "Moore knows that he has hands" is not the proposition the skeptic is denying when she utters "It is not the case that Moore knows that he has hands".

Although I will borrow liberally from contrastivism in what follows, it is not a part of my project to defend extant versions of the theory from objections. However in closing it's worth noting several objections that have been levelled against the theory as well as Schaffer's responses.[26] As with any philosophical theory, counterexamples to the contrastive account have been proposed in the literature. Peter Baumann (2008), for example, believes that the contrastive account cannot plausibly handle mathematical knowledge. Recall that for the contrastivist, propositions that are inconsistent with the putatively known proposition are part of the content of an attribution of knowledge that p. So it should turn out that any meaningful attribution of knowledge should include such contrast propositions. Since it appears that at least some attributions of the form "S knows that 2+2=4" are true (and hence meaningful), the contrastivist ought to be able to account for such knowledge. Baumann thinks that the contrastivist cannot, because in such cases there is no plausible contrast propositions to be found for mathematical propositions like 2+2=4, and that the contrastivist apparatus appears to do no work in such cases.

Schaffer (2012) replies to Baumann in a way that shows the relevance of the contrastivist apparatus in such cases, by showing how in the right situations contrasts can explain why the subject doesn't know the relevant proposition. Consider Ann, a young child who knows the natural numbers up to 10 and can perform simple addition with her knowledge. It appears that some of the contrast propositions to '2+2=4' include those propositions in which the sum is given as any of the natural numbers up to 10 with the exception of 4. Ann can rule out these contrasts, and when they are the relevant ones it will be true to say that Ann knows that 2+2=4 rather than, say, 7. But since Ann is wholly unfamiliar with any natural number greater than 10, its plausible to think that Ann does not know, say, that 2+2=4 rather than 11; presumably she

[26] In the interests of space, I focus here only on objections that Schaffer has responded to in print. Objections to contrastivism that Schaffer has not yet responded to can be found in DeRose (2009), Kelp (2011), Kvanvig (2007), Neta (2008), Pritchard (2008), and Stalnaker (2004).

cannot know something that she doesn't even understand. To make the utility of the contrastivist treatment of mathematical knowledge manifest, Schaffer says the following:

> If you ask me whether 27 x 513 x -1 is -13851 rather than 13851, I can answer that question in a flash, just by seeing that multiplication by -1 will yield a negative number. But if you ask me whether 27 x 513 x -1 is -13851 rather than -13951, I can't immediately answer that question. So I might well know that 27 x 513 x -1 is -13851 rather than 13851, without yet knowing whether 27 x 513 x -1 is -13851 rather than -13951. (Schaffer 2012, p. 412)

Schaffer concludes that although a persuasive objection to contrastivism using mathematical knowledge might yet be generated, Baumann's objection is ineffective.

Other writers have challenged one of Schaffer's central arguments for the question-relativity of 'knows', namely the argument from convergent questions discussed earlier. Brogaard (2009) argues that Schaffer's argument against reductionism is effective only against a defective version of it; once reductionism is properly understood, the view is not committed to the claim that convergent knowledge-wh claims are equivalent, thus defusing the problem that Schaffer took to be a prime motivation for adopting a question relative view. Schaffer (2009) replies by noting an ambiguity in his original use of 'equivalence' whereon it misled Brogaard to suppose that equivalence is to be understood as identity of truth-conditions. But, Schaffer clarifies, 'equivalence' is supposed to be understood as material equivalence, and with this interpretation made clear he goes on to argue that even the alternative formulation of reductionism that Brogaard employs remains committed to the equivalence of convergent knowledge-wh claims, restoring the problem with which his earlier paper began.

Kallestrup (2009) also challenges contrastivism on the point of the convergent questions problem, though in a slightly different way. Unlike Brogaard, Kallestrup accepts that the problem of convergent questions is a genuine one for reductionism. Kallestrup, however,

maintains that we can honor Schaffer's judgment that convergent knowledge claims can be materially inequivalent, while upholding the canonical knowledge relation Ksp that Schaffer is led to reject. In the special case of knowledge-wh claims, Kallestrup proposes a binary relation that has the conjunctive proposition *p*-and–not-*q* in its second argument place as a way to accommodate Schaffer's intuitions without sacrificing a binary knowledge relation. For Kallestrup, then, for Ann to know whether there is a goldfinch or a raven in the garden is for her to simply stand in the canonical binary relation to the conjunctive proposition that there is a goldfinch in the garden and there is not a raven in the garden. Thus Kallestrup can generate the following relation between knowledge-wh and knowledge-that: if p is true, then Ks(whether-p-or-q) reduces to Ks(p-and-not-q). Schaffer (2009) regiments this claim as follows:

> (KAL): s knows-wh iff Ksp, where p is a conjunctive proposition such that (i) for some true answers p1-pm to the indirect question Q of the wh-clause, p1-pm are conjuncts of p, (ii) for all false answers q1-qn to Q, their negations are conjuncts of p, and (iii) there are no further conjuncts of p.

Schaffer replies by levelling several objections against KAL: (i) KAL violates a principle that Kallestrup appears to have to accept, namely the whether-equivalence principle; (ii) KAL allows for easy knowledge that intuition very strongly suggests the subject lacks; and (iii) KAL unfits knowledge for the conceptual roles that knowledge ascription plays in our intellectual economy, namely the roles of identifying experts and indicating the person best suited to answer the question under discussion. As in the case of Brogaard, Schaffer concludes that contrastivism survives as the better approach to the problem of convergent questions.[27]

[27] Other writers have also pressed against Schaffer's approach to the convergent question problem. See Steglich-Petersen (2014) and Aloni and Egre (2010).

4. The project

The substantive chapters that follow are a motley collection. Each chapter is best viewed as a self-contained project that is independent of the others. But they are unified at a broad thematic level by their concern with bringing the related phenomena of questioning and inquiry to bear on standard epistemological concerns. The relatively recent attempt to incorporate questions and related phenomena into epistemological theorizing strikes me as intuitive and interesting, and I hope to advance that program by identifying heretofore unappreciated benefits of that approach, as well as offering what seems to me the best way to implement it.

In Chapter 2 I articulate an original positive proposal concerning the semantics of knowledge attributions. I aim to respect what I think are genuine insights on the part of contrastivism, most notably that the evaluation of knowledge attributions bears very tight ties to the notion of the question under discussion. I argue that the alternative framework I develop has several attractive features; (i) it reveals deep, and to my knowledge heretofore unrecognized, symmetries between the phenomenon of questions and the widely-endorsed apparatus of possible worlds, such that if that apparatus is useful for epistemic theorizing then so is the phenomenon of questions; and (ii) it can be fruitfully applied to what appears to be a stagnating debate between contextualists on the one hand and invariantists on the other, namely the debate over whether the Compatibility Thesis is true. If my proposal is correct, then it has the following quasi-deflationist consequence: resolving the question of whether CT is true does nothing to resolve the question of whether the contextualist's or the invariantist's understanding of epistemic standard is correct. This consequence, in turn, holds the promise of reconceiving the way in which the semantics of knowledge attributions can be sensitive to a wider range of phenomena than has customarily been supposed.

In chapter 3 I identify a deep problem for any theory that has certain general features shared by Schaffer's contrastivism and the theory I articulate in chapter 2. I begin by identifying a class of cases in which, I claim, our intuitive judgment is that the subject clearly has knowledge of the proposition embedded in the knowledge attribution. I then show that the mechanisms of a broadly contrastivist approach commit both Schaffer and I to implausibly denying that the subject knows. I consider some ways in which the contrastivist might respond, argue that none are ultimately viable, and conclude that contrastivism, as conceived by Schaffer and I, simply cannot be made to work in a fully general way.

A second drawback of the theory presented in chapter 2 is that it has no clear applications to the problem of skepticism. In chapter 4 I aim to show how uncontroversial theses concerning the phenomenon of inquiry, which is conceptually tied to the phenomenon of questioning, can be exploited to anti-skeptical effect. Drawing on our earlier distinction between various forms of skepticism along the dimensions of scope and strength, the argument I'll advance is an argument against a kind of skepticism that has the greatest scope and the greatest strength, namely global skepticism about justification. To endorse this brand of skepticism is to endorse the thesis that for any subject S and any proposition p, S is not justified to any degree in believing that p.

To my knowledge, the strategy I'll develop occupies an unexplored niche in the logical space available for responding to skepticism. Extant proposals embody one of several strategies: (i) show that the skeptical thesis is false, (ii) show that the skeptic's position is incoherent or unintelligible, (iii) acknowledge the force of the skeptical challenge but restrict its application, (iv) dogmatic insistence on our possession of the knowledge we ordinarily claim for ourselves, or (v) deflationism- showing that there is no problem to address. The proposal I'll advance is not by any means entirely original; it borrows its spirit from one of the ways that epistemologists

have historically attempted to realize (i), namely via a transcendental style of argument. But I think there is a way to employ a broadly transcendental perspective toward an end that is different from each of (i)-(v), although it bears a distant connection with (v). The position I'll argue for is the following: while skepticism may in fact be neither false nor unintelligible, and although it may in fact be that none of our claims to knowledge are true, it turns out to be the case that skepticism is a position that we have no *motivation* to investigate. Once a small number of unobjectionable theses are in our possession, it is possible to claim, in an entirely principled way, that skepticism is a position that we may permissibly ignore *even if it is in fact true*.

In sum, although the chapters can be conceived as stand-alone projects, I hope that when taken together they help contribute to a recent and innovative research program that aims to connect questions and phenomena associated with them to historically prominent epistemological concerns, and that they do so in a way that is inspired by extant projects but moves beyond them in philosophically fruitful ways.

Chapter 2: The Question-sensitivity of Knowledge Attributions

1. The Compatibility Thesis

Consider a principle that has featured prominently in recent epistemology. I'll call it The Compatibility Thesis and formulate it as follows:

> (CT): It is possible that speaker A1's utterance of 'S knows that p' in one context and speaker A2's utterance of 'S does not know that p' in another context are *both* true, where A1 and A2 are simultaneously speaking of the same subject S who possesses the same justification for the same true proposition p.

Contextualists embrace CT, treating it as the intuition elicited by a suite of well-known cases.[28] Contextualists explain the intuition and evade the spectre of contradiction by proposing a novel semantics for the verb 'knows'. According to the contextualist 'knows' functions as an indexical expression that can select different *epistemic standards* in distinct contexts of use, where epistemic standard is understood as the threshold that a subject's justification for a proposition must meet in order for an attribution of the form 'S knows that p' to be true. For the contextualist, there are many epistemic standards, and each is either more demanding or less demanding than each of the others. Armed with her semantics for 'knows', the contextualist explains the truth of CT by claiming that S's justification for p is good enough to satisfy the less demanding standard selected by A1's utterance, yet it is not good enough to satisfy the more demanding standard selected by the utterance of A2.

Epistemologists who endorse invariantism about knowledge attributions and denials reject CT. Invariantists insist that 'knows' selects the same epistemic standard in every context of use, one that imposes the same demands on any subject at any time. For the invariantist, the

[28] While the cases presented in Cohen (1999) and Fantl and McGrath (2002) are historically important, the Bank Cases presented in DeRose (1992) have drawn the most attention in the literature and are most closely connected with CT.

standard selected by the utterance of A1 must be the very same standard selected by the utterance of A2, and she sees genuine contradiction where the contextualist sees only its appearance.[29]

The debate between contextualists and invariantists over the correct conception of epistemic standard is genuine. However, the fact that CT is often discussed in the context of this debate makes it easy to conflate the question of whether the semantics of knowledge attributions are sensitive to shifts in the epistemic standard with the question of whether CT is true, and I think that contextualists and invariantists alike have in fact tended to conflate them. This is a mistake. The Compatibility Thesis is true, and I will propose an argument for its truth that has nothing at all to do with whether there are more demanding and less demanding epistemic standards. Even if the invariantist is correct that knowledge attributions are not sensitive to shifts in the epistemic standard, there are other shifts to which they *are* sensitive. But we'll not be in a position to understand the phenomena to which attributions are sensitive unless we first appreciate certain facts about justification.

2. An Example

Alex is a supporter of the U.S. Men's National Soccer Team. Today she is in the stadium for a World Cup qualifying match against Costa Rica. She is seated in the first row behind the goal on the U.S. side of the field. The match is being played in a blizzard, and although Alex can see the action on the U.S. side of the field well, she is having some difficulty following the action and

[29] It should be noted that subject-sensitive and interest-relative invariantists depart from their classically invariantist colleagues at points downstream from this statement of the basic invariantist position. However, since my attention in this paper is confined to CT as formulated and the controversy over epistemic standard, the important differences between the contemporary varieties of invariantism are irrelevant for what follows.

identifying players near the Costa Rican goal far downfield[30]. Consider the following episodes from the match:

> Episode 1: U.S. goalkeeper Brad Guzan kicks the ball laterally to U.S. fullback Omar Gonzalez. The two players are about thirty yards from where Alex is sitting. Alex believes the proposition *Guzan kicked to Gonzalez*.
>
> Episode 2: U.S. forwards Clint Dempsey and Jozy Altidore are attacking the Costa Rican goal far downfield from where Alex is sitting. Altidore kicks the ball to Dempsey, and Alex believes the proposition *Altidore kicked to Dempsey*.

There are obvious similarities and obvious differences between Alex's justifiers in the two episodes. One obvious similarity is that the kind of justifier Alex possesses in each episode is the same, namely her perceptions. One obvious difference is that Alex's justifiers in Episode 1 are stronger than her justifiers in Episode 2. The prima facie explanation for this difference is that in Episode 1 the objects of Alex's perceptions are situated in conditions that are well-suited for vision- Guzan and Gonzalez are situated near Alex, and visibility is good. Call these conditions *C*. In Episode 2, by contrast, the objects of Alex's perceptions are situated in conditions poorly suited for vision- both Dempsey and Altidore are situated a great distance from Alex and visibility is poor due to the falling snow. Call these conditions *C**.

The purpose of these observations is revealed by adding one more episode to the example:

> Episode 3: Guzan is taking a goal kick approximately 30 yards from where Alex is sitting. Far downfield, Dempsey breaks free of the Costa Rican defenders as Guzan kicks, and the ball reaches him not far from the Costa Rican goal. Watching the action, Alex believes the proposition *Guzan kicked to Dempsey*.

In Episode 3, Guzan is situated in condition *C*, and Dempsey is situated in condition *C**. If conditions *C* and *C** explain the judgment that Alex's justifiers in Episode 1 are stronger than her justifiers in Episode 2, then it is reasonable to conclude that the strength of Alex's justifiers

[30] While Alex is fictional, the setting is not. The U.S. played Costa Rica in a heavy blizzard on March 22, 2013 in Colorado.

in Episode 3 is *heterogeneous*; some of her justifiers for the proposition *Guzan kicked to Dempsey* are stronger than others.[31]

In what follows I will argue that situations in which the strength of a subject's justifiers for a proposition are heterogeneous are situations that make for the truth of CT. It is important to make clear that in doing so I simply take for granted the viability of modal epistemology. Broadly speaking, modal approaches explicate key epistemic concepts in terms of possible worlds, counterfactual principles, etc..[32] To be sure, there are epistemologists who believe that modal epistemology faces significant difficulties, among them how to make sense of what it is for a subject to rule out or eliminate possibilities incompatible with *p* and how to account for knowledge of necessary truths. While the account I'll propose inherits these problems, it is beyond the scope of this project to argue for modal epistemology or defend it against objections. The project's sole aim is to develop a framework *within* modal epistemology that proves helpful for understanding the context-sensitivity of knowledge attributions.

The framework is little more than a patchwork of familiar ideas, although I will present them in unfamiliar ways. Thinking about these old ideas from the fresh perspective afforded by novel presentations offers a new way to underwrite the view that knowledge attributions are sensitive to shifts in the *question under discussion* in the attributor's context, and a new appraisal of the contextualist-invariantist debate.

3. A Framework

3.1 A new presentation of the modal theory of knowledge

[31] The idea that the strength of a subject's justifiers for a proposition may be heterogeneous is inspired by the kinds of cases presented by Dretske (1970, 1972).

[32] Prominent examples include DeRose (1995), Lewis (1996), Heller (1999), Stine (1976), Goldman (1976) Dretske (1981), Nozick (1981), Sosa (1999), Pritchard (2005), and Becker (2007).

Modal approaches typically refine the widely shared intuition that S's knowing that p is incompatible with the chance that S is mistaken about p by something very like the following principle: 'S knows that p' is true iff S is in a position to rule out every possible world at which p is false.[33] An alternative but extensionally equivalent statement of this theory exploits the fact that, for any contingently true proposition p, there are many different ways that p could have been false. For instance, there are at least three general ways that the proposition *Guzan kicked to Dempsey* could have been false. There are possible worlds at which it is someone other than Guzan who kicked to Dempsey; at one such world, backup goalkeeper Sean Johnson did the kicking. At other possible worlds, Guzan kicked to someone other than Dempsey; at one such world, Jozy Altidore received it. At still more possible worlds, Guzan did something with the ball other than kick it to Dempsey; at one such world, he threw it to Dempsey. At each of these possible worlds, a proposition that is a *contrary* of the proposition *Guzan kicked to Dempsey* is true; the propositions *Johnson kicked to Dempsey*, *Guzan kicked to Altidore*, and *Guzan threw to Dempsey* are each contraries of the proposition *Guzan kicked to Dempsey*.[34] These observations suggest an alternative formulation of the standard modal theory: 'S knows that p' is true iff S is in a position to rule out every possible world at which a contrary of p is true. This alternative formulation is extensionally equivalent to the original only if it selects just the same possible worlds selected by the original. The truth of the following claim ensures that it does: necessarily,

[33] It is often remarked that principles like this are trivially satisfied whenever p is a necessary truth, yielding the unpalatable consequence that every individual knows every necessary proposition. Lewis, who favors a modal principle very similar to this one, replies to this objection: "there is only one necessary proposition. It holds in every possibility; hence in every possibility left uneliminated by S's evidence, no matter who S may be and what his evidence may be. So the necessary proposition *is* known always and everywhere. Yet this known proposition may go *unrecognized* when presented in impenetrable linguistic disguise, say as the proposition that every even number is the sum of two primes." (Lewis 1996, emphases added).

[34] Propositions p and q are contraries iff there is no possible world at which both are true and there is at least one possible world where both are false. I am being somewhat liberal in treating propositions like *Guzan kicked to Dempsey* and *Johnson kicked to Dempsey* as contraries *tout court*; strictly speaking, both propositions could be true together. But if we grant that verbs are associated with temporal variables (see Stanley 2005 for discussion), and assume plausibly that only one person can perform the kick in question, then their status as contraries seems secure; *Guzan kicked to Dempsey (at t)* and *Johnson kicked to Dempsey (at t)* cannot be true together.

p is false iff a contrary of p is true.[35] Consequently, the set of possible worlds at which p is false is the same set as the set of possible worlds at which a contrary of p is true.

The modal theory thus understood is the first element in our framework for the argument for CT.

3.2 A classification scheme for possible worlds

Consider the possible world at which *Johnson kicked to Dempsey* is true. It is just one among many worlds at which it is someone other than Guzan who accomplished that kick to Dempsey. For example, there is a possible world at which *Guzan's twin kicked to Dempsey* is true, another world at which *Gonzalez kicked to Dempsey* is true. Indeed, it could have been anyone in the unrestricted domain of individuals who accomplished that particular kick to Dempsey. Notice, however, that at each of these possible worlds it is Dempsey to whom the ball was kicked. We may express this similarity by saying that each of these worlds is a *Dempsey world*, and we may collect the Dempsey worlds into one of several categories of worlds at which a contrary of *Guzan kicked to Dempsey* is true. This category can be characterized in different ways. It may be characterized enumeratively by a list of contraries, each of which corresponds to a possible world at which it is someone other than Guzan who kicked to Dempsey. A more elegant characterization is available at the next level of generality; the category may be characterized by the open sentence *x kicked to Dempsey*, since each of the worlds in the category is a world at which a substitution instance of this open sentence expresses a true proposition.

[35] The right-to-left conditional is secured by the definition of 'contrary'. One might worry that the left-to-right conditional only comes out true in *almost* all models; the conditional comes out false in the model in which p and each of its contraries are false. But only possible worlds concern us here, and if we're willing to shoulder a little metaphysical baggage we can show that this model describes an impossible world: suppose that true propositions have truth-makers, and that truth-makers are obtaining states of affairs. Since, necessarily, there are no negative states of affairs, the world at which –p is the only true proposition inconsistent with p is an impossible world.

Another category is grounded in a different similarity among some of the remaining worlds. There are many possible worlds at which Guzan kicked to someone other than Dempsey. While none of these worlds are Dempsey worlds, each of them is similar in that each is a *Guzan world*. As above, we can collect the Guzan worlds into a distinct category characterized either enumeratively or by the open sentence *Guzan kicked to y*. We proceed through the remaining worlds in similar fashion, until we have a suite of categories such that each world at which a contrary of *Guzan kicked to Dempsey* is true is a member of one and only one category.[36]

3.3 The contextualist-invariantist debate explained

Begin with a circle. The center point of the circle is the actual world, where p is true. Arrayed between the actual world and the circle are all the worlds at which a contrary of p is true. Our proposed classification scheme imposes order on this array in the following way: imagine a series of radii through the circle so that the result looks like a sliced pie. Each of the categories described above corresponds to one 'slice' of the pie. Each slice of the pie corresponds to one of the general ways in which p could have been false; the category *Guzan kicked to y*, for instance, includes all and only those possible worlds at which *Guzan kicked to Dempsey* is false because Guzan kicked to someone other than Dempsey. The whole pie is an exclusive and exhaustive classification of the worlds at which a contrary of *Guzan kicked to Dempsey* is true.

[36] I should note that my treatment is somewhat idiosyncratic in that the possible worlds I'm appealing to include none of the worlds at which p is true. This idiosyncrasy makes exposition easier, but the reader is free to substitute whatever treatment is preferred. For example, if you prefer a Lewis-style treatment that includes the p-worlds and orders the whole via a similarity relation, I have no objection; my classification scheme can be imposed on it without affecting the arguments to follow.

Adding concentric circles to the pie serves to represent the different conceptions of epistemic standard put forward by invariantists and contextualists. Each of these circles marks a threshold at which a subject's justification with respect to a proposition becomes good enough to make an attribution of knowledge to that subject true.

Invariantism offers limited options. Retracing the original circle yields the standard endorsed by the *strict* invariantist. Strict invariantists insist on a very demanding standard whereby a subject's justification is good enough for knowledge that p only if she can rule out *every* world at which a contrary to p is true. But since this standard seems to imply skepticism, many invariantists prefer a circle contained somewhere within the original. Drawing a circle far enough from the actual world so that it circumscribes many of the worlds at which a contrary of p is true yields a standard likely to be endorsed by *moderate* invariantists. The moderate invariantist's standard is demanding enough to respect our sense that knowledge is elite among epistemic properties, but permissive enough to preserve our sense that we know many things. Although they disagree about where the circle should be located, both kinds of invariantist insist that there is but one circle and, once drawn, it never expands or contracts. The circle drawn represents the single standard that operates in every context for any proposition p.

The contextualist, on the other hand, adopts much more liberal posture. We may draw as many circles as we please; for the contextualist, there are as many standards for knowledge as there are concentric circles, each more demanding than the one nested immediately within it.[37] Which of these many circles governs a particular attribution of knowledge is determined by the various kinds of pressures contextualists appeal to. Whichever mechanism they favor, contextualists unite in their insistence that a given knowledge attribution may be governed by different circles in different contexts of utterance.

[37] This way of explaining the contextualist's approach to epistemic standard is derived from Heller (1999).

This heuristic helps explain what the contextualist-invariantist debate looks like from the perspective of a modal account of knowledge. There are several points of contact between the two theories. Both sides agree that whether a given utterance of 'S knows that p' is true depends on whether S's justification for p meets the epistemic standard. They further agree that the essential feature of an epistemic standard is its distance from the actual world. In a loose but important sense, the epistemic standard *just is* how far out into the possible worlds S's justifiers for p must extend if an attribution is to be true. The two sides disagree only about whether knowledge attributions are sensitive to shifts in the standard's distance. The contextualist insists that they are. The invariantist insists that they are not, for the simple reason that the standard's distance from the actual world doesn't shift.

3.4 The operation of the basic framework

Enough elements of the framework are in place to explain how it works. Additional elements will come on board along the way. Since one aim of my project is to show how the truth of CT is independent of the debate about epistemic standard, in applying the framework I will operate with the assumptions of CT's opponents. In particular, for the remainder of the paper I assume that the standard that operates in every context is the standard of the *moderate* invariantist.

Assume for the sake of exposition that the following is true:

(T1) 'Alex knows that Guzan kicked to Dempsey' is true just in case Alex can rule out many worlds in the category *x kicked to Dempsey,*

where 'many' takes whatever strength moderate invariantists think appropriate. Given T1, the basic framework yields the verdict that 'Alex knows that Guzan kicked to Dempsey' is true. It must be conceded that the weakness of Alex's perceptions of Dempsey puts her in a poor

position to rule out those worlds in which it is someone other than Dempsey whom Guzan kicked to. But no such worlds are included in the category *x kicked to Dempsey*; every world in this category is a world at which it is Dempsey who receives the ball. Consequently, the weakness of Alex's perceptions of Dempsey is entirely irrelevant to the conditions Alex must satisfy in order for the attribution to be true. What *is* required is that Alex's justifiers must permit her to rule out many of those possibilities in which it someone other than Guzan who kicked to Dempsey. Given the strength of Alex's perceptions of Guzan, Alex satisfies the moderate invariantist standard, and 'Alex knows that Guzan kicked to Dempsey' is true.

Assume now that instead of T1 it is the following that is true:

(T2): 'Alex knows that Guzan kicked to Dempsey' is true just in case Alex is in a position to rule out many worlds in the category *Guzan kicked to y*.

Given T2, the basic framework yields the verdict that 'Alex knows that Guzan kicked to Dempsey' is false, even though there is no change in Alex's justifiers for *Guzan kicked to Dempsey*. What has changed are the possible worlds being appealed to. The worlds that T2 specifies as those Alex must rule out are the worlds at which it is someone other than Dempsey whom Guzan kicked to. Certainly the strength of Alex's perceptions of Guzan put her in a good position to rule out possibilities in which it is someone other than Guzan who kicked to Dempsey. But no such worlds are included in the category *Guzan kicked to y*; every world in this category is a world at which it is Guzan who does the kicking. Consequently, the strength of Alex's perceptions of Guzan is entirely irrelevant to the conditions Alex must satisfy in order for the attribution to be true. What *is* required is that Alex's justifiers must permit her to rule out many of the possibilities in which it is someone other than Dempsey whom Guzan kicked to.

Given the weakness of Alex's perceptions of Dempsey, Alex does not meet the moderate invariantist standard, and 'Alex knows that Guzan kicked to Dempsey' is false.[38]

Our business would be nearly concluded but for the inconvenient fact that, as epistemic principles go, neither T1 nor T2 seem to possess much in the way of credibility. At least not yet. By adding to the basic framework I believe that T1 and T2 can both be made to look plausible.

4. The Framework and Knowledge Attributions

4.1 The invariantist's argument

The invariantist would bring the framework to bear on Episode 3 in the following way: the truth-value of a given attribution or denial of knowledge to Alex depends upon the relation between her justifiers with respect to *Guzan kicked to Dempsey* and the epistemic standard; if Alex's justifiers fail to meet this standard, then the attribution is false and the denial 'Alex does not know that Guzan kicked to Dempsey' is true. But the heterogeneity with respect to strength displayed by Alex's justifiers in Episode 3 means that she is not in a position to rule out some of the worlds that she must; as the discussion of T2 revealed, Alex is in no position to rule out many of the worlds in the category *Guzan kicked to y*. Consequently, Alex's justifiers fall short of the epistemic standard, and 'Alex does not know that Guzan kicked to Dempsey' is true.

The invariantist's stated reasons are all true. Yet it does *not* follow that the denial is true. The conclusion follows only if a certain assumption is granted. It is this assumption, and only this assumption, that justifies the appeal to the attribution-falsifying category *Guzan kicked to y*.

[38] It's worth noting how the contextualist judges these cases. Recall that the contextualist treats epistemic standard as responsive to certain features of a particular context; which features those are depends on the contextualist theory one endorses. Since the cases provide no description of such features, the contextualist withholds her verdict pending elaboration. In the case of T1, if the elaborated details are such that her favored mechanism installs the standard very far out into the category *x kicked to Dempsey* she may claim, *contra* our judgment above, that the attribution is false. In the case of T2, if the details are such that the standard crosses the category *Guzan kicked to y* very near the actual world she may claim, *contra* our judgment above, that the attribution is true.

Evaluating the invariantist's reasoning, therefore, requires the identification and evaluation of this assumption.

4.2 Identifying the assumption

One important feature of the heuristic described in 3.3 is the circle representing epistemic standard. But an equally important feature is the *area* of the circle, the entire space of possibilities contained within it. Just as there are principles concerning the circle that impose requirements on subjects if they are to know, there are principles concerning the possibility space it circumscribes. In assuming that the question of the sensitivity of knowledge attributions is settled once the correct number of circles is determined, both contextualists and invariantists assume that a particular principle concerning the space of possibilities is true: for any p, the worlds that S must be in a position to rule out in order for 'S knows that p' to be true are *all* the worlds contained within the circle.

Since the categories of our classification scheme jointly exhaust the possibility space contained within any circle, we can formulate this assumption in the language of our framework as follows:

> (A) For any p, the worlds that S must be in a position to rule out in order for 'S knows that p' to be true includes worlds from every category of the classification imposed on the set of worlds at which a contrary to p is true.

It is this assumption that the invariantist's reasoning depends on; if A is false, the invariantist's appeal to the category *Guzan kicked to y* is entirely *ad hoc*. But there are reasons for thinking that A is a questionable principle.

4.3 Challenging assumption A

If we assume that the set of worlds S must be in a position to rule out in order to know that p is never the empty set, the following principle concerning the space of possibilities is the only rival to A:

> (A*) For any p, the worlds that S must be in a position to rule out in order for 'S knows that p' to be true includes worlds from some, but not necessarily all, categories of the classification imposed on the set of worlds at which a contrary to p is true.

I will challenge A by defending A*. In defending A*, I owe the advocate of A (i) a description of a plausible and independently motivated mechanism that selects and excludes categories in the manner described by A*, and (ii) reasons for thinking that A* rather than A is the correct principle concerning the possibility space contained within the circle.

4.3.1 The mechanism

I claim that the *question under discussion* is the mechanism that selects and excludes categories in the manner A* describes.[39] Questions are suited for this role because, on a plausible account of their semantics, each individual category is semantically connected with a particular question. Semanticists have proposed a number of analyses of the meanings of questions, and Hamblin (1973) offered a very influential early analysis. For Hamblin, what a question denotes is a set of propositions, namely the set of all the propositions that count as a possible answer to the question. The analysis is grounded by the connections between a declarative like 'John walks' and the interrogative *Who walks?*, and the intuition that the wh-word in the latter serves as a placeholder for 'John' in the former. Hamblin's proposal builds

[39] In bringing questions to bear on epistemic evaluation I am following in the footsteps of others; the accounts of Hookway (1996) and Schaffer (2004, 2005a, 2005b, 2007) each connect knowledge attributions to questions in some way. Readers familiar with the literature on contextualism will recognize my considerable debts to Schaffer's work in particular. Although I depart from his work in several respects, one may view this project as an attempt to develop a new way to justify some of the basic insights of his contrastive account of knowledge, and to reveal principles and consequences that I believe are implicit or underappreciated in that account.

questions from the wh-word of an interrogative, with the end result that a question such as *Who walks?* selects all and only those propositions obtained by installing individuals from our unrestricted domain in the place occupied by 'who'.

With this sketch in mind, notice that each category of the classification just is the exhaustive set of possible answers to one of the questions that p is the answer to. Recall, for example, the category *x kicked to Dempsey*. This category features all and only those propositions obtained by substitution of an individual from our domain for the variable in the open sentence that characterizes the category. Hamblin semantics permits us to connect this category to the question *Who kicked to Dempsey?* as being what is denoted by an utterance of the question. Similarly for the category *Guzan kicked to y* and the other categories of our classification. By identifying this direct semantic link between questions and the categories that A* connects to the evaluation of knowledge attributions, we thereby identify an indirect but important link between questions and knowledge attributions.

Extending the previous argument reveals a more direct link between questions and knowledge attributions. Combining the semantic connection between particular questions and particular categories with other features of our classification scheme allows us to extend the chain of extensionally equivalent transformations begun in Section 3.1 to show that the standard modal theory can be equivalently formulated in terms of questions. The standard modal theory links the evaluation of knowledge attributions with the set of –p worlds: 'S knows that p' is true iff S is in a position to rule out every possible world at which p is false. The discussion of 3.1 showed that this set of possible worlds is the same set as the set of worlds at which a contrary to p is true. Now recall that the categories of the classification scheme jointly exhaust the worlds at which a contrary to p is true. Since each category is linked with one of the questions that p

answers, it appears that the set of worlds at which a contrary to p is true is the same set as the set of worlds jointly expressed by the questions that p answers. By transitivity, then, the standard modal theory is extensionally equivalent to the following theory: 'S knows that p' is true iff S is in a position to rule out each of the worlds jointly expressed by the questions that p answers. In consequence, positing a direct connection from questions to the evaluation of knowledge attributions is plausible precisely to the extent that the standard modal theory is itself plausible.

The justification for the inclusion of questions is completed by a final connection between questions and knowledge attributions. A simple extension of an influential account of the pragmatics of assertion provides good reason for thinking that questions are *present* in contexts of knowledge attribution. According to Stalnaker's account of assertion, an assertion is felicitous only if it is uttered in situations in which the participants are trying to distinguish what is actual from among various possibilities. Such situations are *inquiries*, and inquiries are typically directed by a question under discussion that may be explicit or implicit in the context. Consequently, the pragmatics of assertion dictate that assertions are made in view of a question under discussion.[40] But of course knowledge attributions are assertions; so there is good reason for thinking that the pragmatics of knowledge attribution dictate that knowledge attributions too are made in view of a question under discussion.

The question under discussion, then, is a plausible mechanism for selecting and excluding categories in a principled way. Its plausibility is delivered by its semantic connection to both the categories of our classification scheme and the standard modal theory of knowledge, as well as its robust pragmatic connection to knowledge attributions. Our appeal to the question under

[40] This argument is pieced together from several papers collected in Stalnaker (1999). Arguments for a pragmatic connection between questions and assertions can also be found in Hookway (1996) and Fiengo (2007).

discussion is principled because the semantics that underwrite the appeal come out of an independently motivated linguistic research program.

*4.3.2 Comparing A and A**

Distinct lines of evidence suggest that our ordinary attributive practices comply with A* and contravene A, supporting the conclusion that A* rather than A is the correct principle concerning the space of possibilities contained within the circle.

First line: for *reductio*, suppose that A expresses the principle that ordinary speakers employ (perhaps implicitly) in judging whether 'S knows that p' is true. Given the connections between the question under discussion and our classification scheme, the set of worlds that A requires S to rule out is the same set as the set of worlds jointly expressed by the questions to which p is the answer. This implies a criterion by which we'd expect competent attributors to judge whether 'S knows that p' is true: 'S knows that p' is true only if S is able to answer every question to which p is the answer. But there is significant tension between this criterion and the obvious fact that ordinary attributors do not withhold a knowledge attribution until they have interrogated the subject with each of the various questions that p is the answer to. Indeed, speakers are quite free in attributing knowledge that p to a subject who has successfully supplied p as the answer to a *single* question.

Second line: recent experimental results obtained by Schaffer and Knobe can be interpreted to support A* and to impugn A.[41] Schaffer and Knobe presented two groups of respondents with a story vignette and asked them whether they agree or disagree with the attribution 'Mary knows that Peter stole the rubies'. A part of the vignette revealed that Mary

[41] See Schaffer and Knobe (2012). My Episode 3 shares several points of contact with the cases Schaffer and Knobe employ.

possessed very strong evidence that Peter was the thief but *no evidence at all* that it was rubies that were taken. In the vignette presented to the first group the attribution was made in relation to the question *Who stole the rubies?*, while in the vignette presented to the second group the attribution was made in relation to the question *What did Peter steal?* Schaffer and Knobe found that most respondents in the first group agreed with the attribution while most respondents in the second group disagreed.

Schaffer and Knobe's findings align nicely with A* and poorly with A. Especially relevant are the respondents in the first group, who tended to affirm the attribution despite their awareness that Mary possessed no evidence for what Peter stole. Framing these findings in the language of our framework, it is reasonable to conclude that for the majority of respondents in the first group, given the question *Who stole the rubies?* Mary needed to rule out worlds from the category *x stole the rubies* but she did not need to rule out worlds from the category *Peter stole y*. Put another way, affirming 'Mary knows that Peter stole the rubies' requires only that Mary be in a position to rule out worlds from some, but not necessary all, categories of the classification of worlds at which a contrary of *Peter stole the rubies* is true, just as A* prescribes. If A expressed the requirement that ordinary speakers employ, however, then these results are puzzling; Mary's inability to rule out worlds from every category of the classification should have produced a tendency toward disagreement in *both* groups of respondents.

4.4 The rehabilitation of T1 and T2

When last we encountered them T1 and T2 were disreputable principles. Each had come by its bad name through the apparently capricious selection of one category of the classification to the exclusion of others. With the additions of the phenomenon of the question under

discussion and the principle A*, however, their fortunes have taken a turn for the better. The question under discussion in the attributor's context achieves the selection of one category to the exclusion of others, and A* certifies the selection of one category to the exclusion of others. T1 and T2 have acquired plausible epistemic credentials. All that remains is to put them to use.

5. Consequences and Applications

5.1 The framework and the truth of CT

Since according to CT it *can* happen that an utterance of a given knowledge attribution and a simultaneous utterance of its negation are each true, one confirming instance is enough to confirm CT. Episode 3 provides just such an instance when the full framework is brought to bear on its evaluation.

Suppose attributor Smith and his companions are wondering *Who kicked to Dempsey?* The selection of the category *x kicked to Dempsey* by Smith's question under discussion and the certification of that selection by A* means that T1 captures the conditions Alex must satisfy to make an attribution of knowledge true. The initial discussion of T1 shows that Alex does in fact satisfy these conditions. Consequently, Smith's utterance of 'Alex knows that Guzan kicked to Dempsey' is true.

Now suppose that attributor Jones and her companions are at the same time wondering *Who did Guzan kick to?* The selection of the category *Guzan kicked to y* by Jones's question under discussion and the certification of that selection by A* means that T2 captures the conditions that Alex must satisfy to make an attribution of knowledge true. The initial discussion of T2 shows that Alex fails to satisfy these conditions. Consequently, Jones's utterance of 'Alex does not know that Guzan kicked to Dempsey' is *also* true. In consequence,

the Compatibility Thesis is true too, despite the fact that our argument for it employed the moderate invariantist's conception of epistemic standard, a conception commonly thought to be incompatible with it.

5.2 Questions and the evaluation of knowledge attributions

The idea that questions are connected to the evaluation of knowledge attributions is emphatically not among this project's contributions; as noted earlier, credit for that idea belongs with writers who precede me. But my project does contribute new evidence for that connection. Many epistemologists find the apparatus of possible worlds to be a useful tool for thinking about problems of knowledge. They have typically presented this apparatus in terms of the *contradictory* of p, via locutions like '-p worlds' and 'possibilities in which p is false'. The account developed here offers reason to think that for those who deploy such talk the apparatus of questions comes along for free. With the aid of Hamblin-style question semantics I've tried to show how the equivalent presentation of the apparatus of possible worlds in terms of the *contraries* of p, and the novel classification scheme for worlds that this presentation affords, reveal that the apparatus of questions has been implicit in the apparatus of possible worlds all along. Consequently, insofar as the standard way of thinking about possible worlds proves useful for the evaluation of knowledge attributions, the apparatus of questions does too.

5.3 Ramifications for the contemporary debate

The debate between contextualists and invariantists is prominent in contemporary epistemology. I believe that my argument for the truth of CT has modest but surprising implications for this debate. The contextualist's efforts to vindicate CT on the one hand and the

invariantist's attempts to impugn it on the other suggest a shared belief that the truth of CT would be powerful evidence for the contextualist's conception of epistemic standard and powerful evidence against the invariantist's conception. As the initial discussions of T1 and T2 showed, however, the account developed here is consistent both with the invariantist's conviction that there is a single epistemic standard installed at a fixed distance from the actual world and with the contextualist's insistence that are many epistemic standards situated at varying distances from the actual world.[42] Contrary to prevailing assumptions, then, the bare fact of CT's truth won't help resolve the existing debate concerning epistemic standard. Consequently, we find that the importance of the existing debate to the deep question that underlies it is somewhat diminished. Contextualists and invariantists have addressed the deep question of whether the semantics of knowledge attributions are sensitive to non-standard factors by debating the question of whether their semantics are sensitive to shifts in the epistemic standard's distance from the actual world. Rejecting A in favor of A* shows that this debate is not the sole arbiter of the deep question. Knowledge attributions may in fact be sensitive to shifts in the standard's distance, and contextualists and invariantists remain free to pursue the answer. But I've tried to show that even if this question gets settled it is a *further* question whether the semantics of attributions are sensitive to shifts in other phenomena. The adoption of A* and the framework of which it is a part makes room for the claim that their semantics are sensitive to shifts in the categories contained *within* the standard, no matter which account of the standard's distance turns out to be correct. As a result the existing debate between contextualists and invariantists is reduced to but one aspect of a broader inquiry into the semantics of knowledge attributions and the phenomena to which they are sensitive.

[42] q.v. footnote 38 *supra*.

Chapter 3: A counterexample to the contrastive account of knowledge

1. The contrastive account of knowledge

According to Jonathan Schaffer's contrastive account of knowledge (CAK), the verb 'knows' is univocal and lexically ternary. 'Knows' always denotes the contrastive relation *Kspq*: when S knows, S knows that *p* rather than *q*, where *q* is a disjunction of non-*p* propositions. The relation is context sensitive, because it is question-relative: whether S knows depends on which propositions occupy *q*, and which propositions occupy *q* depends on the question under discussion in a given context.

Other features of Schaffer's account fall out of these characterizations. Since lexical structure projects onto syntactic structure, the univocity and lexical ternicity of 'knows' ensures that the logical form of every knowledge ascription displays a contrast variable *q*. The contrastive character of the relation requires an analysis that is explicitly contrastive: *Kspq* iff (i) *p* is true, (ii) S has conclusive evidence that *p* rather than *q*, and (iii) S is certain that *p* rather than *q* on the basis of (ii).

An illustrative example: imagine that the question *Who stole the sapphire?* governs Holmes's and Watson's investigation. The suspects are Black, Scarlet and Mustard. Given that the question under discussion determines the *q*-propositions appropriate for a particular context, the logical form of Watson's announcement "Holmes knows that Black stole the sapphire" is *Holmes knows that Black stole the sapphire rather than Scarlet stole the sapphire or Mustard stole the sapphire*. Holmes has knowledge iff (i) *Black stole the sapphire* is true, (ii) Holmes has conclusive evidence that Black stole the sapphire rather than Scarlet stole the sapphire or

Mustard stole the sapphire, and (iii) Holmes is certain that Black stole the sapphire rather than Scarlet stole the sapphire or Mustard stole the sapphire on the basis of (ii).

CAK promises the reconciliation of ordinary knowledge with skeptical doubt. Holmes can know that Black stole the sapphire relative to contexts that assign ordinary propositions to *q* as eligible answers to the question under discussion, as in the case described above. Holmes won't know that Black stole the sapphire relative to contexts that assign the skeptic's exotic possibilities to *q* as eligible answers to the question under discussion.[43]

2. A counterexample

For the purposes of this paper, call any context in which a speaker attributes knowledge to a subject an *epistemic context*. Call the proposition embedded in the ascription's *that*- clause the *target proposition*, CAK's *q*-propositions *alternative propositions*, and the sets these compose in particular cases *context sets*. A target proposition is *promiscuous* relative to an epistemic context C iff it can be consistently conjoined with each of the alternative propositions induced in C; a target proposition is *non-promiscuous* otherwise. Finally, call any epistemic context a *promiscuous epistemic context* iff its knowledge ascription features a promiscuous target proposition; an epistemic context is a *non-promiscuous epistemic context* otherwise.[44]

---

[43] This overview of the contrastive account of knowledge is drawn from Schaffer 2004, Schaffer 2005, and Schaffer 2007.

[44] The italicized terms are my own with the exception of *context set*, which is borrowed from Schaffer (2004).

CAK rejects the canonical binary knowledge relation *Ksp*, and treats the contrastive ternary relation *Kspq* as the right model for every epistemic context. The theory also presupposes that there are no promiscuous epistemic contexts.[45] This paper argues (i) that such contexts exist, and (ii) that CAK's contrastive ternary knowledge relation is *not* the correct model for them. The alternative propositions that partially constitute *Kspq* are irrelevant to whether a subject knows a promiscuous target proposition. Furthermore, in at least some promiscuous epistemic contexts the apparatus of CAK rules that subjects lack knowledge, despite powerful intuitions that they possess it.

This initial statement of the problem can be illustrated and elaborated upon by giving a contrastive treatment of an example that parallels The Sapphire Case:

> *The Claret Case*. Imagine that Holmes and Watson are investigating a crime that occurred during a meeting attended by Lestrade, Hopkins, LeVillard, and no others. The question *Who drank claret?* is under discussion. Watson announces "Holmes knows that Lestrade drank claret." Given the question under discussion and the facts described, the alternative propositions that partially constitute the knowledge relation are *Hopkins drank claret* and *LeVillard drank claret*. The logical form of Watson's announcement, then, is *Holmes knows that Lestrade drank claret rather than Hopkins drank claret or LeVillard drank claret*. Holmes has knowledge iff (i) *Lestrade drank claret* is true, (ii) Holmes has

[45] In a footnote to his 2005, Schaffer remarks that "The proposition *q* may be glossed as the disjunction of the relevant alternatives. As such, two constraints on *q* are needed: (i) *q* must be non-empty, and (ii) *p* and all the disjuncts of *q* must be pairwise exclusive" (2005: 239). Constraint (ii) clearly rules out the possibility of promiscuous epistemic contexts.

> conclusive evidence that Lestrade drank claret rather than Hopkins drank claret or LeVillard drank claret, and (iii) Holmes is certain that Lestrade drank claret rather than Hopkins drank claret or LeVillard drank claret on the basis of (ii).

The Claret Case presents a promiscuous epistemic context. The target proposition *Lestrade drank claret* is promiscuous, since Lestrade's drinking is compatible with both Hopkins and LeVillard drinking, neither Hopkins nor LeVillard drinking, or only one of Hopkins or LeVillard drinking. Consequently, evidence sufficient for knowledge of the target proposition needn't rule out the alternative propositions. Such evidence, in fact, needn't include any information at all about either Hopkins or LeVillard, nor need it entail anything about their drinking behavior. The alternative propositions upon which CAK makes knowledge depend seem entirely irrelevant to whether Holmes knows that Lestrade drank claret.

Change the facts of The Claret Case slightly, and CAK confronts a more pressing problem. Imagine that Holmes recovers Lestrade's fingerprints from the claret bottles and glasses, discovers Lestrade in the study thoroughly intoxicated, and obtains testimony that the butler witnessed Lestrade drinking claret. If, unbeknownst to Holmes, Hopkins and LeVillard also drank claret, Holmes could not possess evidence that veridically rules out either of the alternative propositions. Consequently Holmes cannot satisfy CAK's second condition on knowledge, and the contrastivist seems compelled to rule that Holmes does *not* know that Lestrade drank claret, despite powerful intuitions that Holmes's evidence is sufficient for knowledge.

These difficulties trace to the insistence on a ternary knowledge relation that is partially constituted by alternative propositions relevant for a given epistemic context. The Claret Case demonstrates that in some epistemic contexts no contextually relevant alternative proposition is relevant for knowledge. A better model for promiscuous epistemic contexts is the canonical binary knowledge relation *Ksp* that CAK rejects.

3. Replies considered

There are two possible strategies for preserving a contrastive account of knowledge. The contrastivist may (i) pursue revisions of CAK that yield a theory that resists the promiscuity problem, or (ii) preserve the existing theory by neutralizing the counterexample.

The first strategy isn't promising. A natural choice for revision is CAK's treatment of the verb. While Schaffer's proposal treats 'knows' as lexically ternary, he acknowledges that a contrastive theory could instead treat 'knows' as semantically or syntactically binary, though incomplete (2004: 77). It might appear that such proposals hold promise for grounding a theory that can accommodate promiscuity. These proposals, however, differ only in their treatments of the verb. Each is committed to the claim that any theory must endorse to count as contrastive: the knowledge *relation* is ternary and contrastive between a target proposition and contextually determined alternative propositions. Consequently, treating the verb as binary won't help contrastivists escape the promiscuity problem.

This commitment also explains why contrastivists can't profitably revise CAK's analysis of the knowledge relation. Any revised theory that preserves the contrastive character of CAK must endorse the ternary, contrastive knowledge relation. So any proper analysis the

contrastivist might construct must explicate it as contrastive between a target proposition and alternative propositions. All analyses of a contrastive ternary knowledge relation are vulnerable to the promiscuity problem.

Revising CAK offers the contrastivist no better protection against promiscuity. There seems little hope of constructing a theory that is both contrastive in character and immune to the problem presented by promiscuous epistemic contexts. Consequently, defending CAK requires challenging the promiscuity of The Claret Case.

The most promising challenge targets the counterexample's question under discussion. Applying research on the pragmatics of discourse structure to The Claret Case, the contrastivist claims that by accepting *Who drank claret?* as the question under discussion, Holmes and Watson incur a commitment to evaluate it with respect to *each* of Lestrade, Hopkins, and LeVillard. This evaluation is accomplished via (possibly implicit) progression through three subsidiary questions of the form *Did x drink claret?* asked about Lestrade, Hopkins, and LeVillard in turn.[46] The target proposition *Lestrade drank claret* is a complete answer to the first of these subsidiary questions, but only a partial answer to *Who drank claret?*. Invoking this fact and CAK's thesis that a knowledge ascription serves to identify the subject as able to answer the question, the contrastivist asserts that Holmes possesses the ability to answer only the subsidiary question *Did Lestrade drink claret?*. Consequently, the purpose and content of Watson's ascription indicate that the question under discussion has shifted from *Who drank claret?* to *Did Lestrade drink claret?*. Finally, since the context set determined by this question is constituted

[46] See Roberts 2004 and Roberts 1996 for detailed discussion of this connection between questions and subsidiary questions.

only by the proposition *Lestrade did not drink claret*, the contrastivist concludes that The Claret Case actually presents a non-promiscuous epistemic context suitable for contrastive treatment.[47]

Clarifying what the ability to answer consists in helps explain why this question-shifting tactic succeeds. Much recent work in the semantics and pragmatics of questions indicates that a speaker's ability to answer should be understood as the ability to *resolve* a question. Whether an answer resolves a question under discussion is determined by contextual features like the goals and interests of the participants.[48] A subject in possession of a partial answer can resolve, and is therefore able to answer, a question for which context fixes a *mention-some* reading, while only the complete answer resolves a question for which context fixes a *mention-all* reading.[49]

The Claret Case as described invites a mention-all reading of *Who drank claret?*. Supposing that the investigators' goal in identifying claret drinkers is to identify suspects in the crime, only a complete answer will resolve the question. What Holmes knows confers the ability to resolve, and therefore answer, *only* the question *Did Lestrade drink claret?*. Consequently, the contrastivist can justifiably maintain that that this question has implicitly superseded *Who drank claret?* as the question under discussion in The Claret Case.

---

[47] I am grateful to the reviewer for *Philosophical Studies* for bringing this question-shifting strategy to my attention.

[48] Ginzburg 1995 is particularly important. See also Asher and Lascarides 1998, Van Rooy 2003, and Boer and Lycan 1975.

[49] The existence of two readings should not be taken to imply any semantic ambiguity in the question. If there is an ambiguity here, it is pragmatic.

Despite this initial success, the question-shifting tactic provides only a temporary remedy to the promiscuity problem. Retain the original statement of The Claret Case and supplement it with details that make explicit how *Who drank claret?* is introduced as the question under discussion:

> Holmes and Watson are accompanied in their investigation by Inspector Bradstreet. The crime scene is grotesque, and poor Bradstreet needs a glass of claret to regain his composure. However, the investigators fear that some of the available liquor has been tainted with a poison that causes certain and immediate death. Assuming that a living claret drinker suffices to establish its purity, Bradstreet asks "Who drank claret?"

While the analysis of the original Claret Case is unaffected by this supplement, the force of the question-shifting strategy is diminished, since the description now invites a mention-some reading of *Who drank claret?*. Given that the investigators' only goal in pursuing the question is to establish the safety of the claret, just one true partial answer can resolve it. What Holmes knows *does* confer the ability to answer this question. Consequently, contrastivists will find it difficult to induce an implicit shift to the subsidiary question by invoking the partial answer status of the target proposition.

It's equally difficult in the supplemented Claret Case to induce the shift by invoking the regularity with which participants address a question by addressing its subsidiary questions. Given the mention-some reading of *Who drank claret?* the investigators do not incur any commitment to evaluate it with respect to each of Lestrade, Hopkins, and LeVillard. Nor should it be assumed that the investigators must nevertheless implicitly or explicitly progress through the subsidiary questions at least until they obtain an answer that resolves the mention-some reading of *Who drank claret?*. If it is plausible at all, this kind of partial progression is plausible

only for mention-some contexts in which an inquiry must be conducted subsequent to the introduction of a question under discussion, e.g., contexts in which each of the participants suffers from ignorance relative to the question. But suppose that Holmes, unaccompanied on his initial inspection of the crime scene, happens upon the visibly impaired Lestrade disporting among an array of empty claret bottles. He resumes his inspection without meditating further on the encounter. When Bradstreet asks "Who drank claret?" sometime after Holmes returns to the group, the question is no sooner posed than resolved by the partial answer Holmes antecedently possesses. No inquiry into the question is needed subsequent to its introduction; consequently, subsidiary questions have no opportunity to explicitly or implicitly arise.

Since neither the partial answer status of the target proposition nor the regularity with which participants address a question by addressing its subsidiary questions can plausibly induce a transition to the question *Did Lestrade drink claret?*, it appears that *Who drank claret?* can justifiably be maintained as the question under discussion in the supplemented Claret Case. Consequently, the question-shifting strategy fails to defend CAK from the promiscuity problem. The strategy does, however, compel an amendment of the characterization of promiscuous epistemic contexts. The divergent outcomes in the two Claret Cases indicate that an epistemic context is promiscuous iff (i) its knowledge ascription features a promiscuous target proposition, and (ii) its question under discussion is one for which context fixes a mention-some reading.

4. Summary and conclusions

The Claret Case demonstrates that promiscuous epistemic contexts are both real and resistant to contrastive treatment. In such contexts, CAK's alternative propositions play no role in the

acquisition of knowledge. The failure of the most promising strategy for converting epistemic contexts from promiscuous to non-promiscuous indicates that the promiscuity problem is resilient, and the prospects for a revised theory that is both contrastive and immune to the problem are dim. Still, nothing in the foregoing arguments speaks against the possibility of a ternary, question-relative view of knowledge that is free of CAK's contrastive emphasis.[50] Given the many appealing features of CAK, this possibility seems worthy of attention.

[50] I thank the reviewer for *Philosophical Studies* for this point.

Chapter 4: A Defense Against Global Skepticism About Justification

1. A Modest Anti-skeptical Strategy

The philosopher who worries that no refutation of skepticism on the skeptic's terms is likely surely has history on her side. The dustbin of philosophy is brimful with attempted refutations of skepticism- even as far back as two centuries ago the dismal record of such attempts drove an exasperated Kant to declare it a scandal.[51] The situation has scarcely improved in the intervening centuries, and while anti-skeptics should persist in searching for a decisive refutation and hope that one will be produced in the fullness of time, prudence counsels them to explore less ambitious strategies in the interim.

This project articulates an anti-skeptical strategy that is minimally ambitious in that it succeeds only in keeping the skeptic at bay. Before presenting that strategy it is useful to provide a bit of stage-setting. The strategy is effective against global skepticism about justification. To endorse global skepticism about justification (henceforward 'GSJ') is to endorse the following thesis: it is not possible that there is some agent A and some proposition p such that A is justified to even the least degree in believing that p.[52] Consequently, when I use the term 'the skeptic' in what follows I should be understood to refer to all and only those who endorse GSJ. Furthermore, the strategy exploits the contours of a particular context at a particular time. The context is the anti-skeptic's very first encounter with the skeptic. If we conceive of this encounter as a series of 'moves' exchanged between them, the defense I'll

[51] In the preface to the *Critique of Pure Reason* (1787) Kant announces that "it always remains a scandal of philosophy and universal human reason that the existence of things outside of us (from which we after all get the whole matter for our cognitions, even for our inner sense) should have to be assumed merely on faith, and that if it occurs to anyone to doubt it, we should be unable to answer him with a satisfactory proof"

[52] While GSJ is less often defended in the literature than other forms of skepticism, contemporary defenses can be found in Unger (1975), who presents it as the consequence of the connections he draws between knowledge and certainty, and Oakley (1976), who offers it as the outcome of the justification regress problem. Something akin to GSJ can also be extracted from the writings of the Pyrrhonian skeptics.

propose on the anti-skeptic's behalf is to be found in the interval of time immediately following the skeptic's first move. Fully understanding the nature of the defense I'll offer requires that we identify the skeptic's first move and the anti-skeptic's prospects for a response.

There is no mystery about what the skeptic's first move must be. The skeptic must somehow draw our attention to GSJ. The means by which the skeptic performs this move is perhaps a matter on which he should get the last word, but the move is most plausibly accomplished by the skeptic's performance of a speech act, a verbal or written act that makes his audience aware of the content of the skeptical thesis. For the anti-skeptic's part, she enjoys a pair of possible responses from which she might choose. What I'll argue now is that although two responses are possible, there is only one appropriate way for the anti-skeptic to respond to the skeptic's first move.

The character of the appropriate response and the anti-skeptic who will make it is best revealed by contrast with the other way for the anti-skeptic to respond. Consider the anti-skeptic who, upon being made aware of the content of GSJ, opts for immediate rejection coupled with a firm refusal to pay the skeptic any further mind. While such a defense exhibits the enviable virtues of utter effectiveness and complete impregnability, these virtues appear to come at the cost of the anti-skeptic's standing as an intellectually responsible agent- such an anti-skeptic is behaving, somehow, in a spirit of bad faith, and her response is one we are unlikely to endorse.

But in what does her irresponsibility consist? Fault is not to be found in her immediate rejection of GSJ. GSJ diverges dramatically enough from the conviction that we do possess at least some justified beliefs that it seems quite clearly incorrect to suggest that, upon learning the content of GSJ, the only good faith response available to the skeptic's audience is immediate suspension of judgment. Rather, the right thing to say seems to be this: anti-skeptics are well

within their rights to cling to the very powerful prima facie intuition that we do have at least some justified beliefs, *until such time as the skeptic succeeds in meeting his burden of proof by showing them that they ought to change their position*. The fault with the anti-skeptical response being considered here is to be found in its failure to honor the qualification captured in the italicized clause. Put another way, the fault of this sort of anti-skeptic lies, not in her immediate rejection of *skepticism*, but rather in her immediate rejection of *the skeptic*- by refusing to permit the skeptic subsequent moves in their interaction, the anti-skeptic denies the skeptic even a chance to try and meet the burden of proof he bears, and it is in this refusal that her irresponsibility consists.

These considerations indicate that only one response is available to the anti-skeptic who is determined to play fair: she will, upon being presented with GSJ, immediately reject it (such rejection is constitutive of being an *anti*-skeptic) but she must nevertheless remain willing to interact with the skeptic, to allow him an opportunity to meet his burden of proof, and to revise her attitude toward GSJ should he succeed. As we'll soon discover, there are good reasons for thinking that further interaction can only consist in an *inquiry* into whether GSJ is true.

Preliminaries in place, I will argue in this paper for the claim that the anti-skeptic is able, in a way that preserves her standing as a fair and intellectually responsible agent, to decline to inquire into whether GSJ is true. The anti-skeptical payoff: even those who enter their first encounter with the skeptic in a spirit of fair play will emerge with a license to ignore skepticism, a license that remains permanently in force unless the skeptic somehow succeeds, against long odds, in showing that their interaction can proceed via some method other than inquiry.

The strategy is nevertheless quite modest in that it succeeds only in keeping the skeptic at bay. It does nothing whatsoever to show that skepticism is false; indeed, every step in the

reasoning that encapsulates the strategy is consistent with the truth of the skeptic's thesis. But even though there is a chance that the anti-skeptical defense that emerges is only temporary, and even though that defense is less robust than the sort of defense that anti-skeptics have historically aspired to construct, there is good reason to think that it provides anti-skeptics with all the defense they'll need.

2. Inquiry

Inquiry is a kind of practice, and practices have a teleological character- a given practice is developed and refined in view of some goal, and performance of that practice is directed at achieving that goal. What is the goal of inquiry? Understood in one way, the question seems to admit of many answers. Someone might participate in an inquiry simply to satisfy a desire for conversation; to merely secure assent with no regard for truth; to make controversy for controversy's sake; or any number of other reasons. Seen in this way, the goals that individual agents aim to achieve through inquiring may be as various as the agents themselves. But interpreted properly, the question isn't answered by providing just any goal that inquiry *could* be put in the service of achieving. The answer will be a statement of the goal that is unique to inquiry, of the end that it is inquiry's *raison d'etre* to achieve. Understood in this way, the motivation for undertaking inquiry is clear- we believe that we are not as informed as we'd like to be with respect to some matter that interests us and we wish to discover the facts of the matter. In short, the goal of inquiry is the goal of *relieving ignorance*.

This answer to the question is perhaps so plausible as to be obvious, but if a defense of this answer is needed we can begin with the observation that, given their teleological character, practices also have a normative dimension- they are associated with a set of rules and maxims

that are fulfilled when a practice is conducted well and breached when it is conducted poorly. Inquiry has many such norms. For example, it is plausible to think that inquiry is governed by the norm that one ought to try and acquire all of the evidence relevant to the question at issue that is in one's power to attain through reasonable effort. Another plausible norm is that one ought to use only those evidence-gathering processes that one believes are reliable (or, more strongly, a process that is in fact reliable). The goal of relieving ignorance explains the existence of such norms, and this goal also explains our tendency to criticize inquirers who violate them- we criticize a fellow inquirer who has gathers evidence in a desultory fashion, using an unreliable process to boot, precisely because in doing so he is unlikely to help us relieve our ignorance with respect to the matter we are inquiring into.

We have identified the goal of the practice of inquiry, and some of the norms that distinguish the good inquiries from the bad. But what of the practices themselves? Reflection upon the range of practices that have been employed to achieve the goal of relieving ignorance suggests that there is little hope of providing a characterization of the practice of inquiry that is independent of the norms and the goals of inquiry. We're likely to agree that the chemist who formulates an hypothesis that he then tests via a carefully controlled experiment is conducting an inquiry into that hypothesis, and that the philosopher who marshals all the evidence she can gather for and against a given proposition and draws conclusions from that evidence in conformity with rigorously tested rules of inference is surely conducting an inquiry into that proposition. But of course such practices as these are not universal across times, cultures, and individuals. So long as we refrain from undue parochialism, we ought to allow that when *The Iliad*'s Polydamas, seer for the Trojan forces, interprets the flight of an eagle, he is conducting an inquiry into the proposition *The Trojans should press their attack all the way to the Greek ships*.

The astrologer is inquiring into the proposition *Mrs. Jones will enjoy a financial windfall this month* when he consults the alignment of the planets. Even a neglectful father who'd prefer to be none the wiser if his teenage son were failing to attend school is nevertheless inquiring into the matter even if the only effort he makes to gather information is to ask his son, whose penchant for self-serving mendacity has long been manifest, "Have you been going to school?" Poor as some of these inquiries may be, poor inquiries are inquiries nevertheless.

These examples suggest a characterization that, while very broad, offers a correct sufficient condition for a practice's being an instance of inquiry- for any practice, whatever its features, if an agent undertakes that practice with the intention of relieving his ignorance with respect to some issue, then that practice is *ipso facto* an instance of inquiry. This claim is of particular importance for what follows. Note that, since the interaction that is to follow the skeptic's first move will be directed toward discovering the truth-value of GSJ, *any* means by which the skeptic and anti-skeptic conduct their interaction will satisfy the sufficiency condition just described. Consequently, the interaction between the skeptic and the non-skeptic can only consist in an inquiry into whether GSJ is true.

The fact that practices are teleological in character means that there is one more critically important link to be drawn between a practice and the goal that is constitutive of that practice. This link concerns the *efficacy* of practice. Clearly enough, a practice that has developed to achieve a particular goal can be efficacious only if this goal is one that it is possible to achieve. The practice of alchemy, for instance, can be efficacious only if it is possible to transmute lead into gold. But of course since lead cannot be transmuted into gold, the practice of alchemy is a *futile* practice. Similarly, if it is impossible for any individual to relieve his ignorance with

respect to any issue whatsoever, the practice of inquiry will be a futile practice. We may express this relationship with the following thesis:

(F1): If relief from ignorance is impossible, then inquiry is futile.

To avoid confusion, the term 'inquiry', as employed here, should be understood to concern not a particular inquiry into this or that proposition, but rather to concern the process of inquiry as a general phenomenon. Put another way, the term 'inquiry' as used in (F) refers not merely to tokens of the process, but to the very process type itself.

But note that since ignorance with respect to some proposition p is relieved only by the acquisition of justification for p, it turns out that the antecedent of (F1) is equivalent to the formulation of GSJ presented in section 1.[53] So (F1) may be reformulated as follows:

(F2): If GSJ is true, then inquiry is futile.

By applying the equivalence rule of contraposition to (F2) we can derive the central thesis of this paper, the thesis upon which the success of my project almost entirely depends:

(T): If inquiry isn't futile, then GSJ is false.

Anti-skeptics who accept (T) have at their disposal a powerful defense against skepticism. It is tempting to suppose that the defense will proceed by arguing for the premise that inquiry is not futile and then combining it with (T) for a simple modus ponens argument. Desirable as that may be, given that the falsity of GSJ is a precondition for the efficacy of inquiry, it is likely that any proof of the claim that inquiry isn't futile will *ipso facto* be a refutation of GSJ. But it was

[53] Though I use 'justification' here and in my formulation of GSJ, readers are free to substitute whatever they think the factor is that converts true belief into knowledge without affecting the argument. Whatever that preferred factor is, it is quite reasonable to view inquiry as the attempt to secure it for our beliefs.

precisely the difficulty of refuting skepticism that led us to seek other anti-skeptical strategies in the first place, so we'll need to find some other way to exploit (T) to anti-skeptical effect. In particular we must find some argument that employs (T) in such a way as to free us from any obligation to support the claim that inquiry isn't futile. The argument to follow does exactly that; the anti-skeptic who accepts little more than (T) can reason her way to the conclusion that she shouldn't inquire into whether GSJ is true, reasoning that allows her to bypass entirely the issue of inquiry's efficacy.

3. The Argument

The context of the argument is the encounter between the skeptic and the anti-skeptic, in the interval of time immediately following the anti-skeptic's first introduction to GSJ. Our anti-skeptic firmly believes that GSJ is false, but being a responsible sort, she is willing to continue the interaction with the skeptic and to change her mind if the interaction recommends it. Since further interaction must consist in an inquiry into whether GSJ is true, suppose that before interacting any further the anti-skeptic pauses to privately deliberate over the following question: *Should I inquire into whether GSJ is true?* She deliberates as follows:

> I accept thesis (T). Since I do not wish to incur the burden of showing that the actual world is a world at which inquiry is not futile, I'll supplement (T) with the weaker claim that the actual world is either a world at which inquiry is futile or it is a world at which inquiry isn't futile. No matter which of these worlds the skeptic and I inhabit, it is true that for any proposition p, three and only three inquiries into whether p is true are possible: (A) an inquiry that leads me to accept p; (B) an inquiry that leads me to accept neither p nor its negation, and (C) an inquiry that leads me to reject p. However, *in the special case in which p is GSJ*, the *status* of each of these possible inquiries very much depends on which of these worlds the skeptic and I inhabit.

I can see this by assuming, first, that the actual world is a world at which inquiry is futile. Of course, if it is futile to inquire into anything it trivially follows that any inquiry into whether GSJ is true would be futile. So if our world is one in which inquiry is futile then the status of each of A-C, no matter which transpires, is that each would be a *futile* inquiry.

Now I assume instead that the actual world is a world at which inquiry is not futile. By (T), such a world is also a world at which GSJ is sure to be false. Now a necessary condition for a given inquiry's success is that the inquiry discover the facts of the matter being inquired into. Consequently, should either of A or B transpire in such a world then each would have the status of a *failed* inquiry- neither would have revealed the facts that the inquiry was supposed to discover. Should inquiry C transpire in such a world, however, it would reveal the facts of the matter being inquired into. While this does not allow me to infer that C would be a successful inquiry, I can conclude that should inquiry C transpire in such a world it is an inquiry that might be successful.

I now know the status of all three inquiries possible in a world at which inquiry is futile. I know the status of two of the three inquiries possible in a world at which inquiry isn't futile. Thus I know that each of these five possible inquiries would be a *defective* inquiry: the first three are defective because each is futile, and the next two are defective because each is a failure. Only one possibility, inquiry C occurring in a world at which inquiry is not futile, remains to be considered. Call this surviving possibility 'CW2'. Of course it could be that, despite the fact that CW2 reveals the facts that obtain in the world in which it occurs, it suffers from some other problem that renders it defective along with the others. But deciding whether CW2 is itself defective in some way is irrelevant: either it is defective or it is non-defective, and since I have discovered that every possible alternative to CW2 is incontestably defective, and that CW2 will lead to the discovery that GSJ is false, it follows that *any inquiry into whether GSJ is true will either be a defective inquiry or it will be a non-defective inquiry that reveals to me that GSJ is false*.

But the truth of this disjunction allows me to answer my question. For although I remain willing to change my mind about GSJ, what the disjunction reveals is that inquiry is not a procedure by which that change can be effected- I *should not* change my mind on the basis of a defective inquiry, and I *cannot* change my mind on the basis of a non-defective inquiry, since I see, *before I begin to inquire*, that the only outcome of a non-defective inquiry (if such an inquiry there be) is my persisting in my already firm belief that GSJ is false. Consequently, it appears that further interaction with the skeptic cannot be conducted by an inquiry into whether GSJ is true.

But it appears that the *only* way that the skeptic and I can interact is by an inquiry into whether GSJ is true. Consequently, though I remain willing to change my mind about GSJ, it appears that I should decline to interact with the skeptic any further, under the following proviso: I will revisit the possibility of interacting with the skeptic if he can return with a way for us to interact that is neither an instance of inquiry nor a process that will lead to inquiry as a consequence.

Our anti-skeptic has arrived at a surprising position. Recall that the irresponsible anti-skeptic discussed earlier had, by refusing the skeptic any opportunity to respond, arrived at the enviable position of being protected by an utterly impregnable anti-skeptical defense. That anti-skeptic, however, paid a heavy price for the defense by losing standing as a fair and intellectually responsible agent. But it appears that at the close of this argument our anti-skeptic has arrived at nearly the same position- she, too, rejects GSJ and (provisionally) refuses further engagement with the skeptic. But she has not paid the same cost- while the earlier refusal was dogmatic and arbitrary, our anti-skeptic refuses on principled grounds that preserves her standing as a fair and intellectually responsible agent. She remains willing to engage with the skeptic at some future time should the skeptic meet certain obligations, and none of the premises of the argument she uses to reach her position are unfairly prejudicial against the skeptic- none of her premises are equivalent to or presuppose that GSJ is false; indeed every premise employed in her reasoning is consistent with the truth of GSJ. The argument articulated here, in short, makes it *permissible* for anti-skeptics who employ it to ignore global skepticism about justification.

4. Comment

Two principles not mentioned earlier make their first appearance in this argument. Since those principles do a fair share of the heavy lifting, comment on each is called for.

The first principle is the principle that if an inquiry is successful, then the inquiry reveals the facts of the matter being inquired into. When considering this principle, one might worry that the necessary condition it identifies is too restrictive. Especially when thinking of fields of inquiry like the natural sciences, we might think that inquiries into very intractable phenomena

don't often succeed in revealing the facts, but they often are able to conclusively eliminate one of that phenomenon's candidate explanations. Surely, it might be thought, an inquiry that manages to produce such a result for a very difficult issue ought to be considered a success. In light of this, it might be thought that a less demanding principle along the following lines is preferable: if an inquiry is successful, then it conclusively eliminates one or more of the situations thought to be candidates for the fact of the matter being inquired into. Note that this principle represents the most minimal requirement there could be on the success of inquiry- any lesser requirement would make room for the claim that an inquiry that makes no progress whatsoever toward resolving the issue it investigates could still be a success, and this is clearly false.

While I don't think this principle enjoys much plausibility, it's not important to argue for its implausibility since, for the purposes of my argument, it doesn't matter whether it or the principle as formulated in the argument is correct. Here's why: in the inquiry at issue in my argument, unlike an inquiry one might typically find in the sciences, there are only two situations that could be candidates for the facts of matter- the situation in which GSJ is true and the situation in which GSJ is false. Consequently, any inquiry into whether GSJ is true that satisfies one of the necessary conditions *ipso facto* satisfies the other, and any inquiry into whether GSJ is true that fails to satisfy one of the necessary conditions *ipso facto* fails to satisfy the other. Since the competing principle identifies the weakest constraint there could be on the success of an inquiry, and the truth of that principle would make no difference to the case at issue, I'll retain the principle as stated.

The second principle is the principle that one ought not to endorse the results of a defective inquiry, a principle that is implicit in the anti-skeptic's conclusion that she should not change her mind on the basis of a defective inquiry into whether GSJ is true. When considering

this principle, one might worry that most inquiries suffer from some minor defect or other. For example, perhaps one of the participants has failed to fully respect just one of the many norms that govern inquiry. But surely we ought not to disavow an inquiry's result on the basis of such trivial defects, as to do so would plausibly deprive us of knowledge that might otherwise be available to us. I concede the point in the case of such minor defects, but the principle is nevertheless true when applied to the sorts of defects identified in the argument. The defects of futility and failure, surely, must be counted among the severest defects that an inquiry can exhibit, and it is obviously reasonable to refrain from endorsing an inquiry's results when such defects are evident.

Thesis (T) and transcendental arguments

Before turning to the applications and consequences of the argument, it is worthwhile to consider one more principle that features prominently in the argument. Given the crucial role that thesis (T) plays, it might seem that a defense of the argument requires a detailed and stout defense of (T). But I'll do no more than recapitulate the defense provided earlier: (T) is equivalent to (F2) by contraposition. (F2) is true, and plainly so. So, (T) is plainly true. This argument, to be sure, won't persuade anyone already opposed to (F2), if such opponents there be. But it's not necessary for me to find some argument that would convince them. It turns out that the success of my project doesn't depend on whether one accepts (F2), as the rejection of (F2) would serve nearly as well as acceptance. To reject (F2) (as the skeptic might be inclined to) is to reject the claim that the goal constitutive of inquiry is relief from ignorance; and to reject this claim is to reject something that we firmly *believe* about what inquiry is fundamentally like. If our belief is true, then the argument of section 3 does show that inquiry and GSJ are in fact incompatible. But of course it might be that our belief is false, in spite of its apparent

obviousness, and there is some other conception of inquiry that better respects the facts. So perhaps we should allow the skeptic his rejection of (F2); but then the skeptic incurs the burden of characterizing inquiry in a way that renders it compatible with his view- call the phenomenon so characterized *inquiry**. But then the modest goal I have set here still succeeds: the anti-skeptic may, for new reasons, keep the skeptic at bay on principled grounds. For it will be true that the anti-skeptic does not (yet) conceive of inquiry in whatever way inquiry* will turn out to be- the anti-skeptic hasn't any idea what inquiry* looks like or how to participate in such a practice, and the skeptic incurs an obligation to describe this practice if the anti-skeptic is to accept his invitation to engage in it. What is the characteristic goal of inquiry*? What procedures may be employed in order to achieve it? Perhaps the skeptic will be able to tell a plausible story that compels a revision of our beliefs about inquiry; but I suggest that the anti-skeptic may decline to interact with the skeptic until such time as answers to questions like these are forthcoming.

Readers familiar with the literature on skepticism are likely to recognize affinities between (T) and a distinctive kind of claim that has, historically, featured prominently in the anti-skeptical project in epistemology. These claims are typically conditionals that exhibit the following features: the antecedent of the conditional is a sentence expressing a proposition that the skeptic is (presumably) committed to, and the consequent is a sentence expressing a proposition that is inconsistent with the skeptic's thesis. If the conditional is true, its force is to reveal that the falsity of the skeptic's view is implicit in certain things that the skeptic himself must accept. The choice for the skeptic seems to be to admit that his thesis is false, or to be saddled with an incoherent belief system.

To see this sort of claim in action, consider external world skepticism. A skeptic of this sort typically denies that our experience *as of* an external world carries justificatory weight sufficient to secure knowledge of the material world around us. But surely the skeptic cannot deny that we in fact *have* experiences as of an external world, and that those experiences have certain undeniable features. One obvious feature of our experience is that it exhibits a temporal ordering: experience E occurs earlier that experience E*, and so forth. Kant (1787) famously argues that, necessarily, if our total experience exhibits a temporal ordering, then our experiences are generally veridical experiences of objects that exist independently of us. In this way Kant purported to have decisively refuted external world skepticism.

Kant's argument and the anti-skeptical strategy it inaugurated have come to be known as *transcendental argument*, and the key feature of such arguments is the transcendental claim captured in the kind of conditional described earlier. Several well-known transcendental claims, and corresponding transcendental arguments, have been advanced by a number of prominent philosophers. A list of the most prominent claims might run as follows (I have also included my thesis (T) for comparison): [54]

(PS): If we experience the world as containing particulars in a unified spatiotemporal system, then particulars must exist even when unperceived.[55]

(SS): If the sentence 'No agent can know whether someone else in is pain' is meaningful, then skepticism about other minds is false.[56]

(HP): If the sentence 'It is possible that I am a brain in a vat' is meaningful,

[54] None of the following theses (with the exception of (T)) are expressed as the original authors expressed them. I have tried to offer simpler paraphrases of each in order to more clearly reveal the affinities between them and my own thesis (T).

[55] This captures the idea of Strawson (1959) who, by arguing for this conditional, is attempting to refute the kind of idealism proposed by Berkeley.

[56] Sydney Shoemaker (1963) appears to think that one of the preconditions for a term like 'pain' to have meaning at all in a given community is that members of that community be able to have knowledge of the mental states of others in the community.

then it is false.[57]

(DD): If the other minds skeptic knows that her own mental states are genuine and contentful, then other minds skepticism is false.[58]

(T): If inquiry isn't futile, then global skepticism about justification is false.

The similarities between (T) on the one hand and each of (PS)-(DD) on the other might lead one to worry that (T) is also subject to the same kind of criticisms that Barry Stroud (1968) levelled against claims such as (PS)-(DD) in a well-known paper, criticisms that many epistemologists take to be decisive against the transcendental strategy. Simplifying considerably, the essence of Stroud's challenge takes the form of a dilemma: either a transcendental argument fails to refute skepticism, or it refutes it in a way that renders the distinctively transcendental features of the argument entirely superfluous. For the first horn, Stroud notes that transcendental claims typically state that the very possibility of our thinking or experiencing in certain ways essentially depend on it being the case that the world *is* a certain way. But, Stroud thinks, nothing in the argument for such claims excludes the possibility that our thinking or experiencing in such ways requires only that we *believe* that the world is a certain way. And although it might be eminently plausible to think that our experiencing in certain ways presupposes certain beliefs about the way the world is, such a claim obviously does nothing to refute the skeptic's position. The only way to exclude this possibility forces the advocate of the transcendental strategy onto the second horn of the dilemma: excluding that possibility requires that the transcendental argument be supplemented with an assumption that expresses some

[57] Hilary Putnam (1981) famously argues against BIV-style external world skepticism by employing an externalist theory of reference- it is impossible for anyone, including someone who really is a BIV, to truly utter the sentence 'I am a BIV'.

[58] Donald Davidson (1989) argues that at least part of what fixes the content of an individual's thought is the way that others are disposed to interpret his responses to events in the world. Thus, if there were no other minds, there are no acts of such interpretation, and the individual's thoughts would, in effect, be lacking in content or indeterminate in content. Since the consequent is obviously false, the antecedent must be too.

version of a verification principle that states that it is possible for us to *know* that the world is a certain way. But of course such a principle would be, all on its own, sufficient to refute skepticism, and the larger transcendental argument in which it is embedded is entirely irrelevant to the refutation.

I believe that while Stroud's dilemma presents a genuine difficulty for the transcendental arguments he targets, there are several reasons for thinking that they do not apply here. Despite the superficial similarities between (T) and (PS)-(DD), there are several crucial differences. First, the argument grounded in (T) that I have proposed doesn't require that the world actually be a certain way in order for the argument to be successful. At no juncture does my argument feature any premise or assumption to the effect that inquiry must not be futile, nor is there any premise or assumption to the effect that someone must in fact have justification for some candidate belief or other. Since this is so, my argument requires no implicit or explicit reliance on any sort of verification principle, and the second horn of Stroud's dilemma poses no threat to it.

Nor is my argument forced onto the first Horn of Stroud's dilemma. All that is required for my argument is that the very possibility that Stroud concedes might be plausible in fact holds, namely that some phenomenon X essentially depends on our *believing* that the world is a certain way. In order to decide to inquire into some matter, it seems that I must believe, at a minimum, that systematic ignorance is *not* my necessary condition. Stroud's challenge might nevertheless still apply if this claim about belief were being exploited to show that GSJ is false, for as Stroud notes claims about what we must believe carry no force against the skeptic. But as I have repeatedly urged, the argument is not offered in the hopes of refuting the skeptic, and is indeed entirely compatible with the truth of GSJ.

Consequently, although the argument I have articulated embodies the spirit of historically prominent transcendental arguments, I see no good reason for thinking that it is susceptible to the challenges that prove fatal to them.

5. Other Potential Applications

It might be that the anti-skeptic is not the only one who can reason as the anti-skeptic has done. Consider first the agnostic; the agnostic is that (presumably rare) soul who, upon being presented with GSJ, forms no opinion either way but is willing to proceed with the skeptic in a spirit of curiosity. Suppose that before proceeding the agnostic pauses to consider the question *Should I inquire into whether GSJ is true?*, and suppose that she proceeds through just the same reasoning the anti-skeptic uses to arrive at the disjunction *any inquiry into whether GSJ is true will either be a defective inquiry or a non-defective inquiry that leads to the discovery that GSJ is false*. Suppose that the agnostic reasons from this point as follows:

> But now I see that I have no *motivation* to inquire into whether GSJ is true. For any proposition p, my motivation for inquiring into whether p is to arrive at a judgment about which claim I should endorse. But in the special case in which p is GSJ, I don't need to actually go through with the inquiry in order to make this judgment. Since I should not endorse the results of a defective inquiry, and since I already see that the belief that GSJ is true is produced only by defective inquiries, I already know that I should not endorse the claim that GSJ is true *on the basis of inquiry*. Furthermore, I have discovered that a non-defective inquiry must produce the belief that GSJ is false. I don't know whether there are any non-defective inquiries, and even if there are I don't know whether the inquiry the skeptic and I will participate in will be one of them, so I don't know right now that I should endorse the claim that GSJ is false. But I *do* know right now that *if* there is a claim that I should endorse on the basis of an inquiry into whether GSJ is true, it can only be the claim that GSJ is false.
>
> So I can't be motivated to address GSJ via inquiry. But, although I remain willing to address both GSJ and the skeptic, I do not see how the skeptic and I can continue to interact through anything other than an inquiry into whether GSJ is true. Consequently, though I remain willing find out whether GSJ is true, it appears that I should decline to interact with the skeptic any further, under the following proviso: I will

> revisit the possibility of interacting with the skeptic if he can return with a way for us to interact that is neither an instance of inquiry nor a process that will lead to inquiry as a consequence.

Most surprising of all is the possibility that even the skeptic can employ the argument articulated here, so long as the skeptic can be induced to swallow thesis (T). Let's call a skeptic anyone who has previously inquired into whether GSJ is true and come to the belief that it is true. Suppose that the skeptic employs just the same reasoning that led to the italicized disjunction in the original argument (changing the present tense constructions in the original argument to past tense constructions as needed). From this point the skeptic can produce nearly the same reasoning produced by the agnostic above, as follows:

> But now I can see that the inquiry that led me to accept GSJ was defective- that inquiry was either futile or it was a failure. But surely it is true that I ought not to endorse the results of an inquiry that I can see was defective. So I should withdraw my assent to the proposition that GSJ is true.
>
> Should I undertake a new inquiry into whether GSJ is true? The answer seems to be no, as it seems I have no motivation to inquire. For any proposition p, my motivation for inquiring into whether p is to arrive at a judgment about which claim I should endorse. But in the special case in which p is GSJ, I don't need to actually go through with the inquiry in order to make this judgment. Since I should not endorse the results of a defective inquiry, and since I already see that the belief that GSJ is true is produced only by defective inquiries, I already know that I should not endorse the claim that GSJ is true *on the basis of inquiry*. Furthermore, I have discovered that a non-defective inquiry must produce the belief that GSJ is false. I don't know whether there are any non-defective inquiries, and even if there are I don't know whether the inquiry I am considering will be one of them, so I don't know right now that I should endorse the claim that GSJ is false. But I *do* know right now that *if* there is a claim that I should endorse on the basis of an inquiry into whether GSJ is true, it can only be the claim that GSJ is false.
>
> So I can't be motivated to address GSJ via an inquiry into whether GSJ is true. But I do not see how to address GSJ through anything other than an inquiry into whether it is true. So unless I can find another way to address the theory, I should refrain from considering it any further.

If the variations attributed to the agnostic and the skeptic enjoy the same strength as the reasoning employed by the anti-skeptic, the central argument of this paper generalizes almost universally. Anti-skepticism, agnosticism, and skepticism together exhaust the stances that are associated with GSJ. If proponents of each stance can come to see that they have good grounds upon which to ignore GSJ in a principled way, then it appears that *everyone* can ignore GSJ in a principled way. But the 'almost' above is important: *permissibly* ignoring the theory requires that one possesses and can provide a principled basis for doing so. The arguments show that such a basis exists, but an individual possesses that basis and is in a position to provide it only if he has run through and accepted the argument himself. It appears that those who are unaware of the argument are not yet in any position to permissibly ignore skepticism.

6. Looking Forward

Although we have explored the possibility that the central argument generalizes to the agnostic and the skeptic, the focus of the project has been the encounter between the skeptic and the anti-skeptic, so I'll close by considering each's prospects moving forward.

The strategy articulated here blocks the skeptic's attempt to transfer the dialectical burden to the anti-skeptic, by showing that the skeptic has more work to do before that transfer can be effected. We have in passing identified two avenues by which the skeptic might complete that work in such a way as to revitalize the possibility of further interaction between the skeptic and the anti-skeptic. It's worthwhile to more closely examine the nature of the labors the skeptic has before him.

One avenue available to the skeptic is to acknowledge that inquiry into whether GSJ is true is in fact the only procedure by which their interaction can proceed, but to deny that the conception of inquiry that underwrites my anti-skeptical argument is the correct one. Carrying this denial through will require the skeptic to articulate a new conception of inquiry that will not make room for the same kind of anti-skeptical defense. And to deny room for that defense, the new conception of inquiry (i) must not include relief from ignorance among its goals, and (ii) must explain how, and in what ways, the anti-skeptic could be motivated to participate in inquiry, especially in light of the fact that the new conception of inquiry must be compatible with the thesis that universal ignorance is our necessary condition.

The other avenue available to the skeptic is the avenue left open at the conclusion of the argument: to find another way in which the interaction between the skeptic and the anti-skeptic can proceed. Earlier, we offered what appear to be very powerful reasons for thinking that there is no other way, that inquiry is the only way in which their interaction can proceed. Should the skeptic nevertheless persist with this strategy, success requires the skeptic to identify an alternative procedure that satisfies two conditions: (i) the alternative must not itself be an instance of inquiry, obviously, but (ii) neither can it be a procedure that will generate inquiry as a downstream consequence- as soon as inquiry becomes a feature of the interaction between the skeptic and the anti-skeptic, the anti-skeptic can successfully employ the defense articulated here to suspend it. But if the sufficiency condition described in section 2 is correct, finding an alternative procedure that satisfies both conditions requires that the skeptic find an alternative procedure which the anti-skeptic will agree to undertake with no intention of discovering the truth-value of GSJ. As soon as this intention becomes a feature of their interaction, the alternative procedure is sure to be an instance of inquiry by the sufficiency condition presented

earlier. But, considering the daunting range of practices discussed earlier that are accommodated by this sufficiency condition, it's very difficult to see what sort of procedure could be available to the skeptic, and should he find such a procedure we might expect it to be exceedingly strange.

But suppose that the sufficiency condition is defective in some way, and that some other condition is preferable. It remains the case that certain procedures are such that any plausible alternative condition must judge them as cases of inquiry. Chief among them is the procedure by which skepticism has, to my knowledge, been exclusively addressed by those who have considered GSJ, and that is the methodology traditionally favored by philosophers- the gathering of all the evidence that bears on a proposition, both the evidence for and the evidence against, and drawing conclusions from that evidence in conformity with the rules of inference. If any practice is paradigmatic of inquiry, this practice is, and the skeptic is thus barred from employing an alternative procedure that has these features.

I will not attempt show in advance that the skeptic cannot successfully complete the work he has ahead. But unless there are other avenues available to the skeptic beyond the pair I have identified here, then the difficulties associated with each stand between the skeptic and success, and no argument is needed to see that these difficulties are severe. In consequence, although the anti-skeptic perhaps cannot be certain that it is so, she (and perhaps the agnostic and even the skeptic himself) can certainly be confident in the belief that her business with global skepticism about justification has permanently concluded.

References

Aloni, M., D. Beaver, B. Clark, and R. van Rooy (2007). The dynamics of topics and focus. In M. Aloni, A. Butler, and P. Dekker (eds.) *Questions in dynamic semantics*. Amsterdam: Elsevier.

Aloni, M. and P. Egre (2010). Alternative questions and knowledge attributions. *Philosophical Quarterly*, 60: 1-27.

Asher, N. and A. Lascarides (1998). Questions in dialogue. *Linguistics and Philosophy*, 21, 237-309.

Baumann, P. (2008). Contrastivism rather than something else? On the limits of epistemic contrastivism. *Erkenntnis*, 69: 189-200.

Beebe, J. and W. Buckwalter (2010). The epistemic side-effect effect. *Mind and Language*, 25: 474-98.

Blaauw, M. (2004). *Contrastivism: reconciling skeptical doubt with ordinary knowledge.* Ph.D. dissertation, Free University of Amsterdam.

Blome-Tillman, M. (2009). Knowledge and presuppositions. *Mind*, 118: 241-94.

Boer, S. and W. Lycan (1975). Knowing who. *Philosophical Studies*, 28, 299-344.

Brogaard, B. (2009). What Mary did yesterday: Reflections on knowledge-wh. *Philosophy and Phenomenological Research*, 78: 439-67.

Brueckner, A. (1983). Transcendental arguments I. *Nous*, 17: 551-75.

Brueckner, A. (1984). Transcendental arguments II. *Nous*, 18: 197-225.

Brueckner, A. (1993). One more failed transcendental argument. *Philosophy and Phenomenological Research*, 53: 633-36.

Brueckner, A. (1996). Modest transcendental arguments. *Philosophical Perspectives*, 10: 265-80.

Buckwalter, W. (2010). Knowledge isn't closed on Saturday: A study in ordinary language. *Review of Philosophy and Psychology*, 1: 395-406.

Cappellen, H. and E. Lepore (2005). *Insensitive semantics.* Oxford: Blackwell.

Cohen, S. (1998) Contextualist solutions to epistemological problems: Skepticism, Gettier, and the lottery. *Australasian Journal of Philosophy*, 76: 289-306.

Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13: 57-89.

Cohen, S. (2001). Contextualism defended: Comments on Richard Feldman's 'Skeptical problems, contextualist solutions'. *Philosophical Studies*, 103: 87-98.

Cohen, S. (2005). Contextualism defended. In M. Steup and E. Sosa (eds.), *Contemporary debates in epistemology* (pp. 56-62). John Wiley and Sons.

Davidson, D. (1991). Three varieties of knowledge. In A. Phillips-Griffiths (ed.) *A.J. Ayer: memorial essays*. Cambridge: Cambridge University Press.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52: 913-29.

DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104: 1-52.

DeRose, K. (2009). *The case for contextualism: Knowledge, skepticism and context vol. 1*. Oxford: Oxford University Press.

DeRose, K. (2011). Contextualism, contrastivism, and X-phi surveys. *Philosophical Studies*, 156: 81-110.

Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67: 1007-1023.

Dretske, F. (1972). Contrastive statements. *The Philosophical Review*, 81: 411-437.

Dretske, F. (2003). Skepticism: what perception teaches. In S. Luper (ed.) *The skeptics*. Hampshire: Ashgate Publishing.

Fantl, J. and M. McGrath (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111: 67-94.

Feltz, A. and C. Zarpentine (2010). Do you know more when it matters less? *Philosophical Psychology*, 23: 683-706.

Fiengo, R. (2007). *Asking Questions*. Oxford: Oxford University Press.

Ginzburg, J. (1992). *Questions, queries and facts: A semantics and pragmatics for interrogatives*. Ph.D. thesis, Stanford University Department of Linguistics.

Ginzburg, J. (1995). Resolving questions, I. *Linguistics and Philosophy*, 18, 459-527.

Ginzburg, J. (2005). Abstraction and ontology: Questions as propositional abstracts in type theory with records. *Journal of Logic and Computation*, 15: 113-30.

Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73: 771-91.

Gram, M. (1971). Transcendental arguments. *Nous*, 5:15-26.

Gram, M. (1977). Must we revisit transcendental arguments? *Philosophical Studies*, 31: 624-26.

Groenendijk, J. and M. Stokhof (1984). *Studies in the semantics of questions and the pragmatics of answers.* PhD. Dissertation, University of Amsterdam.

Groenendijk, J. and M. Stokhof (1997). Questions. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language* (pp. 1055-1124). Amsterdam.

Hagstrom, P. (2003). What questions mean. *GLOT International* 7: 188-201.

Hamblin, C.L. (1958). Questions. *Australasian Journal of Philosophy*, 36: 159-68.

Hamblin, C.L. (1973). Questions in Montague English. *Foundations of Language,* 10: 41-53.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.

Hawthorne, J. (2005). The case for closure. In M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary Debates in Epistemology* (pp. 40-56). John Wiley and Sons.

Heim, I. (1994). Interrogative semantics and Karttunen's semantics for *know*. In R. Buchalla and A. Mittwoch (eds.) *Proceedings of IATL 1*. Jerusalem: Akademon.

Heller, M. (1999). Relevant alternatives and closure. *Australasian Journal of Philosophy*, 77: 196-208.

Heller, M. (1999). The proper role for contextualism in an anti-luck epistemology. *Philosophical Perspectives*, 13: 115-29.

Hookway, C. (1996). Questions of Context. *Proceedings of the Aristotelian Society*, 96: 1-16.

Johnsen, B. (2001). Contextualist swords, skeptical plowshares. *Philosophy and Phenomenological Research*, 62: 385-406.

Kallestrup, J. (2009). Knowledge-wh and the problem of convergent knowledge. *Philosophy and Phenomenological Research*, 78: 468-76.

Kant, I. (1787). *Critique of Pure Reason* (trans. P. Guyer and A. Wood). Cambridge: Cambridge University Press, 1997.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds.),

Themes from Kaplan (pp. 481-563). Oxford: Oxford University Press.

Karttunen, L. (1977). Syntax and semantics of questions. *Linguistics and Philosophy*, 1: 3-44.

Kelp, C. (2011). A problem for contrastivist accounts of knowledge. *Philosophical Studies*, 152: 287-92.

Korner, S. (1966). Transcendental tendencies in recent philosophy. *Journal of Philosophy*, 63: 551-61.

Korner, S. (1967). The impossibility of transcendental deductions. *Monist*, 51: 317-31.

Krifka, M. (2001). For a structured meaning account of questions and answers. In C. Fery and W. Sternefeld (eds.) *Audiatur vox sapientia: A festschrift for Arnim von Stechow*. Berlin: Akademie Verlag.

Kvanvig, J. (2007). Contextualism, contrastivism, relevant alternatives, and closure. *Philosophical Studies*, 134: 131-40.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8: 339-59.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74: 549-67.

Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In G. Preyer and G. Peter (eds.), *Contextualism in philosophy: Knowledge, meaning and truth*, (pp. 11-

50). Oxford: Oxford University Press.

Luper, S. (2006). Dretske on knowledge closure. *Australasian Journal of Philosophy*, 84: 379-94.

MacFarlane, J. (2007). The assessment sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1: 197-233.

May, J., W. Sinnott-Armstrong, J. Hull, and A. Zimmerman (2010). Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1: 265-73.

Morton, A. and A. Karjalainen (2003). Contrastive knowledge. *Philosophical Explorations*, 6: 74-89.

Neta, R. (2008). Undermining the case for contrastivism. *Social Epistemology*, 22: 257-70.

Nozick, R. (1981). *Philosophical explanations*. Cambridge: Harvard University Press.

Oakley, I.T. (1976). An argument for skepticism concerning justified beliefs. *American Philosophical Quarterly*, 13: 221-28.

Pritchard, D. (2008). Contrastivism, evidence, and skepticism. *Social Epistemology*, 22: 305-23.

Putnam, H. (1981). *Reason, truth and history*. Cambridge: Cambridge University Press.

Recanati, F. (1986). On defining communicative intentions. *Mind and Language* 1: 213-42.

Reinhart, T. (1998). Wh-in-situ in the framework of the minimalist program. *Natural Language Semantics*. 6: 29-56.

Richard, M. (2004). Contextualism and relativism. *Philosophical Studies*, 119: 215-42.

Roberts, C. (1996). Information structure: towards an integrated theory of formal pragmatics. In J. Yoon and A. Kathol (Eds.), *OSU Working Papers in Linguistics, Volume 49: Papers in Semantics* (pp. 91-136). The Ohio State University Department of Linguistics.

Roberts, C. (2004). Context in dynamic interpretation. In L. Horn and G. Ward (Eds.), *Handbook of Contemporary Pragmatic Theory*, (pp. 197-220). Oxford: Blackwell.

Rooth, M. (1985). *Association with focus.* Ph.D. dissertation, UMass Amherst.

Rullmann, H. and S. Beck (1998a). Reconstruction and the interpretation of *which*-phrases. In G. Katz, S. Kim, and H. Winhart (eds.) *Reconstruction: Proceedings of the 1997 Tubingen workshop.* University of Tubingen.

Rullmann, H. and S. Beck (1998b). Presupposition projection and the interpretation of *which*-questions. In D. Stolovitch and A. Lawson (eds.) *Proceedings of SALT 8.* Ithaca: Cornell University Department of Linguistics.

Sauerland, U. and F. Heck (2003). LF-intervention effects in pied-piping. In M. Kadowaki and S. Kawahara (eds.) *Proceedings of the northeast linguistics society*. Amherst, MA: GLSA Publications.

Schiffer, S. (1996). Contextualist solutions to skepticism. *Proceedings of the Aristotelian Society*, 96: 317-33.

Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, 119: 73-103.

Schaffer, J. (2005a). Contrastive knowledge. In T. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology, Volume 1* (pp. 235-271). Oxford: Oxford University Press.

Schaffer, J. (2005b). What shifts? Thresholds, standards, or alternatives? In G. Preyer and G. Peter (Eds.), *Contextualism in Epistemology*, (pp. 115-130). Oxford: Oxford University Press.

Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75, 383-403.

Schaffer, J. (2007). Closure, contrast and answer. *Philosophical Studies*, 133: 233-55.

Schaffer, J. (2009). Knowing the answer redux: Replies to Brogaard and Kallestrup. *Philosophy and Phenomenological Research*, 78: 477-500

Schaffer, J. (2012). Contrastive knowledge: Reply to Baumann. In A. Tolksdorf (ed.) *The concept of knowledge*. DeGruyter.

Schaffer, J. and J. Knobe (2012). Contrastive knowledge surveyed. *Nous*, 46: 675-708.

Shoemaker, S. (1963). *Self-knowledge and self-identity*. Ithaca: Cornell University Press.

Sinnott-Armstrong, W. (2004). *Pyrrhonian skepticism*. Oxford: Oxford University Press.

Stalnaker, R. (1999). *Context and content*. Oxford: Oxford University Press.

Stalnaker, R. (2004). Comments on 'From contextualism to contrastivism'. *Philosophical Studies*, 119: 105-17.

Steglich-Petersen, A. (2014). Knowing the answer to a loaded question. *Theoria*, 80 (1).

Stern, R. (2000). *Transcendental arguments: Answering the question of justification*. Oxford: Oxford University Press.

Stine, G. (1976). Skepticism, relevant alternatives and deductive closure. *Philosophical Studies*, 29: 249-61.

Strawson, P. (1959). *Individuals*: *An essay in descriptive metaphysics*. London: Methuen

Strawson, P. (1966) *The Bounds of Sense: An essay on Kant's 'Critique of pure reason'*. London: Methuen.

Strawson, P. (1985) *Skepticism and naturalism: Some varieties*. London: Methuen.

Stroud, B. (1968). Transcendental arguments. *Journal of Philosophy*. 65: 241-56.

Stroud, B. (1994). Kantian arguments, conceptual capacities, and invulnerability. In P. Parrini (ed.) *Kant and contemporary epistemology*, (pp. 155-76).

Dordrecht: Kluwer.

Stroud, B. (1999). The goal of transcendental arguments. In R. Stern (ed.) *Transcendental arguments: problems and prospects*, (pp. 203-23). Oxford: Oxford University Press.

Unger, P. (1975). *Ignorance*. Oxford: Oxford University Press.

Van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26, 727-763.

Von Stechow, A. (2000) Some remarks on choice-functions and LF-movement. In K. von Huesinger and U. Egli (eds.) *Reference and anaphoric relations*. Dordrecht: Kluwer.

Williamson, T. (2002) *Knowledge and Its Limits*. Oxford: Oxford University Press.

Vita

Jason M. Rourke

jmrourke@syr.edu

Education:

B.A., University of Scranton, 1996

M.S., Syracuse University, 2010

Publications:

Rourke, J. (2013). A counterexample to the contrastive account of knowledge. *Philosophical Studies*, 162: 637-643.

Paper Presentations:

"Prosody and Implicature" presented to the Syracuse University Language and Linguisitics Colloquium Series, March 2010.