Georgia State University

ScholarWorks @ Georgia State University

Content and Contrastive Self-Knowledge

Vincent G. Abruzzo
Georgia State University

Follow this and additional works at: https://scholarworks.gsu.edu/philosophy_theses

CONTENT AND CONTRASTIVE SELF-KNOWLEDGE

by

VINCENT ABRUZZO

Under the Direction of Andrea Scarantino

ABSTRACT

It is widely believed that we have immediate, introspective access to the content of our own thoughts. This access is assumed to be privileged in a way that our access to the thought content of others is not. It is also widely believed that, in many cases, thought content is individuated according to properties that are external to the thinker's head. I will refer to these theses as *privileged access* and *content externalism*, respectively. Though both are widely held to be true, various arguments have been put forth to the effect that they are incompatible. This charge of incompatibilism has been met with a variety of compatibilist responses, each of which has received its own share of criticism. In this thesis I will argue that a contrastive account of self-knowledge is a novel compatibilist response that shows significant promise.

INDEX WORDS: Content externalism, Anti-individualism, Contrastivism, Contextualism, Skepticism, Self-knowledge, Privileged access, A priori knowledge, McKinsey paradox, Epistemic closure, Warrant transmission failure

CONTENT AND CONTRASTIVE SELF-KNOWLEDGE

by

VINCENT ABRUZZO

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2012

CONTENT AND CONTRASTIVE SELF-KNOWLEDGE

by

VINCENT ABRUZZO

Committee Chair: Andrea Scarantino

Committee: Neil Van Leeuwen

Daniel Weiskopf

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

August 2012

DEDICATION

This thesis is dedicated to the memory of Sergeant Michael Kashkoush.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

1 INTRODUCTION

It is widely believed that we have immediate, introspective access to the content of our own thoughts. This access is assumed to be privileged in a way that our access to the thought content of others is not. It is also widely believed that, in many cases, thought content is individuated according to properties that are external to the thinker's head. I will refer to these theses as *privileged access* and *content externalism*, respectively. Though both are widely held to be true, various arguments have been put forth to the effect that they are incompatible. This charge of incompatibilism has been met with a variety of compatibilist responses, each of which has received its own share of criticism. In this thesis, I will present a novel compatibilist response that shows significant promise.

I will argue that a contrastive account of self-knowledge helps to clarify what we do and do not have privileged access to in regards to the content of our thoughts. In chapter 2, I will briefly discuss content externalism and privileged access. I will then present an argument to the effect that content externalism and privileged access are incompatible, namely, the McKinsey paradox. I will also briefly discuss various proposed solutions to the McKinsey paradox. In chapter 3, I will introduce epistemic contrastivism, distinguish it from epistemic contextualism, and offer several reasons why contrastivism is the superior theory. In chapter 4, I will present my argument that a contrastive account of self-knowledge solves the McKinsey paradox and that it does so more persuasively than the warrant transmission failure solution, which is currently the most prominent solution to the paradox. I will conclude by considering some of the philosophical consequences of a contrastive account of self-knowledge and privileged access.

2   THE MCKINSEY PARADOX

It has been argued that content externalism and privileged access are incompatible. Martin Davies (1998) has identified two types of incompatibilist arguments, one unveiling what he calls the *achievement problem* and another unveiling the *consequence problem*. An example of the former is given by Paul Boghossian's (1989) argument to the effect that if content externalism is true, then we can't *achieve* the sort of privileged access to our self-knowledge that we think we have, hence the achievement problem. An example of the latter is given by Michael McKinsey's (1991) argument to the effect that the conjunction of content externalism and privileged access has an absurd *consequence*, hence the consequence problem. From here on, I will focus on the consequence problem. As I will suggest in the conclusion, a contrastivist solution to the consequence problem can also help solve the achievement problem.

I will now discuss content externalism and privileged access in order to explicate the most prominent form of the consequence problem: the McKinsey paradox.

2.1   Content Externalism

In a seminal paper, Hilary Putnam writes that "'meanings' just ain't in the *head*" (1975, 227). Putnam arrives at this conclusion by considering his now famous Twin Earth thought experiment. We are asked to consider two physically identical individuals, Oscar and Toscar, who both utter the phrase 'water is wet.' Oscar utters the phrase on Earth, where the clear liquid called 'water' is comprised of H_2O molecules. Toscar utters the phrase on Twin Earth, where the clear liquid called 'water' is comprised of XYZ molecules. According to Putnam, the word 'water' means H_2O on Earth and XYZ on Twin Earth. But since Oscar and Toscar are physical duplicates, meaning must be determined in part by environmental factors. The focus of Putnam's argument is on the meaning

of words, hence he establishes *semantic externalism*. Colin McGinn (1977) argues that it is natural to extend this line of reasoning to thought contents. Oscar and Toscar both think to themselves and believe that 'water is wet.' Because the word 'water' means something different to Oscar and Toscar, their thoughts and beliefs are about different things. That is, the content of their thoughts is determined in part by environmental factors. As McGinn writes, "in specifying their relational beliefs we must, if we are to report adequately, refer these beliefs to the substances causal interaction with which makes them of the substances they are. We should not let this elementary observation be obscured by the fact that they suppose their mental states to be identical" (531).[1] And this holds for all propositional attitudes including beliefs, desires, hopes, etc. This form of content externalism is commonly referred to as *natural kind externalism*.

Further extending the scope of the argument, Tyler Burge (1979) has argued that the individuation of mental content depends on the practices of the linguistic community within which the individual is embedded. The first step in Burge's argument is to suppose that Oscar goes to his doctor and tells her that he thinks that he has arthritis in his thigh. The doctor informs Oscar that he can't have arthritis in his thigh because arthritis is a rheumatoid ailment of the joints only. Arthritis cannot spread to muscles. Oscar realizes his misuse of the word 'arthritis' and learns something about the concept *arthritis*.

The second step in Burge's argument is to imagine a counterfactual scenario in which a physically identical Oscar goes to his doctor and tells her that he has arthritis in his thigh. In this counterfactual scenario, Oscar's physical history and brain states are exactly the same as in the actual scenario. The only difference is that, in the counterfactual scenario, the linguistic community within which Oscar is embedded uses the word 'arthritis' to refer to rheumatoid ailments of both joints and mus-

---

[1] McGinn (1977) anticipates the problem that content externalism poses to privileged access. He writes in a footnote to this remark that, "this point shows what is wrong with certain doctrines of 'privileged access.' Since we are often less well placed to individuate our own mental states than another is, it can hardly be maintained that we enjoy authoritative access to our own mental descriptions" (531).

cles. According to Burge, in the actual scenario Oscar has a false belief about the concept *arthritis* while in the counterfactual scenario Oscar has a true belief about the concept *tarthritis (Twin Earth arthritis)*. Since the actual and counterfactual Oscars are historically and physically identical, it must be some other factor that determines the different mental contents that they have. This difference is the linguistic communities that they belong to. Hence, the individuation of mental content is determined in part by factors outside the thinker's head, in this case socio-linguistic factors. This form of content externalism is commonly referred to as *social externalism*.

To sum up, the general thesis underlying all forms of content externalism is that many of our thoughts imply some external condition E. In the case of a thought about *water*, the external condition implied is that the thinker has interacted with water (H_2O).[2] In the case of a thought about *arthritis*, the external condition implied is that the thinker is embedded within a linguistic community that has the concept *arthritis*. Content externalism is widely accepted and I will assume, for the purposes of getting the McKinsey paradox off the ground, that it is true. This will allow us to examine the supposed incompatibility of it and privileged access. It is to this latter thesis that I now turn.

2.2   Privileged Access to Self-Knowledge

It is widely accepted that we have privileged access to the content of our own thoughts. We can illustrate this point by comparing third-person thought ascriptions to first-person thought ascriptions. If Oscar wants to know what Mary is thinking, then he must infer the content of her thoughts based on his experiences of her actions, her speech, or third-person testimony of her thoughts, actions or speech. If Mary schedules a doctor's appointment with an arthritis specialist to address the pain in her thigh or says 'I think I have arthritis in my thigh,' then Oscar may infer that Mary has the thought 'I think I have arthritis in my thigh.' This is not the case when Oscar wants to

---

[2] This does not hold for all thought contents, e.g., thoughts about kryptonite.

know the content of his own thoughts. Oscar does not need to observe his own actions or speech in order to infer what he himself is thinking. Rather, he is in a position to gain immediate knowledge of his own thoughts as he thinks them. That is, he has privileged access to his thought contents.

Philosophers often characterize this sort of self-knowledge as *a priori* in the sense that it is prior to any particular experience of the world outside the thinker's head. When Oscar thinks the thought 'I have arthritis in my thigh,' he is immediately aware of the thought; it is not necessary to infer from external world experience the content of his own thought. It is important to note that this use of the term *a priori* does not indicate that the knowledge is knowable without relying on any evidence from experience. Rather, the term in this context is meant to indicate that the knowledge is knowable via introspection alone.

This construal of privileged access sidesteps the question of whether or not introspection itself is associated with experience. Our knowledge of our thoughts is *a priori* insofar as this knowledge is not based on experience *of the external world*, regardless of whether or not introspection is associated with a distinctive phenomenology. As Jessica Brown (2004) has pointed out, this notion of privileged access is modest in its claims and is often assumed in the literature. This moderate construal of the *a priori,* privileged access that we have of our thought content is sufficient to elucidate the problem that it poses to content externalism.

2.3   The Incompatibility of Content Externalism and Privileged Access

Both content externalism and privileged access to content are widely accepted. However, some have argued that these two theses are incompatible. One version of the problem is referred to as the McKinsey paradox. The McKinsey paradox states that three commonly held theses form an inconsistent triad. The three theses are:

(1) individuals have *a priori* knowledge of the contents of their own thoughts (privileged

access)

(2) mental content is determined in part by factors outside the thinker's head (content externalism)

(3) it is not possible to gain *a priori* knowledge of the external world

I have already discussed theses (1) and (2). Thesis (3) is the intuitive position that we cannot gain knowledge of the external world without experiencing or inquiring about whatever it is we are attempting to gain knowledge about. If I want to know whether or not there is water in my environment, for instance, I have to look around to see if there is any water, as picked out by a set of stereotypical 'watery' properties (being transparent, being found in lakes, being potable, etc.). I also have to make sure that the liquid endowed with these stereotypical properties is H_2O and not XYZ or some other chemical compound.

Michael McKinsey (1991) argues that even though these three theses are commonly held to be true, they are in fact inconsistent. Specifically, the conjunction of (1) and (2) implies the negation of (3). According to (1), the instantiation of specific mental content implies that specific environmental factors obtain, since mental content is individuated according to those factors. But according to (2), we come to know the content of our thoughts *a priori*. This, according to the McKinsey paradox, has the consequence that we can know *a priori* the environmental factors responsible for the individuation of the content of the thoughts that we have. But gaining knowledge of specific environmental factors *a priori* is the negation of thesis (3), according to which we can't gain knowledge of the external world *a priori*.

McKinsey considers the concrete case of Oscar and the following three propositions:

"(1) Oscar knows *a priori* that he is thinking that water is wet.

(2) The proposition that Oscar is thinking that water is wet necessarily depends upon E.

(3) The proposition E cannot be known *a priori*, but only by empirical investigation" (1991, 12).

Proposition (1) is entailed by privileged access. When Oscar thinks to himself that 'water is wet,' he has immediate, introspective access to that thought. Proposition (2) is entailed by content externalism. The proposition E that McKinsey refers to is the environmental condition that individuates the content of the thought as specified by content externalism. In the case of a thought about *water*, the external condition implied is that the thinker has interacted with water (H_2O). McKinsey argues that, according to content externalism, a thought about *water* implies this external condition.

The conjunction of propositions (1) and (2) implies the negation of proposition (3). If Oscar can know the content of his own thought *a priori*, and if he knows that thinking that thought implies proposition E, then he is in a position to know proposition E *a priori*. But that is the negation of proposition (3).

Two of the most recent articulations of the McKinsey paradox differ little from McKinsey's original formulation as I have just presented it. For example, Crispin Wright presents the following line of reasoning:

"(A) I believe that water is wet

(A+) If I believe that water is wet, then I or others in my speech community have had a history of interaction with water.

Hence

(B) I or others in my speech community have had a history of interaction with water" (2011, 81).

Jesper Kallestrup (2011) presents a general form of the McKinsey paradox and another version of the paradox involving thoughts about water. The general form of the paradox is:

"(1) S has mental property M;

(2) S meets non-mental condition C if she has mental property M;

(3) S meets non-mental condition C" (157).

Kallestrup's specific version of the paradox involving a thought about water is as follows:

"(1) S is thinking thoughts containing the atomic, natural kind concept water;

(2) S has causally interacted with water if S is thinking thoughts containing the atomic, natural kind concept water;

(3) S has causally interacted with water" (158).

In *all* of these more recent articulations of the McKinsey paradox the reasoning is the same. If the statements in (1) and (2) are knowable *a priori*, then (3) is knowable *a priori* as well. But, according to all three writers, (3) is a fact about the external world that is not knowable *a priori*, so something has to give.

2.4   Exploring the Solution Space

There are various proposed solutions to the McKinsey paradox. I will discuss three solutions that I will refer to as the externalist solution, the closure denial solution, and the warrant transmission failure solution. I will discuss each in turn.

2.4.1   The Externalist Solution

One of the earliest direct responses to McKinsey is due to Anthony Brueckner (1992). Brueckner claims that McKinsey misinterprets the commitments of content externalism. He correctly points out that McKinsey interprets content externalism as a thesis that holds that some of our thoughts imply facts about the world, such as that there are natural kinds present in the environment or that the thinker is embedded in a certain social community. However, he believes that McKinsey incorrectly interprets this implication as *conceptual* implication. For example, Oscar's thought that 'Fred is a bachelor' conceptually implies that Oscar is thinking that 'Fred is an unmarried man.' But, according to Brueckner, the correct interpretation of the externalist's commitments is that some of

our thoughts *metaphysically* imply facts about the world. Furthermore, if the sort of metaphysical implication that characterizes externalism is true, then it doesn't pose a threat to privileged access.

McKinsey (1994) agrees with Brueckner's analysis of the paradox but argues that Brueckner has misunderstood his aim in his original paper. McKinsey did not just assume that externalism is a thesis regarding the conceptual implications regarding thoughts about the world. Rather, he argues that it is a thesis regarding either the conceptual implications or the metaphysical implications of thoughts about the world. If it is the latter, then the thesis is trivialized. As McKinsey writes:

> If the anti-individualist adopts a sense of 'wide state' according to which metaphysical entailment of external objects is sufficient for a state to be wide, then it turns out that absolutely *every* psychological state will be wide, no matter how 'narrow' that state might intuitively have appeared (126).

McKinsey refers here to *anti-individualism*, another commonly used term for content externalism. Content externalism is a thesis about the properties of some of our thoughts; it is a thesis explaining the special wide properties of thoughts about natural kinds or socially dependent concepts, for example. However, if content externalism is a thesis about the *metaphysical* implications of those thoughts, then it is not really saying anything about those types of thoughts. This is the case because every mental state we have, given some commonly held materialist assumptions, has some sort of metaphysical implication. As McKinsey points out, every thinker's existence metaphysically implies the existence of his or her parents, which is to say that every thought metaphysically implies the existence of the thinker's parents. But if all thoughts have metaphysical implications, then content externalism isn't saying anything additional about the types of thoughts that the thesis purports to make interesting claims about.

To sum up: the externalist solution holds that if content externalism is understood as a thesis regarding the metaphysical implications of certain thoughts, rather than the conceptual implications of those thoughts, then it is not a threat to privileged access. However, if it is understood as a thesis regarding the metaphysical implications of certain thoughts, then it is not saying anything interesting

about those thoughts, because *all* thoughts have such implications. Therefore, we should understand externalism as a thesis regarding the conceptual implications of certain thoughts, and thus as a threat to privileged access as exemplified by the McKinsey paradox.

2.4.2 The Closure Denial Solution

Another possible solution to the McKinsey paradox is to deny epistemic closure (hereafter closure).[3] Closure maintains that if *s* knows that *p*, and if *s* knows that *p* entails *q*, then *s* is in a position to know that *q*. Dretske's (1970) zebra case illustrates the motivation behind denying closure. If you walk up to a pen at the zoo and see a zebra, then it seems like you know that you're standing in front of a zebra. And it also seems that if you know that you're standing in front of a zebra, then you know that it is not the case that you're standing in front of a cleverly disguised mule. If closure holds, then you are at least in a position to conclude, based on your knowledge of the conditional just stated and the antecedent of that conditional, that you are not standing in front of a cleverly disguised mule. However, it does not seem like you have enough evidence to know that you're not standing in front of a cleverly disguised mule. Dretske's response to this dilemma is to deny closure. Denying closure allows one to maintain all of the seemingly plausible premises of the argument while avoiding the counterintuitive conclusion.

Similarly, denying closure avoids the absurd conclusion of the McKinsey paradox. The form of the argument underlying the McKinsey paradox is the same as the argument regarding Dretske's zebra: they are both a modus ponens. If we deny closure, then it is possible for an individual to know the conditional premise of the McKinsey paradox and the antecedent of that conditional *without* being in a position to know the consequent of the conditional. That is, if we deny closure, then

---

[3] No one has proposed this solution to the McKinsey paradox itself. However, it is often mentioned as a possible solution to philosophical problems of the same form in which *plausible* premises imply an *implausible* conclusion of a valid argument (for example, Dretske 1970).

we cannot arrive at the conclusion of the McKinsey paradox, that we can gain knowledge of the external world prior to experiencing it, based on our knowledge of the premises of the argument underlying it. Those premises being privileged access and content externalism.

However, I do not think that we should deny closure if we do not have to. It is a drastic solution to be chosen only if other, less drastic solutions are not available. One of the virtues of contrastivism is that it enables us to avoid denying closure in zebra-like cases. However, contrastive closure looks different than the current closure principle under discussion, due to the third argument of the knowledge relation. There is also still much work needed in constructing a problem-free contrastive closure principle and the debate is currently ongoing.[4]

2.4.3   The Warrant Transmission Failure Solution

The most prominent response to the McKinsey paradox is to argue that the *warrant* for the premises fails to transmit to the conclusion. The warrant for a proposition is the evidence in its favor. For example, my warrant for my claim that 'the sky is blue' is my visual experience of the blue sky. This response is known as the warrant transmission failure response and it has been put forth most forcefully by Martin Davies (1998, 2000) and Crispin Wright (2000, 2003, 2011). I will briefly provide an example of both warrant transmission success and warrant transmission failure.

Suppose that I see a glass of water and proclaim that 'there is a glass of water on the table.' My warrant for this claim is my perceptual experience of the glass of clear liquid on the table. Suppose that I also proclaim that 'if the glass is filled with water, then it's not filled with orange juice.' I then conclude that 'the glass on the table is not filled with orange juice.' This is a case in which the warrant for the premises successfully transmits to the conclusion. That is, my warrant for the claim that the glass is not filled with orange juice is my perceptual experience of the clear liquid in the

---

[4] See Schaffer (2007).

glass. I can conclude that there is no orange juice on the table *based on my warrant for the premises of the argument.*

Now suppose that I again see a glass of water and proclaim that 'there is a glass of water on the table.' But suppose that this time I proclaim that 'if the glass is filled with water, then it's not filled with vodka.' It seems that I can conclude that 'the glass is not filled with vodka.' However, my warrant for the claim that the glass is filled with water is my perceptual experience of the clear liquid on the table. Since vodka is also a clear liquid, the warrant for my claim that the glass is filled with water does not transmit to the conclusion that it is not vodka, since it does not rule out the possibility that it is vodka. There is no warrant for the conclusion, even though the argument is valid and I have warrant for the premises. That is, the warrant fails to transmit from the premises to the conclusion of the argument.

Warrant transmission occurs when the warrant for the premises of an argument warrants a belief in the conclusion of the argument as seen in the water/no orange juice example. Warrant transmission failure occurs when there is no warrant for belief in the conclusion of an argument whose premises are warranted. Different reasons are given as to why warrant sometimes fails to transmit; for example, there might be an implicit instance of question begging. For example, in the water/vodka case just discussed, in order to conclude that the liquid is not vodka based on the warrant for the premise that the liquid is water, we must be able to antecedently exclude the possibility that it is not vodka, since the fact that the liquid is clear is consistent with it being either water or vodka.

Warrant transmission failure is closely related to denying closure, but it is a weaker claim. As Brown (2004) points out, denying closure entails warrant transmission failure, but warrant transmission failure does not entail the denial of closure. A case of warrant transmission failure is not necessarily a counterexample to closure because it might be the case that anyone who knows the premises

of the argument still knows its conclusion for independent reasons, just not in virtue of the warrant for the premises of the argument. Davies and Wright have applied warrant transmission failure more generally to cases such as Dretske's zebra case as well as the McKinsey paradox. Regarding the latter, the claim is that the warrant for the belief that Oscar is thinking about water does not transmit to the conclusion that Oscar is in an environment with water because that warrant is compatible with Oscar being in an environment with twin water. The warrant fails to transmit in the argument underlying the paradox.

After I have explained the contrastivist solution to the McKinsey paradox I will compare it to the Warrant Transmission Failure solution. I believe that the contrastivist solution is based on many of the same principles but that contrastivism does a better job of *explaining why* we do and don't know what it seems like we do and don't know.

To sum up, each of the three solutions I have considered so far has flaws. We can scale back externalism and make it compatible with privileged access, but then we trivialize the thesis and make it philosophically uninteresting. We can deny closure, but this seems like a drastic solution and we should avoid it if we can. Warrant transmission failure is a promising solution, but if another solution shares its insights *and* is explanatorily superior then we ought to prefer it. I believe that a contrastive solution shares the insights of warrant transmission failure and provides a clearer account of what exactly we have privileged access to. It is not the case that we either have an absurd access to facts about the external world or that we don't have any privileged access to the content of our own thoughts. Rather, we have privileged access to some, but not all, of the aspects of our thoughts.

I will now introduce contrastivism, and then show how it can be used to solve McKinsey's paradox.

3   CONTRASTIVISM

According to the traditional view, knowledge is a *two-place* relation that takes the form K*sp*, where subject *s* knows that proposition *p* is the case. According to the contrastivist, knowledge is a *three-place* relation that takes the form K*spq*, where *s* knows that *p* rather than the contrast set *q*. This theory has its roots in Fred Dretske's relevant alternatives view of knowledge, according to which, "to know that *x* is *A* is to know that *x* is *A* within a framework of relevant alternatives, *B, C,* and *D*. This set of contrasts, together with the fact that *x* is *A*, serves to define what it is that is known when one knows that *x* is *A*" (1970, 1022). That is, when we say that *s* knows that *p*, we are implicitly asserting that *s* knows that *p* rather than some contrast set *q*.[5]

Contrastivism also shares an important insight with contextualism, the theory that the truth conditions of knowledge ascriptions can vary from context to context. For example, consider the case of Moore and his knowledge claim that he has hands. Under normal circumstances, our pre-theoretic intuitions lead us to ascribe to Moore the knowledge that he has hands when he raises his hands and says that 'I have hands.' In virtue of their general acceptance, knowledge ascriptions under circumstances such as these are what Jonathan Schaffer (2004) refers to as *dogmatic* ascriptions. However, after considering the possibility that Moore is just a brain in a vat, we might conclude with the *skeptic* that Moore does not know that he has hands. Moore's knowledge that he has hands seems to vary depending on the context. When at home and not philosophizing, he knows that he has hands; when in the philosophy classroom, he does not know that he has hands. Since we can take either the dogmatic or skeptical position, we are left with the contradictory knowledge claims K*mh* and ~K*mh*, where *m* is Moore and *h* is the proposition that he 'has hands.' The contextualist explanation of this problem is that we make different knowledge ascriptions in different contexts. The

---

[5] I will often use a single variable to specify a contrast set but it should be kept in mind that the singular term is meant to denote a set of contrasts. For example, if we say that Moore knows that he has hands rather than stumps, the contrast can be construed as a set of alternatives including: stumps, hooks, fins, etc.

contrastivist explanation is that there is an implicit third argument in the knowledge relation and that this third argument can vary from context to context.

According to the contrastivist, when we ascribe to Moore the knowledge that he has hands we are actually ascribing to him the knowledge that he has hands rather than stumps (or hooks, fins, etc.). This is the case because Moore can perceptually distinguish between hands and stumps. That is, when he raises his arm, he sees a hand and not a stump. In this case there is an immediate percep- tual justification. James Pryor characterizes this type of justification as that which, "whenever you have an experience as of p, you thereby have immediate *prima facie* justification for believing p" (2000, 536). However, when we deny to Moore the knowledge that he has hands in light of the skeptical challenge, what we are asserting is that it is not the case that Moore knows that he has hands rather than vat-images of hands. The possibility of being a brain in a vat and looking at vat-images of hands is perceptually indistinguishable from looking at real hands. Because Moore can't distinguish between the two, he can't know whether he has actual hands or vat-images of hands. This is the force of the skeptical position. Formalizing these claims contrastively we get Kmhs and ~Kmhv, where s is stumps and v is vat-images of hands. According to contrastivism, there is no contradiction because Moore's knowledge that he has hands rather than stumps is consistent with his not knowing that he has hands rather than vat-images of hands. Whereas the contextualist purports to explain our contradictory knowledge ascriptions in cases such as these, the contrastivist eliminates the contradic- tion altogether.

Schaffer (2004) has made a compelling case that all of our knowledge ascriptions are actually contrastive, and that contrastivism is superior to contextualism. Schaffer has argued on linguistic grounds that the term 'knows' is lexically ternary and that for this reason, "all knowledge ascriptions contain a syntactically real contrast variable q in their logical forms" (77). He and Joshua Knobe (forthcoming) have also conducted experimental philosophy surveys of people's intuitions regarding

knowledge claims in classic cases that have been previously used to support a contextualist model of knowledge ascriptions. They conclude that the results support a contrastive model of knowledge ascriptions and are in fact inconsistent with a contextualist model of knowledge ascriptions.[6] Taken together, the linguistic and empirical data makes a strong case that knowledge is a ternary relation.[7]

Furthermore, Schaffer argues that contrastivism is superior to contextualism on two further grounds. This first difference is that contextualism is characterized by *equivocationism* whereas contrastivism is characterized by *compatibilism*. Both terms denote strategies that are, according to Schaffer, "distinct philosophical accounts of how dogmatic knowledge is shielded from skeptical doubt" (2004, 82). According to contextualist equivocationism, in the case of Moore we *equivocate* between the dogmatic attribution that Moore knows he has hands and the skeptical denial of that knowledge. According to contrastivist compatibilism, it is *compatible* to hold that Moore knows that he has hands and that he does not know that he has hands. This is because in the former case the contrast is with stumps and in the latter case the contrast is with vat-images of hands. These positions, writes Schaffer, have different consequences for epistemic closure.

As we have seen, according to the principle of epistemic closure, given p and $p \rightarrow q$, q follows. The equivocationism characteristic of contextualism, combined with this principle, leads to what Schaffer calls *immodest knowledge claims* and *immodest skeptical denials of knowledge*, neither of which are acceptable. For example, according to the dogmatist, Moore knows that he has hands. Additionally, if Moore knows that he has hands, then he knows that he is not a brain in a vat. But knowledge that he is not a brain in a vat is contrary to the skeptical possibility that he is a brain in a vat. That is, in a context in which the dogmatic position is assumed, we can gain the immodest knowledge of not being a brain in a vat. But when the skeptical position is assumed, we can never gain any knowledge

---

[6] Schaffer and Knobe (forthcoming) found that people do not attribute knowledge based on the stakes of the attribution, as contextualists have presupposed. Rather, knowledge attributions are contrast sensitive. For a critical discussion of the X-Phi results referenced see DeRose (2011).

[7] For a critical discussion of the case for contrastivism see Neta (2008).

of having hands, and this is an immodest ignorance. As Schaffer writes, "the equivocationist is forced to swing from a manic dogmatism to a depressive skepticism" (2004, 91).

On the other hand, the compatibilism that characterizes contrastivism enables us to avoid the immodest claims that follow from the deductive closure principle. Moore's knowledge that he has hands rather than stumps does not entitle him to make the claim that he knows that he has hands rather than vat-images of hands. And the denial that he does not know that he has hands rather than vat-images of hands does not entail that he does not know that he has hands rather than stumps. A modest dogmatism and a modest skepticism end up being felicitously compatible with one another.

The second difference between contextualism and contrastivism is that the former is characterized by relevance whereas the latter is characterized by saturation. As Schaffer writes, "these are distinct linguistic mechanisms for factoring alternatives into the truth-conditions, which differ as follows. By relevance, alternatives enter into the truth-conditions via the semantical rule of relevant alternatives (as triggered by the 'knows' indexical). Whereas by saturation, alternatives enter into the truth-conditions via such mechanisms as 'rather than'-arguments, interrogatives, focusing, and free variables (as serving to saturate the q slot)" (2004, 87).

As I will argue, contrastivism also achieves this sort of modesty with regard to the McKinsey paradox. To foreshadow my central claim, the immodest knowledge claim at the root of the McKinsey paradox is that we can gain *a priori* knowledge of the external world. This immodesty follows from privileged access and content externalism, so long as we assume the epistemic closure principle. As in the case of Moore and his knowledge claim that he has hands, if we make explicit the contrasts in each knowledge ascription of the McKinsey paradox, we can avoid the immodest, *a priori* knowledge claim regarding the external world.

To summarize the case for contrastivism: there is evidence, both linguistic and experimental, that supports the conclusion that the contrastive model of knowledge ascriptions is the model best suited to explain how we actually make knowledge ascriptions. And, though it is similar to contextualism, there are a variety of reasons to prefer it instead. In addition to the virtues of contrastivism that I have laid out in this section, I will now explain in some detail how it can help us solve the McKinsey paradox.

4   A CONTRASTIVIST SOLUTION TO THE MCKINSEY PARADOX

I believe that a contrastive account of self-knowledge can solve the McKinsey paradox. It does so by specifying exactly what privileged access entails. Roughly, Oscar can know that his thought is about *water* rather than some subjectively distinguishable thought content; he cannot know that his thought is about *water* rather than some subjectively indistinguishable thought content. As we will see, this allows us to remove the seeming inconsistency between the supposed inconsistent triad of propositions underlying the McKinsey paradox.

4.1   Solving the Paradox

The basic form of the McKinsey paradox is that if content externalism is true, and if privileged access is true, then we can gain *a priori* knowledge of the external world. But *a priori* knowledge of the external world is absurd. To reiterate, the three propositions that McKinsey (1991) argues are inconsistent are:

(1) Oscar knows *a priori* that he is thinking that water is wet.

(2) The proposition that Oscar is thinking that water is wet necessarily depends upon E.

(3) The proposition E cannot be known *a priori*, but only by empirical investigation.

Recall that the proposition E that McKinsey refers to is the environmental condition that individuates the content of the thought as specified by content externalism. In the case of a thought about *water*, the external condition is that the thinker has interacted with water (H_2O). He also has privileged access to his own thought content. That is, he can know *a priori* what the content of his thoughts are. According to the McKinsey paradox, Oscar can know *a priori* facts about the external world. We can see this by way of the following argument:

(1) If Oscar knows that he has a thought about *water*, e.g. that 'water is wet,' then he can infer that he or someone in his community has interacted with H_2O

(2) Oscar can know *a priori* that he has a thought about *water*

(3) Therefore, Oscar can know *a priori* that he or someone in his community has interacted with H_2O

The first premise is supported by content externalism. Through philosophical theorizing and a familiarity with the sorts of externalist arguments put forth by Putnam, Burge, and McGinn, Oscar knows that the content of his thought is partly determined by factors outside of his head. In this particular case, Oscar can learn from the armchair that thoughts about *water* imply that he or someone in his community has interacted with H_2O. The second premise is supported by the thesis of privileged access. Oscar can know what he is thinking simply by introspecting his own mental state. The conclusion (3) follows from the conditional in premise (1) and the affirmation of the antecedent of that conditional in premise (2). But this conclusion is contrary to the claim that we cannot have *a priori* knowledge of the external world. And it seems absurd that Oscar can know *a priori* that his thoughts are about H_2O and not a perceptually indistinguishable chemical substance, XYZ. Hence, McKinsey's paradox: Oscar can know something about the external world *a priori*, even though he cannot know anything about the external world *a priori*.

Now suppose that knowledge is contrastive. If so, we can reformulate each premise of the argument that Oscar can gain *a priori* knowledge of the external world by making explicit the contrast set of each knowledge attribution. The conclusion of McKinsey's argument is that Oscar can gain *a priori* knowledge of the external world. In this case the knowledge that Oscar supposedly gains is that he or someone in his environment has interacted with H_2O. That is, Oscar can supposedly eliminate the possibility that he is on Twin Earth, which has an abundance of XYZ. The only way for Oscar to eliminate this possibility, without empirical investigation, is by being able to distinguish between thoughts about *water* and thoughts about *twater*. This would allow him to eliminate the possibility that his thought is about *twater* and hence that he is on Earth as opposed to Twin Earth. Accordingly, the contrast set in the antecedent of premise (1) needs to include *twater* and the contrast set in the consequent needs to include XYZ. Therefore, my reformulation of the first premise is:

> (1*) if Oscar knows that he has a thought about *water* rather than *twater*, then he can infer that he or someone in his community has interacted with H_2O rather than XYZ

However, when we ascribe to Oscar the knowledge of the content of his own thought, we *ought not* attribute to him the ability to distinguish between *water* and other possible content that is subjectively indistinguishable from *water*. For example, if Oscar was secretly transported to a Twin Earth environment and was there long enough for his 'water' thoughts to be about *twater,* he should not be able to know *a priori* that his thoughts are about XYZ and not H_2O. This sort of *a priori* access to facts about the external world is the absurdity that we are trying to avoid. But this is not to deny that Oscar has privileged access to his thought content. He can still eliminate the possibility that his thought is about a large set of subjectively distinguishable, possible contents. That is, when we ascribe to Oscar the knowledge that his thought is about *water* we are ascribing to him the

knowledge that his thought is about *water* rather than some subjectively distinguishable concept, such as *rye*.[8] The second premise can be reformulated as follows:

(2*) Oscar knows that he has a thought about *water* rather than *rye*

The reason that Oscar knows that his thought is about *water* rather than *rye* is because he can distinguish between these two thoughts. They are subjectively distinguishable due to their different subjective, introspectively accessible characteristics, such as the functional roles that they play in Oscar's mental life. For example, if Oscar wants to quench his thirst he might think to himself that 'I need water;' whereas if he wants to relax and go out dancing he might think to himself that 'I need rye.' He can distinguish between thoughts about *water* and *rye* because they play different roles in his reasoning, action, and speech. This is analogous to the perceptual justification that Moore has for his assertion that he knows that he has hands. The perceptual justification in Moore's case allows him to know that he has hands rather than stumps. Moore does not know that he has hands rather than vat-images of hands because hands and vat-images of hands are perceptually indistinguishable. In Oscar's case, he is able to subjectively distinguish between a thought about *water* and a thought about *rye*. The instantiation of the subjective characteristics of a thought about *water* provides Oscar the justification for his claim that he knows that his thought is about *water* rather than *rye*. I will briefly discuss a distinction often made between types of content to elucidate this point.

According to content externalism, mental states that have content, such as thoughts about water or arthritis, are *wide* states. For example, a thought about *water* is a wide state. Many philosophers believe that there are also *narrow* states. Consider again Putnam's Twin Earth thought experiment. Because Oscar and Toscar are physical duplicates, it is thought that they share some internal mental states. For example, they both have the same phenomenal experience when they look at the

---

[8] The claim is not limited to rye. The "rather than" clause indicates a set of alternatives including rye, but also whiskey, cider, mashed potatoes, etc. This is the same for Moore and his knowledge claim that he has hands. The set of contrasts in Moore's knowledge claim that he has hands is not limited to stumps; it includes hooks, fins, etc. The important point is that the contrast set that includes rye *does not* include twater.

clear liquid that they call 'water.' Their beliefs, though different, also play the same role in their mental lives. If they are thirsty, their belief that there is a glass of either *water* or *twater* on the table might cause them to pick up the glass and drink the clear liquid that it contains. These states that Oscar and Toscar share, the phenomenal experience of the liquid and the functional role that it plays in their mental lives, are often thought to be narrow states. However, the existence of narrow states and the relationship between narrow and wide states is controversial.[9]

One way to understand the relationship between wide and narrow mental states is to hold a two-factor or two-dimensional theory of mental content. According to this view, which has been defended by Ned Block (1994) and David Chalmers (2002) among others, the content of a thought is comprised of two distinct types of content, one narrow and one wide. For example, according to Block's two-factor conceptual role semantics,

> there are two components to meaning, a conceptual role component that is entirely "in the head" (this is the narrow meaning) and an external component that has to do with the relations between the representations in the head (with their internal conceptual roles) and the referents and/or truth-conditions of these representations in the world (93).

While two-factor theories offer valuable insight into the nature of the mental states that Oscar and Toscar share, they are inconsistent with the main claim of content externalism; namely, that certain thoughts are merely wide, not both narrow *and* wide. In order to deviate as little as possible from the main claim of content externalism, while nevertheless keeping in mind the insights of two-factor theories, I suggest that we allow for the sorts of states that Block and Chalmers describe while remaining agnostic towards the status of those states. Gilbert Harman (1999) and Robert Stalnaker (1999) have each argued that the mental states that are often referred to as narrow are in fact wide and are individuated according to environmental factors. As Stalnaker writes, "the externalist can account both for the role of content in psychological explanation and for privileged access and

---

[9] For example, see Segal (2000) and Stalnaker (1999).

transparency of self-interpreting character of thought, without invoking a notion of narrow content" (209).

Stalnaker acknowledges that denying that mental states have a narrow component may cause problems for privileged access. In anticipation of a solution to this problem he suggests that "the strategy that seems to me most promising...recognizes that claims to knowledge are essentially contrastive and context-dependent" (1999, 208-9). Stalnaker's proposed strategy is the strategy that I am currently pursuing. I maintain that Oscar does indeed have privileged access to the content of his own thoughts that is *a priori* in the sense that he does not have to infer that content based on his experience of the external world. However, he is not in a position to eliminate *all* contrast sets *a priori*. He can only eliminate contrast sets that have different subjective characteristics than the thought that he introspects.

To sum up my diagnosis of the McKinsey paradox: the contrast set in premise (2*), which is that Oscar knows that he has a thought about *water* rather than *rye,* is not the same contrast set that is in the antecedent of premise (1*), which is that Oscar knows that he has a thought about *water* rather than *twater*. In order to gain *a priori* knowledge of the external world, Oscar must be able to know *a priori* that his thought is about *water* rather than *twater*. But Oscar is incapable of knowing this because thoughts about *water* and *twater* are subjectively indistinguishable. However, this is not to deny that Oscar has privileged access to his thought content. Oscar can know that his thought is about *water* rather than *rye* since thoughts about *water* and *rye* are subjectively distinguishable. This diagnosis is similar to the contrastivist diagnosis of Moore and his knowledge claim that he has hands. Moore does not know that he has hands rather than vat-images of hands because the two are perceptually indistinguishable. But Moore does know that he has hands rather than stumps.

We are now in a position to see how a contrastive view of self-knowledge can solve the McKinsey paradox. The argument can be reformulated as follows:

(1*) if Oscar knows that he has a thought about *water* rather than *twater*, then he can   infer that he or someone in his community has interacted with H_2O rather than XYZ

(2*) Oscar knows that he has a thought about *water* rather than *rye*

(3*) therefore, Oscar can know *a priori* that he or someone in his community has interacted with H_2O rather than XYZ

According to a contrastive view of self-knowledge, the argument that leads to the McKinsey paradox is invalid. It is not the case that the conjunction of content externalism and privileged access implies the absurd conclusion that we can gain knowledge of the external world prior to experiencing it. However, we have not denied any of the three theses that form McKinsey's inconsistent triad. Nor have we denied closure. I will now consider the similarities between the contrastivist solution and the warrant transmission failure solution, concluding that the contrastivist solution captures the insights of warrant transmission failure and does a better job identifying our epistemic position in regards to the paradox.

4.2   Explaining Away Warrant Transmission Failure

As mentioned earlier, one prominent explanation of the McKinsey paradox is known as warrant transmission failure. Warrant transmission failure has also been used to explain cases such as Dretske's zebra case and the immodest knowledge claim regarding Moore not being a brain in a vat. Brown has characterized what she calls the *transmission principle* as the claim that the belief in the conclusion of an argument "constitutes knowledge in virtue of the fact that it is inferred from the known premises of an inference that is known to be valid" (242). Crispin Wright (2000) and Martin Davies (1998) have both presented arguments that this principle fails in various cases. I will focus on Wright's version of the argument, which is based on the claim that the arguments in cases such as the McKinsey paradox are circular (sometimes implicitly). The arguments are circular in that one

cannot have warrant for the premises of the argument *before* one has warrant for the conclusion of the argument. Wright provides a template that helps to identify such cases.

The template applies to arguments of the form A, A implies B, therefore B. An argument is said to be a case of warrant transmission failure if it meets these conditions:

"(i) A entails B.

(ii) There is a proposition C that is incompatible with A.

(iii) My warrant for A consists in my being in a state that is subjectively indistinguishable from a state in which C would be true.

(iv) C would be true if B were false" (Wright 2000, 155).

If we plug in the McKinsey paradox into this template we get:

(A) I am thinking that water is wet.

(A implies B) If I am thinking that water is wet, then I have interacted with water.

(B) I have interacted with water.

(C) I am thinking that twin-water is wet.

This argument meets condition (i) because A does entail B. Regarding condition (ii), the proposition C that is incompatible with A is the proposition that I am thinking that twin-water is wet. Either the thought is about water or twin-water, but not both. Regarding condition (iii), my warrant for A consists in the introspective access I have to the thought I am thinking that "water is wet." I am thinking about the clear flavorless liquid called 'water' that quenches thirst. I know that I am thinking this as opposed to "rye is wet." The thought was caused by seeing a glass of clear flavor-less liquid, and it might have the effect that I drink the liquid in order to quench my thirst. If I was transported to a twin-Earth environment long ago that contains no water but plenty of twin-water, then I'd still be in the same exact subjective state. I would still think to myself that "water is wet." But in that case, B would be false and C would be true. Condition (iii) is met because my warrant for thinking about water is subjectively indistinguishable from the state I would be in if I were thinking

about twin-water. Finally, condition (iv) also holds. If it is false that I have interacted with water, then I must be thinking that some other clear flavorless liquid is wet, that is, I must be thinking that twin-water is wet.

Since the warrant for A and C is the same, I cannot move beyond the disjunction of knowing that (A or C) without having independent reason to rule out C. The independent reason could be that B is true, but this means that the warrant for A depends on the warrant for B, and hence the argument is circular. As I have demonstrated above, the McKinsey paradox fits this warrant transmission failure template. That is, since the warrant for the premises of the argument does not transmit to the conclusion. Hence, the absurd conclusion that we can gain *a priori* knowledge of the external world is unwarranted.

There are two significant similarities between the warrant transmission failure solution and the contrastivist solution to the McKinsey paradox. First, both solutions rely on the notion of what is or is not subjectively distinguishable. Second, both solutions concede that what we have privileged access to in the immediate *a priori* sense is a disjunction of possible, subjectively indistinguishable thought contents. However, the contrastivist solution offers a deeper and clearer diagnosis of the problem.

One significant advantage of the contrastivist solution is that it does not simply identify implicit circularities in arguments. By forcing us to fill in the third argument of the knowledge relation, we confront some of the issues that cause the problem right away. For example, when forced to provide the third argument of my knowledge regarding the zebra in front of me, I quickly realize that I know it is a zebra as opposed to an elephant but not a zebra as opposed to a cleverly disguised mule. As we have seen, determining whether or not an argument is implicitly circular and a case of warrant transmission failure requires a somewhat elaborate process wherein we need to determine whether or not the argument satisfies conditions i-iv of Wright's template for warrant transmission

failure. Contrastivism, on the other hand, provides a surface indication that the arguments we have been discussing are invalid. As we have seen, once we are forced to fill in the third place of the knowledge relation, it becomes clear that the problem occurs due to illicit shifting of the contrast variable. This surface explanation of problems like the zebra case and the McKinsey paradox is simpler and more elegant than the warrant transmission failure solution.

Contrastivism also allows us to gain a deeper understanding of our epistemic position. According to the warrant transmission failure solution, the warrant for the premises does not transmit to the conclusion because that warrant itself depends on independent warrant for the conclusion. This diagnosis of the problem clearly states what we do not know, but it fails to highlight what we do know. Contrastivism, on the other hand, identifies the contrast set that can be eliminated during the knowledge claim. For example, while both warrant transmission failure and contrastivism conclude that we don't know if we're thinking about water or twin-water, only contrastivism makes the positive claim that we do know that we're thinking about water rather than rye (or mashed potatoes, gasoline, etc.).

The contrastivist solution to the McKinsey paradox exhibits the positive attributes of the warrant transmission failure solution while avoiding its negative attributes. It is also able to generalize to cases such as the zebra case and the Moore hand knowledge case. This is sufficient reason for preferring the contrastivist solution.

4.3   A Final Objection

One final objection to my proposal might run as follows: "How can Oscar know that his thoughts are about *water* rather than *rye*? If he can't distinguish between *water* and *twater*, then how does he know that his thought is about *water* rather than *rye*, as opposed to *twater* rather than *rye*?"

I agree that Oscar cannot know that his thought is about *water* rather than *rye*, as opposed to *twater* rather than *rye*. However, I do not believe that this counts against my proposal. Because Oscar cannot distinguish between thoughts about *water* and *twater*, a thought about either of them is indeterminate regarding which it actually is. It will require empirical investigation of some sort to settle the matter. That is, it is technically the case that when Oscar knows that his thought is about *water* rather than *rye,* what Oscar knows is that his thought is about (*water* or *twater*) rather than *rye*. Since *water* and *twater* play the same functional role for Oscar, a thought about one could possibly be a thought about the other. Again, Oscar will have to investigate his environment to settle the matter. In light of this, we can reformulate the second premise as follows:

> (2**) Oscar knows that he has a thought about (*water or twater or anything else that has the same functional role*) rather than *rye*

However, it is still the case that Oscar has privileged access to the fact that his thought is about one of a set of possible concepts that all share the same functional role and subjective character. But this privileged access does not enable Oscar to gain *a priori* knowledge of the external world.

5   CONCLUSIONS

I have tried to show that privileged access and content externalism are compatible. The conjunction of these two theses does not entail the ability to gain *a priori* knowledge of the external world. However, my proposal does force us to reconsider the nature of privileged access in light of a contrastive account of self-knowledge. This reconsideration has allowed us to discern exactly what we have access to. If two concepts play the same functional role in our mental lives, then we cannot determine which of the two concepts we are thinking about. Realizing that we do not have this sort

of privileged access to our thought content is a small price to pay for avoiding the McKinsey paradox.

One of the upshots of this proposal is that it is consistent with empirical considerations regarding the access that we have to our own mental states. There is an abundance of data that suggests that we do not always introspect our own mental states accurately. For example, it has been argued that we often wrongly attribute emotional states and intentional states such as beliefs and desires to ourselves.[10] If this is the case, then the theoretical implications that a contrastivist solution to the McKinsey paradox has for our notion of privileged access should be anticipated and welcomed. Additionally, contrastivism itself is a model of knowledge attribution that is, as I have discussed, gaining much empirical and theoretical support. It is also being applied to philosophical problems that are not directly epistemological including: moral theory, linguistics, probability theory, causation, and free will.[11] If contrastivism accurately captures the knowledge relation, and if we know that we often don't have complete privileged access to our mental states, then a contrastive account of self-knowledge is exactly what is called for.

---

[10] For an overview of the empirical research on which these claims are based see Daniel Haybron (2007) and Eric Schwitzgebel (2008).

[11] Blaauw (forthcoming) is a collection of essays on contrastive approaches to these and other topics.

6 REFERENCES

Blaauw, Martijn. Forthcoming. *Contrastivism in Philosophy*. Routledge.

Block, Ned. 1994. "Advertisement for a Semantics for Psychology." In *Mental Representations: A Reader,* edited by Stephen Stich and Ted Warfield, 81-141. Cambridge, MA: Blackwell.

Boghossian, Paul. 1989. "Content and Self-Knowledge." *Philosophical Topics* 17 (1): 5-26.

Brown, Jessica. 2004. *Anti-Individualism and Knowledge*. Cambridge, MA: MIT Press.

Brueckner, Anthony. 1992. "What an Anti-Individualist Knows *A Priori*." *Analysis* 52 (2): 111-18.

Burge, Tyler. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4 (1): 73-121.

Chalmers, David. 2002. "The Components of Content." In *Philosophy of Mind: Classical and Contemporary Readings*, edited by David Chalmers, 608-33. Oxford: Oxford University Press.

Davies, Martin. 1998. "Externalism, Architecturalism, and Epistemic Warrant." In *Knowing Our Own Minds*, edited by Cynthia Macdonald, Barry Smith, and Crispin Wright, 321-62. Oxford: Oxford University Press.

---. 2000. "Externalism and Armchair Knowledge." In *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, 384-414. Oxford: Oxford University Press.

DeRose, Keith. 2011. "Contextualism, Contrastivism, and X-Phi Surveys." *Philosophical Studies* 156 (1): 81-110.

Dretske, Fred. 1970. "Epistemic Operators." *The Journal of Philosophy* 67 (24): 1007-23.

Harman, Gilbert. 1999. "(Nonsolipsistic) Conceptual Role Semantics." In *Reasoning, Meaning, and Mind*, 206-34. Oxford: Oxford University Press.

Haybron, Daniel. 2007. "Do We Know How Happy We Are? On Some Limits of Affective Introspection and Recall." *Noûs* 41 (3): 394-428.

Kallestrup, Jesper. 2011. "Recent Work on McKinsey's Paradox." *Analysis* 71 (1): 157-71.

McGinn, Colin. 1977. "Charity, Interpretation, and Belief." *The Journal of Philosophy* 74 (9): 521-35.

McKinsey, Michael. 1991. "Anti-Individualism and Privileged Access." *Analysis* 51 (1): 9-16.

---. 1994. "Accepting the Consequences of Anti-individualism." *Analysis* 54 (2): 124-28.

Neta, Ram. 2008. "Undermining the Case for Contrastivism." *Social Epistemology* 22 (3): 289-304.

Pryor, James. 2000. "The Skeptic and the Dogmatist." 34 (4): 517-49.

Putnam, Hilary. 1975. "The meaning of 'meaning.'" In *Mind, Language, and Reality*, 215-71. Cambridge: Cambridge University Press.

Schaffer, Jonathan. 2004. "From Contextualism to Contrastivism." *Philosophical Studies* 119 (1/2): 73-103.

---. 2007. "Closure, Contrast, and Answer." *Philosophical Studies* 133 (2) 233-55.

Schaffer, Jonathan, and Joshua Knobe. Forthcoming. "Contrastive Knowledge Surveyed." *Noûs*.

Schwitzgebel, Eric. 2008. "The Unreliability of Naive Introspection." *Philosophical Review* 117 (2): 245-73.

Segal, Gabriel. 2000. *A Slim Book about Narrow Content*. Cambridge, MA: MIT Press.

Stalnaker, Robert. 1999. "Narrow Content." In *Context and Content: Essays on Intentionality in Speech and Thought*, 194-209. Oxford: Oxford University Press.

Wright, Crispin. 2000. "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof." *Philosophical Issues* 10 (1): 140–63.

---. 2003. "Some Reflections on the Acquisition of Warrant by Inference." In *New Essays on Semantic Externalism and Self-Knowledge*, edited by Susana Nuccetelli, 57–78. Cambridge, MA: MIT Press.

---. 2011. "McKinsey One More Time." In *Self-Knowledge*, edited by Anthony Hatzimoysis, Oxford: Oxford University Press.