

following the superevaluationist policy of coping with semantic indeterminacy by saying what's true on all resolutions of the indeterminacy.

22. The real lesson of Kripke's alleged examples of *a posteriori* necessity and *a priori* contingency (Kripke 1980, 97ff., 128ff.), is that when we have reference fixing, we have two different ways for a sentence to express a proposition. It may happen that the proposition expressed one way is necessary and knowable *a priori* but the proposition expressed the other way is contingent and knowable only *a posteriori*.

See Chalmers 1996, 56ff.; Stalnaker 1978; Jackson 1998a, 46ff. Let sentence S contain terms subject to reference fixing. (Ignore other forms of indexicality.) The *primary intension* of S (Chalmers), or the *diagonal proposition* (Stalnaker), or the *A-intension* (Jackson) is the proposition that is true at just those worlds W such that S is true at W with respect to reference fixing as it is at W. The *secondary intension* of S (Chalmers), or the *horizontal proposition* (Stalnaker), or the *C-intension* (Jackson) is the proposition that is true at just those worlds W such that S is true at W with respect to reference fixing as it is in actuality. If, for instance, the reference of 'heat' is fixed on molecular motion at the actual world but on other things at other worlds, the primary intension of 'Heat is molecular motion' is contingent but the secondary intension is necessary; and *vice versa* for 'Heat is the occupant of the heat-hole'.

So far, when I have spoken of the propositions expressed by the answer-sentences, I've meant the secondary intensions. If instead we took the primary intensions, it would be just as if we'd omitted the rigidifier 'actual'. As already noted, that would not be a way to ask the question we meant to ask. We would know the primary intensions of the answer-sentences, but these would not be the answer-propositions that identify the properties that occupy the roles.

23. See Maxwell 1978; and for a sympathetic discussion of the view, see Chalmers 1996, 135ff. Just as our pansychist thinks it is because fundamental properties are qualia that we can know their identities by acquaintance, so others think it is because (or insofar as) they are *not* qualia that we cannot know their identities. See, *inter alia*, Russell 1927, 497; Unger 1999.

10 A Partial Defense of Ramseyan Humility

Dustin Locke

1 Introduction

In "Ramseyan Humility" (this volume), David Lewis argues that "we are irremediably ignorant about the identities of the fundamental properties that figure in the actual realization of the true final theory" (214). Of the three published responses to Lewis's essay (Langton 2004; Schaffer 2005; Whittle 2006), each argues that even if we accept Lewis's metaphysical assumption (the thesis known as 'quidditism'), we need not accept his epistemic conclusion (hereafter, 'Humility'). My aim in this paper is to defend Lewis against these critics.

Ann Whittle claims that Humility rests on an implausibly strong account of *identification*—namely, that to identify X, one must be able to distinguish X from all other actual *and possible* entities. Accordingly, Whittle attempts to refute Humility by an appeal to a more lenient account of identification. In what follows I defend Lewis by showing that Lewis's demanding account of identification is a perfectly good account of at least one perfectly legitimate sense of 'identification'.

Jonathan Schaffer and Rae Langton each claim that Humility rests on an implausibly strong account of *knowledge*—namely, noncontextual infallibilism. Accordingly, Schaffer and Langton attempt to refute Humility by appealing to more lenient epistemic principles. In what follows I defend Lewis, not by defending noncontextual infallibilism, but by showing that even on the alternative epistemic principles of Schaffer and Langton, Lewis's conclusion still follows.

2 The Argument for Humility

Lewis's argument can be conveniently reconstructed in five stages.

The stage-setting stage. Let T be a true and complete final theory of our world. Lewis claims that the language of T can be divided into two parts: the terms that have meaning independently of T (the 'O-terms') and the terms that are implicitly defined by T (the 'T-terms'). Writing T as a single sentence and replacing the T-terms with variables, we get the realization formula of T—the open sentence ' $T(x_1, x_2, \dots, x_n)$ '. Any n -tuple of properties that *could have* satisfied the realization formula of T is called a 'possible realization of T'; any n -tuple of properties that *does* satisfy the realization formula of T is called an 'actual realization of T'.

*The metaphysical/linguistic stage.*¹ Lewis assumes that there is a realization of T, say, $\langle P_1, P_2, \dots, P_n \rangle$ such that at least two of P_1, P_2, \dots, P_n are fundamental properties of the same adicity (1-place, 2-place, etc.) and category (all-or-nothing properties, real-valued magnitudes, etc.). He further assumes that if the O-language refers to the properties that realize T, it is only via descriptions of their nomological/locational roles in nature. Finally, Lewis assumes quidditism—the thesis that for any fundamental property P and nomological/locational role of appropriate adicity and category R, P might have realized R.² Given these assumptions, it follows that there is some permutation of the actual realization of T that is a possible realization of T. Since the actual realization of T is of course also a possible realization of T, it follows that T has *multiple possible* realizations.

The first epistemic stage. Lewis assumes that the O-language of T suffices to describe all possible observations. Now consider the Ramsey sentence of T— $\exists x_1, x_2, \dots, x_n T(x_1, x_2, \dots, x_n)$. Since the Ramsey sentence of T implies all and only the O-language sentences of T, it follows from the above assumption that our (observational) evidence will always be consistent with the proposition expressed by the Ramsey sentence of T. Thus, for any world w such that the proposition expressed by the Ramsey sentence of T is true at w , w will never be eliminated by—that is, inconsistent with—our evidence.

The second epistemic stage. Although Lewis doesn't say so, the most straightforward way to get from the conclusion of the previous stage to the conclusion of Ramseyan Humility is via Lewis's own account of knowledge. In Lewis 1999, Lewis defends the following infallibilist account of knowledge: S knows that P iff S's evidence eliminates all possibilities in which not-P (psst!—see this note).³ This analysis, together with the conclusion of the last stage of the argument, implies that we will never know p , where p is the proposition that is true at all and only those worlds where the actual realization of T is the realization of T (= the proposition that is true at all and only those worlds where $\langle P_1, P_2, \dots, P_n \rangle$ realizes T).⁴

The third epistemic stage. The final stage of the argument takes us from the claim that we will never know p (from above) to the claim that we will never know *which* n -tuple of properties realizes T. As in the last stage of the argument, Lewis leaves the key assumption here implicit, but the path seems fairly clear. We first assume what Jonathan Schaffer (2007) has called the "Standard Analysis of knowing-wh," according to which S knows-wh (who, which, what, etc.) iff S knows P, where P is the true answer to the indirect question of the wh-clause. If we assume also that p is the true answer to the question 'Which n -tuple of properties realizes T?' it follows from the conclusion of the previous stage of Lewis's argument that we will never know which n -tuple of properties realizes T.

As I said above, Whittle, Schaffer, and Langton all argue that even if Lewis's metaphysical assumption (quidditism) is correct, his conclusion (Humility) can be blocked. We have just seen that Lewis's conclusion *cannot* be blocked if we assume the noncontextual version of Lewis's infallibilist account of knowledge and the account of knowing-wh given in the third epistemic stage. But Whittle, Schaffer, and Langton are by no means obligated to grant Lewis these accounts. After all, the account of knowledge is an *infallibilist* account of knowledge and the account of knowing-wh is a quite strong one. Accordingly, we might have expected these accounts to imply some pretty serious epistemic depravity. The interesting question, then, is whether quidditism implies Humility on some more *lenient* epistemic principles. This is exactly the question that these authors attempt to address, and they all answer it in the negative. We will look at their arguments, and why I think they're wrong, in a moment. For now, we must pause to clear something up.

3 Interlude: Trivial Knowledge and Two Dimensions of Intension

In the last section I argued that Lewis's thesis that we will never know which properties realize T is equivalent to the thesis that we will never know p , where p is the proposition that is true at all and only those worlds where the actual realization of T is the realization of T. Now this may seem a bit odd, for this is just the thesis that

(1) We will never know that the actual realization of T is the realization of T.

But how could (1) possibly be true? What could be easier than knowing that the actual F is the F?

This response hinges on an ambiguity in (1). The ambiguity concerns which proposition is expressed by the embedded sentence 'the actual realization of T is the realization of T'. One way to understand the ambiguity is with the resources of "nonambitious" two-dimensional semantics.⁵ There are several particular versions of two-dimensional semantics, but only the general idea is needed here. This is the idea that at least some statements have two intensions: a C-intension and an A-intension.⁶ The A-intension of a statement, roughly, is the proposition that is true at all and only those worlds *considered as actual* where the statement is true ('A' for 'actual'). The C-intension, on the other hand, is the proposition that is true at all and only those worlds *considered as counterfactual* where the statement is true ('C' for 'counterfactual').

Let's take an example. Consider the statement 'The actual US president is the US president'. Now consider a world w_1 where the US president is Al Gore. Is the statement 'The actual US president is the US president' true at w_1 ? Well, that depends on whether we're considering w_1 as actual or as counterfactual. Consider it first as counterfactual. In other words, ask yourself what would have been true had Al Gore been the US president. In particular, ask yourself if the actual US president would have been the US president. The answer, it seems, is 'no'—the actual president, George W. Bush, would not have been the president had Al Gore been the president. But now consider w_1 as actual. In other words, ask yourself what is true if Al Gore is the US president. In particular, ask yourself if the actual US president is the US president if Al Gore is the US president. The answer in this case, it seems, is 'yes'. Thus, the A-intension and C-intension of 'The actual US president is the US president' are not identical: the former is true at w_1 , but the latter is not.

Exactly the same thing happens in the case of 'The actual realization of T is the realization of T'. The A-intension of this statement is true at all and only those worlds where T is uniquely realized, regardless of which properties realize it at that world. This is simply because for any world w where T is uniquely realized, the statement 'The actual realization of T is the realization of T' is true at w considered as actual. However, the C-intension of 'The actual realization of T is the realization of T' is not true at every world where T is uniquely realized. In particular, it is false at any world where T does not have the same realization that it does in this (the actual) world. This allows us to explain the ambiguity in (1). If the embedded statement in (1) expresses its A-intension, then (1) is false, provided only that we will come to know that T is uniquely realized. However, only if the embedded statement in (1) expresses its C-intension is (1) equivalent to Lewis's thesis.

This concludes my explication of Lewis's thesis and his argument for it. I turn now to the criticisms of Lewis and my defense.

4 Whittle's Response to Humility

Ann Whittle argues that even if Lewis's metaphysical assumptions are correct, "we can [still] know which property occupies a certain role, because we are able to *identify* the property in question" (2006, 469, emphasis added). I disagree.

Whittle's argument explicitly assumes a version of Russell's principle: being able to identify A is having 'discriminating knowledge' of A, where one has discriminating knowledge of A if one has a description that uniquely picks out A. Thus, since we can uniquely pick out the property that realizes such-and-such role with the description 'the property that realizes such-and-such role', it follows by Whittle's version of Russell's principle that we can identify the property that realizes such and such role.

The first thing to notice about Whittle's version of Russell's principle is just how *cheap* it makes identification. An example should make this clear. Suppose we're in a courtroom and I'm on the stand. An attorney asks me, "Do you know which person stole your chicken?" Suppose that all I know about the person who stole my chicken is that the person who stole my chicken is the person who stole my chicken. However, it turns out that there is a unique person who stole my chicken. It follows by Whittle's version of Russell's principle that I can identify the person who stole my chicken.

Clearly Lewis had a stronger notion of 'identification' in mind when he stated his thesis. Whittle seems to agree. She writes, "Lewis, however, seems to think that something more is required: that we be able to single out the entity from all other actual *and possible* entities" (*Ibid.*, 470). But she goes on to argue that

[S]uch a condition on identification seems too strong. Suppose, for instance, that I am a historian who knows all there is to know about Napoleon and this knowledge, we can safely assume, allows me to single out one person in history, namely Napoleon. Granted a certain view about the transworld identities of particulars, namely Haecceitism, this exhaustive knowledge of Napoleon would, nevertheless, fail to identify him. Why? Haecceitism states that two possibilities can differ just in the permutation of individuals. So everything could be qualitatively identical in two possible worlds, yet those worlds differ in that the person we name 'Napoleon' instantiates all the same properties that Nelson does in this world, and vice versa. (*Ibid.*)

Whittle's mistake, it seems to me, is in searching for *one* (or perhaps a *few*) notions that deserve the title 'identification'. As many theorists have noted, expressions such as 'identification', 'knowing what', and 'knowing which' are all highly context-sensitive.⁷ Lewis himself notes that in addition to all the common and less demanding senses of these expressions, there is "an uncommonly demanding and literal sense" (Lewis 1995, 142). Moreover, Lewis explicitly notes (this volume, 215 and endnote 20) that it is this "uncommonly demanding sense" that he has in mind when he says that we will never know which properties realize T.

Above I stated that according to the standard analysis of knowing-*wh*, S knows-*wh* iff S knows P, where P is the true answer to the indirect question Q of the *wh*-clause. The different senses of 'knows-*wh*' can be modeled using this analysis by varying what counts as the true answer to Q. If we want a more demanding sense, we simply require that only more specific propositions count as answers to Q. Taking this process to its limit, we get the "uncommonly demanding and literal sense" of knowing-*wh*. In this sense of knowing-*wh*, to know-*wh* is to know the maximally specific answer to the question. In our case, the question is "Which possible realization of T is the realization of T?" and the maximally specific answer is the proposition that $\langle P_1, P_2, \dots, P_n \rangle$ realizes T, which, by the second epistemic stage of Lewis's argument, is a proposition that we will never know.⁸

In any case, whether or not Lewis is correct to state his thesis in terms of 'knowing which' or 'identification'—that is, whether the third epistemic stage of Lewis's argument is sound—the more interesting question would seem to be whether the argument through the second epistemic stage is sound. After all, the second epistemic stage itself concludes with a claim about our irremediable ignorance. Let us then move on to a discussion of Schaffer (2005) and Langton (2004), whose criticisms appear to be directed precisely at this earlier stage of the argument.

5 This Ain't Your Daddy's Skepticism

In "Quiddistic Knowledge," Jonathan Schaffer argues that Lewis's thesis "is just a species of skepticism about the external world" and that, therefore, "whatever answer one offers to skepticism about the external world will thereby answer quiddistic skepticism" (Schaffer 2005, 19).

I believe that Schaffer is wrong that Humility is "just a species of skepticism about the external world." Moreover, I believe that the way in which Humility differs from traditional external-world skepticism is precisely

what makes Humility *immune* to almost all of the typical responses to traditional skepticism.⁹ Let me first illustrate the difference between Humility and traditional skepticism by discussing two of Schaffer's attempts to apply a standard response to traditional skepticism to Humility. After I discuss these particular responses, I'll attempt to say in somewhat more precise terms what exactly the difference is between Humility and traditional skepticism and why most standard responses to the latter will therefore be unsuccessful responses to the former.

5.1 Deductionism (aka the Moorean Response)

Consider what Schaffer calls the "deductionist" response to external-world skepticism. A deductionist response takes the following form:

- (2) I know that *p* (by common sense).
- (3) If I know that *p*, then I am in a position to know that *q*.
- (4) Therefore, I am in a position to know that *q*.

where *p* is such that it is "common sense" that I know that *p*, and *p* entails *q*. If we add to the above argument the assumption that I have deduced *q* from *p*, it follows that I know that *q*.

For example, the deductionist might respond to the brain-in-a-vat skeptic as follows:

- (5) I know that I have hands (by common sense).
- (6) If I know that I have hands, then I am in a position to know that I am not a handless brain in a vat.
- (7) Therefore, I am in a position to know that I am not a handless brain in a vat.

Might we offer a deductionist response to Humility? Schaffer thinks so: "[Just] as the deductionist claims that by starting with one's knowledge that one has hands, one can come to know that the external world is real, so she should claim that by starting with one's knowledge that, for instance, this brick has mass, one can come to know the quiddities. The same deductive moves are available in both cases" (2005, 21). So Schaffer thinks that the deductionist should offer something like the following reply to Humility:

- (8) I know that this brick has mass (by common sense).
- (9) If I know that this brick has mass, then I am in a position to know that mass realizes the mass role.¹⁰
- (10) Therefore, I am in a position to know that mass realizes the mass role.

Will this line of argument work? If so, it must be "common sense" that I know that this brick has mass. Suppose, on the deductionist's behalf, that it is common sense that I know that this brick has mass. What proposition do I thus know?

I stated earlier that T implicitly defines the T-terms. According to Lewis (1970, 1972, and this volume) a T-term t_i is equivalent to the definite description 'the x_i such that there exist unique $x_{i1}, x_{i2}, \dots, x_i$ such that $T(x_{i1}, x_{i2}, \dots, x_i)$ '—for short, 'the property that realizes the t_i role'. Such definitions are commonly known as 'Ramsey-style definitions'. Here is one variation of the story that Lewis often told to motivate the idea that T-terms have Ramsey-style definitions:

We are assembled in the drawing room of the country house; the detective reconstructs the crime. That is, he proposes a *theory* designed to be the best explanation of the phenomena we have observed: the death of Mr. Body, the blood on the wallpaper, the silence of the dog in the night, the clock seventeen minutes fast, and so on. He launches into his story:

X, Y and Z conspired to murder Mr. Body. Seventeen years ago, in the gold fields of Uganda, X was Body's partner... Last week, Y and Z conferred in a bar in Reading... Tuesday night at 11:17, Y went to the attic and set a time bomb... Seventeen minutes later, X met Z in the billiard room and gave him the lead pipe... Just when the bomb went off in the attic, X fired three shots into the study through the French windows....

And so it goes: a long story. Let us pretend that it is a long conjunctive sentence.

The story contains three names, 'X', 'Y' and 'Z'. The detective uses the new terms without explanation, as though we knew what they meant. But we do not. We never used them before, at least not in the senses they bear in the present context. All we know about their meanings is what we gradually gather from the story itself. Call these... *T-terms*... because they are introduced by a theory. Call the rest of the terms in the story *O-terms*....

In telling his story, the detective set forth three roles and said that they were occupied by X, Y and Z. He must have specified the meanings of the three T-terms 'X', 'Y' and 'Z' thereby; for they had meanings afterwards, they had none before, and nothing else was done to give them meanings. They were introduced by an implicit functional definition, being reserved to name the occupants of the three roles....

If, as I claim, the T-terms are definable as naming the first, second and third occupants of the unique triple that realizes the story, then the T-terms can be treated like definite descriptions.... (Lewis 1972, 249-252)

And just in case you were thinking that the descriptions merely fix the referents of the T-terms, without actually giving their semantic contents, Lewis goes on to say

If the story is uniquely realized, they name what they ought to name; if the story is unrealized or multiply realized, they are like improper descriptions. If too many triples realize the story, 'X' is like 'the moon of Mars'; if too few triples—none—realize the story, 'X' is like 'the moon of Venus'. Improper descriptions are not meaningless. Hillary Putnam has objected that on this sort of account of theoretical terms, the theoretical terms of a falsified theory come out meaningless [Putnam 1962]. But they do not, if theoretical terms of unrealized theories are like improper descriptions. 'The moon of Mars' and 'The moon of Venus' do not (in any way) name anything here in our actual world; but they are not meaningless, because we know very well what they name in certain alternative possible worlds. (Ibid., 252-253)

If Lewis's semantics of T-terms is correct, and if 'mass' is a T-term as we are supposing, then the term 'mass' is equivalent to the definite description 'the property that realizes the mass role'.¹¹ Let us then reconstruct the deductionist's argument ((8), (9), and (10)) in light of this equivalence.

(8*) I know that this brick has the property that realizes the mass role (by common sense).

(9*) If I know that this brick has the property that realizes the mass role, then I am in a position to know that the property that realizes the mass role is the property that realizes the mass role.

(10*) Therefore, I am in a position to know that the property that realizes the mass role is the property that realizes the mass role.

According to (10*), what I am in a position to know is that *some unique* property realizes the mass role. But Lewis's thesis is that we will never know *which* property (in the demanding sense) realizes the mass role—that is, we will never know the proposition that is true at all and only those worlds where the property that *actually* realizes the mass role, is the property that realizes the mass role.

What about rigidifying our descriptions? Consider

(8**) I know that this brick has the property that *actually* realizes the mass role (by common sense).

(9**) If I know that this brick has the property that actually realizes the mass role, then I am in a position to know that the property that actually realizes the mass role is the property that realizes the mass role.

(10**) Therefore, I am in a position to know that the property that actually realizes the mass role is the property that realizes the mass role.

The conclusion of this argument (10**) indeed contradicts Lewis's thesis, but *only* if we read the sentence containing the definite description 'the

property that actually realizes the mass role' as expressing its C-intensions (see section 3 above). Is this a move the deductionist can make? No, because in that case the plausibility that the first premise is "common sense" has gone right out the window. What is common sense (if anything here is) is that we know that this brick has mass, which, by Lewis's semantics for T-terms, is knowing that this brick has the property that realizes the mass role.

5.2 Abductionism (aka Inference to the Best Explanation)

The abductionist response to external-world skepticism claims that (1) the hypothesis that the actual world is roughly as we perceive it to be is a better explanation of our experiences than the hypothesis that, say, we are disembodied dreamers and (2) if the hypothesis that p is the best explanation of one's experiences, then one can come to know that p by an *inference to the best explanation*. Schaffer claims that the abductionist ought to have a parallel response to Humility: "[J]ust as the abductionist claims that skeptical scenarios constitute poor explanations of the appearances, so she should claim that quiddity swapping scenarios constitute more complex, less conservative, or at least somehow inferior explanations of the powers" (Schaffer 2005, 22). In a footnote to this passage, Schaffer admits that the abductionist will have a difficult time making the case that the "quiddity swapping" scenarios (i.e., the alternative realization scenarios) constitute inferior explanations of the powers. However, he seems to be suggesting that since these same difficulties arise for the abductionist response to traditional skepticism, there is no special difficulty for the abductionist response to Humility.

I disagree. Although there are indeed difficulties for the abductionist's claim that the normal external-world hypothesis is a better explanation of our experiences than the skeptical hypothesis, at least in that case one can *distinguish* various features of the competing hypotheses that then figure in the arguments for which is the better explanation. For example, the external-world hypothesis postulates the existence of many things outside of our minds, whereas the dreaming hypothesis does not. But are there any such distinguishing features of the competing hypotheses in the case of Humility?

Let R_1 and R_2 be two possible realizations of T , such that the first and second members of R_1 are permuted in R_2 . Now consider H_1 , the hypothesis that R_1 realizes T and H_2 , the hypothesis that R_2 realizes T . What are the distinguishing features of H_1 and H_2 ? Well, simply the fact that according to H_1 , R_1 realizes T whereas according to H_2 , R_2 realizes T . Is this difference between H_1 and H_2 relevant to which is the better explanation? There

seems to be absolutely no way to tell, and the reason for this is a fact about how we are forced to *represent* these two hypotheses.

If R_1 is the actual realization of T , it may seem that we can represent H_1 with the sentence ' t_1, t_2, \dots, t_n ' realizes $T(x_1, x_2, \dots, x_n)$ ', where this sentence uses each T-term ' t_i '. But if we learned our lesson from above, then we'll know that that won't do the trick. According to Lewis, the sentence just mentioned is equivalent to the sentence 'the n -tuple that realizes T is the n -tuple that realizes T '. This sentence of course is trivial (save its existential and uniqueness presuppositions) and obviously does not express the hypothesis H_1 . What we'll have to do, if we want to use some sentence like the above to represent H_1 , is *rigidify* our descriptions. Thus, the sentence 'the n -tuple that *actually* realizes T is the n -tuple that realizes T ' manages to express the hypothesis H_1 . Similarly, the sentence 'the n -tuple formed by permuting the first and second members of the n -tuple that *actually* realizes T realizes T ' manages to express the hypothesis H_2 .

Our question, remember, is whether H_1 and H_2 have any distinguishing features such that we can infer to H_1 on the grounds that it is a better explanation than H_2 . But now seeing how we are forced to represent these hypotheses, it should be fairly clear that if there are features of H_1 and H_2 that make one a better explanation than the other, we are going to be completely ignorant of any such features. As Lewis would say, our rigidified descriptions enable us to represent the hypotheses in question but only in such a way that we are "blinded" to their distinguishing features.¹²

Consider the detective in Lewis's story above. Suppose that Jack actually realizes the X role and that Bridget actually realizes the Y role. Now suppose that the *only* means that the detective has of representing the hypothesis that Jack realizes the X role and Bridget realizes the Y role is with the sentence 'The person who actually realizes the X role realizes the X role and the person who actually realizes the Y role realizes the Y role'. Can the detective thus infer that proposition (over the proposition that Mary realizes the X role and John the Y role) by an inference to the best explanation? On pain of allowing too much knowledge too easily, we had better say "no." Through his ineliminable use of rigidified descriptions, the detective is blinded to any features of that hypothesis that make it a better explanation than the alternative. The same goes for us in the case of Humility.

5.3 The Diagnosis

At this point we should be asking ourselves questions along the following lines. Don't we, or might we someday, have a language rich enough to express H_1 *without* the use of rigidification? If so, couldn't we then infer to

that proposition via inference to the best explanation? Moreover, if we have such a language, then perhaps it contains some sentence *S*, such that it is (will be) common sense that we know *S* (here *S* is used) and such that *S* expresses the proposition that this brick has the property that actually realizes the mass role. If so, couldn't we then deduce the proposition that Lewis claims we don't know?

The crux of the matter is whether we have or ever will have a language that is expressive enough to rigidly refer to the properties that realize *T* without the use of rigidified descriptions. Lewis seems committed to saying "no." The theory *T*, recall, is supposed to be the "true and complete final theory of our world." It would thus seem that the language of *T* is as expressive of a language as we will ever have. Thus, by Lewis's assumption that the language of *T* is not expressive enough to rigidly refer to the properties that realize *T* without the use of rigidified descriptions, it seems to follow that we will never have a language that is that expressive.¹³

In my opinion, it is no accident that Lewis's paper is titled "Ramseyan Humility." Humility is Ramseyan because it's due to a certain alleged fact about our language—namely, that our only terms for the fundamental properties that realize *T* have Ramsey-style definitions. This, I believe, is really what distinguishes Humility from traditional skepticism.

5.4 Contextualism

It wouldn't be right to leave our current subject without discussing the prospects of applying Lewis's own response to traditional skepticism—contextualism—to Humility. I have chosen to discuss this type of response separately since I think it fails as a response to Humility for a reason independent of the one I've just been going on about.

When I gave Lewis's analysis of knowledge above, I left out his famous *sotto voce* proviso. Here is Lewis's analysis in its entirety:

Lewisian Knowledge "S knows that *P* if and only if *S*'s evidence eliminates all possibilities in which not-*P*—Psst!—except for those possibilities that we are properly ignoring" (Lewis 1999, 425).

Lewis argues that, in certain contexts, we can truly be said to 'know' (as uttered in those contexts) that we have hands, because we are properly ignoring all the possibilities in which we don't have hands that our evidence does not eliminate (e.g., the possibility of being a handless-brain-in-a-vat).

Can we make this sort of response to Humility? Again, Schaffer thinks so: "[u]s[ing] as the contextualist allows that claims to know that one has hands count as true when skeptical scenarios are not salient, so she should

allow that claims to know which properties exist count as true when quiddity-swapping scenarios are not salient" (Schaffer 2005, 23). Rae Langton (2004) agrees: "if Lewis shows us how we can 'properly ignore' the skeptical possibility, perhaps the same strategy could show us how we can 'properly ignore' for example, the possibilities of permutation, or replacement by 'fiddlers', that the argument for Ramseyan Humility exploits. The good news for my reader, then, is that your knowledge of things in themselves may be safe—at least as far as the argument for Ramseyan Humility is concerned" (Langton 2004, 134–135).

I don't think that a contextualist response to Humility has much plausibility. Lewis's brand of contextualism, in particular, is committed to this being the case. Here is Lewis's explanation of why we do not know that a given lottery ticket will lose, no matter how low the odds (a version of the so-called lottery paradox):

Suppose one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other It is the Rule of Resemblance that explains why you do not know that you will lose the lottery, no matter what the odds are against you and no matter how sure you should therefore be that you will lose. For every ticket, there is the possibility that it will win. These possibilities are saliently similar to one another: so either every one of them may be properly ignored, or else none may. But one of them may not properly be ignored: the one that actually obtains [by Lewis's Rule of Actuality]. (Lewis 1999, 429–430)

(Lewis also uses the Rule of Resemblance to explain why the subjects of Gettier cases do not know what they truly and justifiably believe.)

If Lewis is right that we may not properly ignore possibilities that saliently resemble actuality, then it is difficult to see how there could be a context in which we are properly ignoring alternative realization scenarios. As before, suppose R_1 is the actual realization of *T*, and R_2 is just like R_1 , except that the first two members of R_1 have been permuted in R_2 . What could more saliently resemble actuality (the possibility that R_1 realizes *T*) than the possibility that R_2 realizes *T*? This case seems exactly parallel to the lottery case. Suppose the winning ticket is T_1 and not T_2 . Lewis explains that we cannot ignore the possibility that T_2 is the winning ticket because it saliently resembles actuality (the possibility that T_1 is the winning ticket). Since the possibilities here merely differ over which ticket realizes the winning-ticket role, and this seems to explain their salient similarity, I don't see why we shouldn't say the very same thing about the possibilities of R_1 and R_2 realizing *T*.

I don't think that the argument against appealing to contextualism in the case of Humility applies only to Lewis's particular version of

contextualism. It seems to me that *any* version of contextualism that gets the lottery case right will be useless against Humility. The trouble is simply that, as long as we're talking about realizations of T, at least some realization of T other than the actual realization is going to be salient.¹⁴

This concludes my discussion of the responses to ordinary skepticism that won't work as responses to Humility.

6 Two Dogs That Might Hunt

In the preceding sections we have seen that three of the most popular responses to ordinary skepticism are unsuccessful when it comes to Humility. I will now briefly discuss two responses that do have some initial plausibility as responses to Humility.

But first one qualifier. My intension here is not to defend these responses as successful responses to Humility. To do that, I would need to defend the principles upon which they, respectively, rest. My point here is simply to show that *if* one of these principles is right, then Humility is false (or, rather, we have no reason to think that it is true).

6.1 Antiquidditism

The first response to consider is that of denying the metaphysical assumption upon which the argument for Humility rests, namely, quidditism. If, contrary to quidditism, possibility is not preserved by freely permuting properties with respect to their nomological/locational roles, then Lewis's claim that any permutation of T is a possible realization of T is false, and so Lewis's argument for Humility is unsound. Since this way of responding to Humility has been discussed at length in other places, I won't spill any more ink over this (difficult) issue here. However, since many take the epistemic consequences of quidditism—namely, Humility—as reason to reject quidditism, I hope to have shed some light on this issue by getting clearer on just what Humility amounts to and whether it follows from quidditism.

6.2 Direct Realism

The other response to Humility that has at least some initial plausibility is direct realism. Lewis himself considers what I believe qualifies as a direct realist response to Humility:

A friend of phenomenal qualia might speculate that all the actually instantiated fundamental properties are qualia. That would not preclude them from also occupying

physical roles—a sort of panpsychism. So, even if our true final theory is a physical theory, they might be the members of its actual realization. He might also accept the Identification Thesis: anyone acquainted with a quale knows just which property it is. Now it may seem that we can know the identities of the fundamental properties after all—we need only become acquainted with them. (This volume, 217)

Interestingly, Lewis rejects this response to Humility not because it requires that the fundamental properties be qualia, but because he rejects the Identification Thesis.

I do not want to debate the Identification Thesis here. My point is simply to note that (A) the response Lewis is considering seems to be something of a direct realist response and (B) the direct realist response is at least an appropriate response to Humility.

Why is the direct realist response an appropriate response to Humility? The direct realist will deny Lewis's assumption that the O-language refers to the fundamental properties only by means of descriptions of their roles. Consider the fundamental property referred to by the T-term ' t_1 '. Now imagine the direct realist becoming directly acquainted with t_1 (as a quale) and, according to the Identification Thesis, thereby knowing which property she is acquainted with. Now imagine her "baptizing" this property with the O-term ' o_1 '. This is enough to block Lewis's arguments that a permutation of this property within the actual realization of T is also a possible realization of T. Recall that that argument assumed that (1) the O-language refers to fundamental properties only by means of role/locational descriptions and (2) the O-language suffices to describe all possible observations. If the direct realist is right about this case, then it seems that ' o_1 ' does not refer to o_1 by means of a role-description. Thus, if the term ' o_1 ' is made part of the O-language, then (1) is not true. If, on the other hand, the term ' o_1 ' is not made part of the O-language (because, say, terms of direct acquaintance such as ' o_1 ' can only be part of a private language) then (2) is not true, because there is a possible observation that the O-language does not describe—namely, that of becoming directly acquainted with o_1 . (As Lewis notes, this second possibility is rather dubious: there seems to be no reason why agents would have a private language that was more expressive than the language of T, which is, recall, the true and complete *final* theory of our world.)

7 Capitulation

The final response to Humility we should consider is capitulation. This is of course Lewis's own response. He asks, rhetorically, "[Why] should I want

to block that argument? Why is Humility 'ominous'? Who ever promised me that I was capable in principle of knowing everything?" (this volume, 211). If there is one thing that all parties to the dispute over Humility are agreed to, it is that Humility is probably the most benign form of ignorance one could have. Not knowing which properties realize T is nothing, to say the least, like not knowing that one has hands, or not knowing that one is talking to other people.

What has been less than clear, I think, is exactly *why* Humility is less ominous than ordinary skepticism. I think that things here have been confused because there are two ways in which Humility is less ominous than ordinary skepticism. The first, and fairly well-recognized way, is that Humility claims that we will never know something that we never believed we did know (consider again the deductionist response from above). The ordinary skeptic, on the other hand, claims that we do not know things that we certainly do believe we know.

However, it seems to me that there is a second, and perhaps more important way in which Humility is less ominous than ordinary skepticism. A full discussion of this reason will have to wait for another time, but for now, let me just give a taste. Ignorance can be ominous either because knowledge is good for its own sake, or because knowledge is useful. Insofar as knowledge is good for its own sake, then there's just no getting around it: Humility implies that we will forever lack whatever intrinsic value there is to knowing what Humility says we will never know. However, it seems that Humility (unlike traditional skepticism) is in no way ominous in the *other* sense: the knowledge denied by Humility is utterly and completely *useless* knowledge.¹⁵ The reason for this is the same as the reason why Humility is immune to the standard responses to traditional skepticism: Humility postulates ignorance of propositions that can be expressed, according to Lewis, *only* via the use of rigidified descriptions. Given the nature of the decision problems we face—in particular, given that we never differentiate actions or differently value outcomes that can be distinguished only via rigidified descriptions—it seems that learning the proposition that Humility claims we will never know is, in principle, quite useless.

Acknowledgments

I would like to thank Jim Joyce, Allan Gibbard, Andy Egan, Ivan Mayerhofer, Jonathan Schaffer, David Braddon-Mitchell, Robert Nola, Aaron Bronfman, Boris Kment, Howard Nye, Shamik Dasgupta, Elka Shortleeve, and all the participants in the 2007 NYU/Columbia Graduate Student

Philosophy Conference and the 2007 Southeast Graduate Philosophy Conference, for their thoughtful comments on earlier drafts of this essay.

Notes

1. In this stage, Lewis actually gives three arguments for the claim that T has multiple possible realizations. For the sake of simplicity, I present only what seems to be the least controversial argument here—the so-called permutation argument.

However, I would like to briefly point-out what seems to be a mistake in another of Lewis's argument—namely, the replacement-by-idlers argument (this volume, 212–214). An idler is a property that is instantiated in our world yet "plays no role in the workings of nature." As such, idlers will not be mentioned by the theory T. Lewis gives an argument similar to the permutation argument, but which involves replacing members of the actual realization of T with idlers (rather than permuting them) to form a distinct possible realization of T. Lewis claims that as a modal realist, he has plenty of reason to believe in the existence of idlers, but he admits that actualists have little reason to believe in them. Nevertheless, Lewis asks

But how much reason has he [the actualist] to disbelieve in idlers? I grant that Occam's razor justifies him in having a low degree of belief in idlers. But can he claim to know that there are no idlers? I think not. To say they don't exist because we can have no reason to believe they do exist seems nothing better than an appeal to verificationism.

The upshot is that whatever you think about the metaphysics of modality, you should at least agree that for all you know, there may be an abundant supply of idlers. So for all you know, the replacement argument using idlers may succeed. But that is good enough. Humility is, after all, a thesis of irremediable ignorance. If we are irremediably ignorant about whether the replacement argument using idlers establishes Humility, then either the argument does establish it, and it is true; or else the argument does not establish it, but we cannot know that. In both cases alike, we are irremediably ignorant about the identities of the fundamental properties that figure in the actual realization of the true final theory (this volume, 213–214)

The move at the end seems to be a mistake: even if we cannot know that the replacement argument does not establish Humility, it does not follow that we are irremediably ignorant about the identities of the fundamental property that figure in the actual realization of the true final theory. If there are no idlers, then, setting aside the other arguments for the moment, there is only one possible realization of T—regardless of whether we *know* that there are no idlers and so regardless of whether we *know* that there is only one possible realization of T. But if there is only one possible realization of T, then any world where T is realized is a world where the actual realization of T realizes T. Thus, since Lewis allows that we may come to know that T is realized, it follows that we may come to know that the actual realization of T is the realization of T.

The trouble with Lewis's argument is that he seems to be assuming some sort of "KK principle"—i.e., that knowing that P implies knowing that one knows that P. This principle holds on some accounts of knowledge, but not on others. Interestingly enough, this principle does not hold on Lewis's own account of knowledge

(see below), since on that account what one knows depends in part upon what possibilities *there are* (e.g., whether there is possible world where T is realized by some property that is actually an idler) regardless of whether one *knows* that these are the possibilities.

2. My definition of 'quidditism' diverges from that of Robert Black (2000) which is the definition cited by Lewis in "Ramseyan Humility." Black defines 'quidditism' as the thesis that "there is primitive identity between fundamental properties across possible worlds" (Black 2000, 92). However, as Lewis says, quidditism is to be to properties as haecceitism is to individuals. Accordingly, I have constructed my definition of 'quidditism' to be analogous to Lewis' definition of 'haecceitism' in (1986a), where he was at pains to *not* identify haecceitism with the thesis that "there is primitive identity between individuals across possible worlds". I believe my definition also squares better with the remainder of Lewis's discussion of quidditism in "Ramseyan Humility" (see in particular this volume, 209–212).

3. I here allows Lewis's contextualist *sotto voce* proviso ("Psst!—except for the possibilities that we are properly ignoring") to remain *sotto voce*. As I argue below, the appeal to contextualism won't help in blocking the argument for Humility.

4. Here and henceforth I follow Lewis in assuming that propositions are sets of possible worlds (or, more generally, sets of *possibilia*). No doubt much of what I say in this paper would need to be revised to take on the assumption that propositions are more finely individuated.

5. The 'nonambitious' qualifier on two-dimensional semantics comes from Soames 2005. The two-dimensionism appealed to here is nonambitious in that it only claims that *some* statements (in particular, those containing the term 'actually') have two intensions (please see the next note).

6. The labels 'C-intension' and 'A-intension' are the ones used by Frank Jackson in his particular version of two-dimensional semantics (see, e.g., Jackson 1998a). In other versions, different labels have been used. My choice to use Jackson's labels should not be taken as an endorsement of his particular view of two-dimensional semantics. Indeed, nothing more than the relatively conservative two-dimensionalism of Davies and Humberstone (1982) is required for what I say here.

7. See, e.g., Aloni 2005, Gerbrandt 1997, and van Rooy 2003.

8. I should note that Whittle does believe that *some* epistemic conclusion follows from Lewis's argument—namely, that we do not know the "intrinsic natures" of the properties that realize T. I object to this reading of Lewis's argument on the grounds that *fundamental* properties do not *have* intrinsic natures, and thus there simply isn't anything there to know. But of course this subject surely needs more treatment than I can give it here.

9. However, I do believe that Schaffer is right in that *one* of the responses to external-world skepticism (direct realism) does stand a chance at refuting Humility. I will discuss this potentially successful response in section 6.

10. When I speak of 'the mass role', I mean the role of being the *i*th member of the realization of T, where 'mass' is the *i*th T-term in the postulate of T. If 'mass' is not a T-term of T, then Schaffer's example will need to be adjusted accordingly.

11. This isn't circular, just shorthand. Please see the previous note.

12. We can of course distinguish the hypotheses in trivial ways such as *the proposition expressed by such and such sentence vs the proposition expressed by such and such other sentence*. But I take it that those distinctions are irrelevant to which is the better explanation.

13. For my part, I don't yet wish to take a position on this assumption. My point here is merely that (most of) the traditional responses to skepticism will not work as responses to Humility *unless* they are coupled with responses to this semantic thesis.

14. What if we're not talking about realizations of T? I admit that there might be some contexts in which we are properly ignoring possibilities in which T has some alternative realization. However, as soon as we even *mention* Lewis's thesis, we are no longer ignoring those possibilities and so it is true. It might be responded that the contextualist response to traditional skepticism works in the same way: as soon as we mention the traditional skeptic's thesis that we do not know that we are not brains in vats, the possibility that we are brains in vats has become salient and so the skeptic's thesis is true. However, the traditional skeptic *also* claims, for example, that we do not know that we have hands. But just mentioning *that* thesis does not put us in a context in which the possibility that we are brains in vats has become salient. Thus, the traditional skeptic sometimes makes claims that are refuted by contextualism *when he is making them*. Not so in the case of Humility.

15. Again, this isn't to say that it isn't intrinsically valuable to have such knowledge.