

Three Arguments For Humility

Ramseyan humility is the thesis that we cannot know which properties realize the nomic roles specified by the laws of completed physics. Lewis' argument for this thesis looks like that of an evidential sceptic, in which he first argues that on a Humean conception of fundamental properties, there are multiple possible realizations of completed physics, which by definition accommodates all the available evidence. He concludes, on the assumption that empirical knowledge requires evidence, that we cannot know which of these realizations is actual. However, Lewis also appeals to a range of familiar semantic principles when framing his argument, which leads some authors to suppose that he also offers a purely semantic argument for humility. Getting clear about this argument is important not just for Lewis scholars, but also for those who embrace anti-Humean ontologies in which fundamental properties have their causal roles essentially. The epistemological consequences of Humeanism are among the primary reasons given by anti-Humeans for rejecting it,¹ and Lewis' is by far the most detailed account of those consequences.

I shall not attempt to defend a particular interpretation of Lewis, and grant that he offers both sceptical and semantic arguments for humility. My purpose is to argue that Lewis is also committed to principles that license a purely *metaphysical* argument for humility. I do not claim that the metaphysical argument I offer is what Lewis had in mind—it is not—but given his commitments in the metaphysics of mind, it is a natural argument to run on his behalf. It is also, I shall argue, a significantly better argument than the others. The plan of the paper is as follows. In §1 I give a standard sceptical interpretation of Lewis, and in §2 discuss the semantic arguments suggested on his behalf by Leuenberger and Kelly.² In §3 I detail the metaphysical

¹ See for instance Shoemaker (1980); Whittle (2006); Bird (2007).

² Leuenberger (2010); Kelly (2013).

argument for humility and its relation to Lewisian metaphysics of mind; and in §4 I argue that this metaphysical argument is the strongest of the three.

1. A Sceptical Argument for Humility³

Suppose we have a true and complete final theory of everything, and label it T. The language of T, Lewis assumes, has two kinds of term: T-terms and O-terms. The T-terms refer to fundamental properties, and there are no such terms referring to either alien (uninstantiated) or idle (instantiated but inert) properties. The T-terms are those that are introduced by T and implicitly defined by it, and the O-terms are those that are understood independently of T. The O-language is taken to be rich enough to express every possible observation, but fundamental properties are not named in it except as the occupants of roles. Because T is complete, there is a T-term for every instantiated fundamental property that is not an idler. T consists of all the logical consequences of a sentence Lewis terms the *postulate* of T, which may be written: $T[t_1, \dots, t_n; o_1, \dots, o_m]$. We get the Ramsey-sentence R(T) of T by replacing the T-terms t_1, \dots, t_n with bound variables: $\exists x_1, \dots, x_n T[x_1, \dots, x_n; o_1, \dots, o_m]$. R(T) logically implies all and only those O-language sentences that are logical consequences of the postulate of T, from which it follows that an observation confirms T iff it confirms R(T).

R(T) says that there is at least one n-tuple that satisfies the open sentence $T[x_1, \dots, x_n; o_1, \dots, o_m]$. Following Lewis, call such an n-tuple a *realization* of T. If there are *multiple* possible realizations of T, no evidence could bear on the question of which one is actual. Because T is complete, it is capable of accommodating any evidence we could possibly gather, but R(T) accommodates that evidence just as well, and R(T) is true whichever realization of (T) is actual. It follows that if T is multiply realizable, which n-tuple actually realizes T is empirically

³ I focus on Lewis' *permutation argument* for brevity, but my conclusions apply to the other arguments he offers.

undecidable. Why suppose T is multiply realizable? According to Lewis' *combinatorialism*, we can 'take apart the distinct elements of a possibility and rearrange them', to yield a possibility, which entails that permutation of fundamental properties yields a possibility.⁴ According to *quidditism*, fundamental properties stand in primitive transworld identity relations, so permutation of such properties results in a *distinct* possibility. As Lewis has it, '[t]wo different possibilities can differ *just* by permutation of fundamental properties. *They do not differ in whether T is realized, or in what we observe*'.⁵

I need not argue here that Lewis is an evidential sceptic, for my purposes in this paper are not exegetical; I shall instead argue that Lewis is committed to an epistemic premise from which a sceptical argument can be constructed. In order to permute two fundamental properties, we must swap both their nomic roles as given by T, and their pattern of instantiation throughout spacetime.⁶ Depending on what position we take on laws of nature, our stock of fundamental properties may need to include second-order lawmaking relations between first-order properties.⁷ In order to fix ideas, it will be helpful to employ Leuenberger's notion of *fundamental structure*, which will inform much of the discussion to follow.⁸

Consider two worlds w_1 and w_2 , where F_1 and F_2 are the sets of fundamental properties and relations instantiated at each world respectively, D_1 and D_2 their domains of fundamental individuals. Adapting Leuenberger's account so as to allow for fundamental second-order lawmaking relations, I shall say that w_1 and w_2 have the same fundamental structure iff there exists a 1-1 function f from $D_1 \cup F_1$ to $D_2 \cup F_2$ such that (a) for every $x \in D_1$ and every $X \in F_1$: x

⁴ Lewis (2009), p. 208.

⁵ Op. cit., p. 209 (all my italics).

⁶ Op. cit., pp. 207-8.

⁷ Armstrong (1983); see also Lewis (2009), nn. 11-12.

⁸ Leuenberger (2010), pp. 331-2.

has X at w_1 iff $f(x)$ has $f(X)$ at w_2 , and (b) for every $(X, \dots, X_n) \in F_1$: $L_n(X, \dots, X_n)$ at w_1 iff $L_n(f(X), \dots, f(X_n))$ at w_2 , where ‘ L_n ’ denotes an arbitrary n -ary lawmaking relation between fundamental properties. By way of illustration, consider a toy world w_1 at which just two simple fundamental properties A and B are instantiated at various locations throughout spacetime, and let the only fundamental law be that A s cause B s. The distribution of A and B and the law that A s cause B s define the roles of A and B at w_1 . Now consider a world w_2 such that every instance of A is replaced by an instance of B , and *vice-versa*; and at which the only law is that B s cause A s. Intuitively, w_1 and w_2 differ as to which properties occupy the distributional and nomic roles, but not as to which roles are occupied. It is this similarity between w_1 and w_2 that the notion of fundamental structure (hereafter abbreviated ‘FS’) is intended to capture.

The conjunction of combinatorialism and quidditism entails that there is a distinct possible world with the same FS as actuality, at which mass and charge are permuted as to their roles. However, that a world w has the same FS as actuality does *not* entail that $R(T)$ is true at w —for that Lewis needs the additional premise that *evidence supervenes on FS*. Conversely, if $R(T)$ is true at a world w , then w not only has the same FS as actuality, but also the same total evidence. Let me explain why this is so. The theorems of T include O -language sentences truly describing the totality of available evidence. Now evidence, for Lewis, is perceptual experience or memory thereof, and its epistemic role consists in the elimination of possibilities. Lewis maintains that “...a possibility w is *uneliminated* by [S ’s] perceptual experience and memory iff [S ’s] perceptual experience and memory in w exactly match [S ’s] perceptual experience and memory in actuality”.⁹ On this account, a possibility w is eliminated by S ’s experience E iff ‘ w is a possibility in which [S] is not having E ’.¹⁰ Now suppose for the sake of argument that the

⁹ Lewis (1996), p. 553.

¹⁰ Op. cit. p. 553. I draw the reader’s attention to the fact that Lewis explicitly uses the locution *having E*. I shall return to this in §3, where I consider the relationship between Lewisian evidence and metaphysics of mind.

phenomenal character or intentional contents of perceptual states are partially determined by the primitive identities of the fundamental role occupants.¹¹ It follows that there is a world w with the same FS as actuality, but at which my perceptual experiences do *not* exactly match my actual experiences, on account of having different phenomenal character or content. Hence, my actual experiences rule out w , and the available evidence differs between actuality and w .

Because the O-language is by hypothesis rich enough to describe all the available evidence, actuality and w differ as to O-language truths. It follows that $R(T)$ is false at w , so T is not realized there. A world w 's having the same FS as actuality therefore fails to guarantee the realization of T at w . The crucial point is this: T is realized at all possible worlds with the same FS as actuality only if evidence supervenes on FS. Since Lewis clearly *does* think that T is realized at all possible worlds with the same FS as actuality, it follows that he is committed to the supervenience of evidence on FS. We must therefore distinguish the following claims: (1) the FS of actuality is multiply realizable; (2) T is multiply realizable.

The supervenience of evidence on FS enables us to give a straightforward sceptical argument for humility. The conjunction of combinatorialism and quidditism entails that there are distinct worlds with the same FS as actuality. Given that evidence supervenes on FS, this entails that there are other possible worlds at which the totality of available evidence does not differ from actuality. On the assumption that knowledge requires evidence, Ramseyan humility follows. On this reconstruction, the Ramsey-Lewis semantics for theories are seen merely as encoding the key sceptical premise via the claim that T is realized at all worlds with the same FS as actuality, which in turn presupposes that evidence supervenes on FS. Once that presupposition is exposed, however, there is no obvious reason to express the argument in semantic terms at

¹¹ I return to these issues in §§3-4.

all. The argument presented above seems to be a typical sceptical argument stated in semantic terms. For that reason, some authors suppose that humility can be defused by traditional anti-sceptical rebuttals.¹² However, were this the only argument on offer, it would be difficult to make sense of Lewis' reliance on semantics to state it; accordingly, some authors suppose that even if Lewis does offer a sceptical argument for humility, that cannot be the whole story.

2. Semantic Arguments for Humility

In addition to remarks that suggest the kind of sceptical reasoning discussed in §1, Lewis also makes remarks that suggest a further argument for humility based on gaps in our knowledge of *meaning*. Here is a key passage:

There are alternative answer-propositions, to be sure. For each fundamental property F...there is a contingent proposition true at all and only the worlds where F occupies the role in question. But we do not have alternative answer-sentences that express those alternative answer-propositions and do so in such a way that we can know which sentence expresses which proposition.

“Which property occupies the role?—The occupant of the role, whatever that is.” A true answer, sure enough, but not an answer to the question we meant to be asking. Indeed, not an answer to any question we'd be likely to ask: the only information it conveys is that the role is uniquely occupied.¹³

Lewis clarifies these remarks by appealing to two-dimensional semantics. The answer-sentence ‘charge occupies the charge role’ has the primary intension, given the Ramsey-Lewis account of theoretical terms, that the unique actual occupant of the charge role occupies the charge role. This proposition, however, is true at any world, considered as actual, at which the charge role is uniquely occupied. The proposition we need to know in order to know which property occupies the charge role is a proposition that is true at our world but false at worlds with the same FS as actuality but a different occupant of the charge role—the *secondary* intension of

¹² Schaffer (2005) argues this way, and Leuenberger (2010) concurs; see Locke (2009) for opposition.

¹³ Lewis (2009), p. 216.

‘charge occupies the charge role’. Hence, although we know which answer-sentences are true, we cannot know which answer-propositions those sentences express. Leuenberger and Kelly attribute to Lewis distinct semantic arguments for humility based on conjoining semantic principles with combinatorialism and quidditism. I cannot do full justice to these interpretations here for reasons of space, but a brief summary will help to facilitate the discussion of §4.

Leuenberger takes Lewis’ argument to depend, *inter alia*, on the following premises:¹⁴

Expressibility: If p is knowable, then p is expressible in O.

Structuralism: If p is expressible in O, then p supervenes on FS.

Combinatorialism: Fundamental properties are recombinable in such a way that there are distinct possible worlds that have the same FS.

In the present context, a proposition is *expressible* in a language L iff it is the *primary* intension of a sentence of L: ‘expressible’ is thus shorthand for the notion ‘expressible in such a way that we can know which proposition is expressed.’ I have simplified Leuenberger’s *Expressibility* premise, which has ‘ p is entailed by a proposition that is expressible in O’ as consequent; nothing turns on this for present purposes. *Expressibility* states that we can only know a proposition p if p is the primary intension of an O-language sentence. Now a proposition supervenes on FS iff its truth value cannot differ between worlds with the same FS. *Structuralism* is therefore the claim that worlds with the same FS do not differ as to O-language expressible truths. The only differences between worlds with the same FS are differences in fundamental role occupancy, so because the O-language refers to fundamental properties only as the occupants of roles, it follows that it lacks the resources to express any propositions whose truth-values differ between worlds with the same FS, including the answer-propositions.

¹⁴ Leuenberger (2010), p. 330; simplified for exposition.

Leuenberger's *Combinatorialism* follows from the conjunction of Lewis' combinatorialism and quidditism. Because he is concerned with Lewis' replacement argument, Leuenberger focuses on the possibility of alien properties occupying the actual roles, but that need not concern us. What is important is that given combinatorialism and quidditism, there are worlds with the same FS as actuality but different fundamental role occupants. This in turn entails that the answer-propositions do not supervene on FS. By *Structuralism*, the answer-propositions are not expressible in O; from *Expressibility*, it follows that they are not knowable.

Leuenberger is not sympathetic to the argument he finds in Lewis. *Expressibility*, he suggests, requires a commitment to the additional claims (i) if p is knowable, then p is entertainable, and (ii) if p is entertainable, then p is expressible in O.¹⁵ The conjunction of (ii) with *Structuralism* entails that we can only entertain propositions that supervene on FS. Leuenberger's case against Lewis then depends on constructing counterexamples to this latter claim. What is important for my purposes is the fact that Leuenberger's reconstruction depends on principles restricting our ability to cognize the answer-propositions: we can only know what we can entertain, we can only entertain the primary intensions of O-language sentences, but these latter supervene on FS, so we cannot know the answer-propositions, which do not so supervene.

Like Leuenberger, Kelly thinks that Lewis offers an argument for humility based on our limited knowledge of meaning.¹⁶ Kelly gives a Lewisian account of propositional grasp, and argues on this basis that we cannot grasp the answer-propositions. Kelly's account of grasp is based on Lewis' two-dimensionalist distinction between knowing that a sentence expresses a true proposition, and knowing which proposition it expresses. Kelly suggests that we grasp the

¹⁵ Again, simplified. See Leuenberger (2010), pp. 338-41 for details.

¹⁶ Kelly (2013).

proposition expressed by a sentence iff we are ‘...able to evaluate the proposition at any given counterfactual world, given all there is to know about that world...with the proviso that the representation of facts about the counterfactual world must not include any indexical reference to the actual world’.¹⁷ Take the primary intension of the concept water to be ‘the unique actual occupant of the water-role’. I know the proposition expressed by ‘water occupies the water-role’ iff I know that H₂O is the unique actual occupant of the water-role, and am therefore able to evaluate its truth at counterfactual worlds without indexical reference to actuality: for any world *w* considered as counterfactual, the proposition expressed by ‘water occupies the water-role’ is true at *w* iff H₂O occupies the water-role at *w*. Those who do not know that water is H₂O cannot assign truth-conditions in this way, and so fail to grasp the proposition expressed.

Given that we can only think of fundamental properties as the unique actual occupants of their various causal roles, it follows from the two-dimensional account of grasp that we cannot grasp the answer-propositions. Grasping an answer-proposition requires that we can *identify* its referent—in a strong Lewisian sense that requires knowledge of essence¹⁸—but that is impossible in the case of fundamental properties, which we can only know via their inessential causal roles. On the assumption that knowing *p* requires grasping *p* in the relevant sense, humility follows. It is worth noting that if H₂O is a *structural* kind composed of elements with certain fundamental properties, and standing in certain relations to each other, then knowing that water=H₂O is at most *partial* knowledge of the essence of water. As Lewis recognises,

¹⁷ Kelly (2013), p. 718.

¹⁸ Here is Lewis on identification: ‘I spoke of “an uncommonly demanding and literal sense of ‘knowing what’”. Let me elaborate. I say that according to the Identification Thesis, the knowledge I gain by having an experience with quale Q enables me to know what Q is—identifies Q—in this sense: any possibility not ruled out by the content of my knowledge is one in which it is Q, and not any other property instead, that is the quale of my experience. Equivalently, when I have an experience with quale Q, the knowledge I thereby gain reveals the essence of Q.’ Lewis (1995), p. 142. Lewis rejects the identification thesis for qualia, but there is evidence that he takes identification of the referent(s) to be necessary for singular propositional grasp, which is correspondingly very rare; see for instance the resolution offered to Kripke’s puzzle about belief in Lewis (1981).

humility *spreads* from fundamental properties to structural compounds thereof.¹⁹ Knowing which compound H₂O is seems to require that we know at least some of the answer-propositions. Still, perhaps we know *enough* about water—the purely structural aspects of its nature, for example, assuming spatiotemporal relations are not themselves subject to humility—to grasp the relevant propositions in the way Kelly suggests. Or perhaps our grasp of such propositions, like our knowledge of H₂O, is partial. Or perhaps H₂O is a functional kind individuated above the level of fundamental physics, and beyond the reach of humility.

I need not pursue these matters further here. Suffice it to say that there are clear affinities between the semantic interpretations of Lewis we find in Leuenberger and Kelly. Both suggest arguments based on conjoining conditions on propositional grasp with combinatorialism and quidditism. Leuenberger's Lewis claims we can only entertain propositions that are the primary intensions of O-language sentences, but those propositions supervene on FS; Kelly's Lewis has it that grasping a proposition requires at least some knowledge of the essential nature of its referent(s), which is impossible in the case of fundamental properties. I think both these interpretations are plausibly Lewisian, and there is clearly much more to be said about them, and the relationship between them. I need not say it here, however, for I have already said enough to facilitate the discussion to follow.

3. A Metaphysical Argument for Humility

I shall now suggest a third argument for humility that is neither sceptical nor semantic; because it is based solely on Lewisian metaphysics—combinatorialism, quidditism and metaphysics of mind—it is natural to refer to it as a metaphysical argument. Let us refer to the supervenience of a set of properties or facts on FS as their *structural supervenience*. As Leuenberger argues,²⁰

¹⁹ Lewis (2009), pp. 214-5.

²⁰ Leuenberger (2010), pp. 344-5.

structural supervenience must not be confused with Lewis' familiar doctrine of *Humean* supervenience. Here is a canonical expression of Humean supervenience:

I hold, as an *a priori* principle, that every contingent truth must be made true, somehow, by the pattern of coinstantiation of fundamental properties and relations....If two possible worlds were exactly isomorphic in their patterns of coinstantiation of fundamental properties and relations, they would thereby be exactly alike *simpliciter*.²¹

It follows from this definition that the pattern of coinstantiation of fundamental properties at a world is not its FS: *not* all worlds isomorphic as to FS are “exactly alike *simpliciter*”, for they differ as to which properties realize their FS. Patterns of coinstantiation are partially determined by the *identities* of the fundamental properties. Lewis' appeal to pattern *isomorphism* as a sufficient condition on the identity *simpliciter* of worlds is somewhat infelicitous. Each of *n* distinct worlds with the same FS has *a* pattern of coinstantiation of fundamental properties and relations, and it is surely correct to describe the *n* patterns as isomorphic, despite the fact that they are different patterns. Suffice it to say that structural supervenience is not Humean supervenience. Crucially, Humean supervenience is consistent with, and entailed by, structural supervenience. A set of facts that supervenes on the coarse-grained FS *a fortiori* supervenes on the fine-grained Lewisian pattern, because structural difference entails difference of pattern.

Lewis clearly holds that mental facts supervene on the fine-grained pattern of instantiation of fundamental properties, because he holds that everything so supervenes. However, as noted above, this does not rule out the structural supervenience of mental properties. I shall now argue that there is very good reason to suppose that Lewis endorses this latter supervenience thesis. For reasons of exposition, I shall first draw attention to an apparent tension between Lewis'

²¹ Lewis (1994), p. 292.

metaphysics of mind and his views on evidence, then argue that the structural supervenience of mental properties is the key to resolving it.

For Lewis, folk psychology is a theory, with mental state terms such as ‘pain’ and ‘belief’ its theoretical terms, defined by their places in it. The Ramsey sentence of folk psychology enables us to provide a list of analytic truths of the form: mental state M = the occupant of the M -role. The occupant of the M -role turns out, on empirical investigation, to be brain state B , so it is contingently true that $M=B$.²² Now recall that for Lewis, a possibility w is eliminated by S 's perceptual evidence E iff w is a possibility in which S is not having E . Given that experiential states are identified with physical states, however, it seems as though our actual experiences eliminate worlds isomorphic as to FS at which the fundamental roles are differently occupied. In such worlds, the brain states that occupy our counterparts' folk psychological roles cannot be type identical to ours, because they involve the instantiation of different fundamental physical properties. It seems then that our experiences eliminate such possibilities, for they are possibilities in which those experiences would be different. However, this is contrary to Lewis' explicit claim that our perceptual evidence *fails* to rule out such worlds.

The key to understanding why there is no contradiction here lies in the notion of *having E*, to which Lewis appeals when giving his account of the conditions under which experiences rule out possibilities. By way of illustration, let us suppose, with Lewis, that human pain = C-fibre activity, whereas Martian pain = foot cavity inflation.²³ Because C-fibre activity is not identical to foot cavity inflation, does it not follow right away that human pain is not identical to Martian pain? It is natural to suppose that Lewis is identifying the mental property of feeling painful

²² See for instance Lewis (1970, 1980, 1994b).

²³ Lewis (1980).

with C-fibre activity in humans, and with foot cavity inflation in Martians, which is why he is often miscast as a type identity theorist. Lewis' theory is actually far more nuanced than that, for he does *not* identify the distinctively mental properties of a mental state with its physical properties. Human pain and Martian pain *feel the same*, on Lewis' theory, despite the fact that they are not the same state. Here is an oft-quoted passage from Lewis:

The definitive characteristic of any (sort of) experience as such is its causal role, its syndrome of most typical causes and effects. But we materialists believe that these causal roles which belong by analytic necessity to experiences belong in fact to certain physical states. Since those physical states possess the definitive characteristics of experience, they must be the experiences.²⁴

When Lewis says that the definitive characteristic of pain *qua* experience is its syndrome of typical causes and effects, he does not mean to deny that the way pain feels is definitive of pain. Rather, his claim is that *the way pain feels* is defined in terms of a characteristic causal role. In humans, C-fibre activity has the definitive characteristic of pain; in Martians, foot cavity inflations have it. This, however, is just to say that in humans, C-fibre activity feels like pain, whereas in Martians, foot cavity inflations feel like pain. The phenomenal character of human and Martian pain states are the same, despite the fact that the states themselves are not identical. The states in question have the definitive character of pain in virtue of their contingent causal roles, which is to say they *feel like pain* contingently. What makes human pain the state it is—what individuates it *qua* state—is the physical properties that constitute C-fibre activity. However, what makes it *pain* is that in us, C-fibre activity occupies the pain-role. Lewis refers to the second-order property that humans and Martians share—being in *some* state that occupies the pain-role—as *being in* or *having* pain. It is easily overlooked that Lewis' theory identifies *the way pain feels*—although not, of course, pain itself—with this shared functional

²⁴ Lewis (1966), p. 17.

property.²⁵ Lewis is a type identity physicalist about pain *qua state*, which is to say he regards mental states as individuated by their physical properties. However, he is a functionalist about the distinctively mental properties—irrelevant to their individuation *qua states*—that such states sometimes possess, in virtue of which they count as mental.

Lewis' functionalism about mental properties covers intentional content as well as phenomenal character. Lewis summarises his position on intentional content as follows:

The contentful unit is the entire system of beliefs and desires.... That system is an inner state that typically causes behaviour, and changes under the impact of perception (and also spontaneously). Its content is defined, insofar as it is defined at all, by constitutive rationality on the basis of its typical causal role.²⁶

There is much more to be said about Lewis' theory of content than I can say here, but a sketch will suffice for present purposes.²⁷ The folk psychological roles are not just causal, but also rational—the role of the belief that *p* involves not only what typically causes subjects to believe that *p* and how such subjects typically behave given their other beliefs, but also the manner in which the attribution of holistic networks of belief including the belief that *p* enable radical interpreters to make sense of agents. Thus constitutive norms of rationality play a crucial role in determining the contents of belief in Lewis' theory. However, a network of inner states *fills* the folk psychological roles that define a given network of intentional states solely in virtue of their causal roles in relation to each other, perceptual input and behavioural output. To have a network of intentional states, for Lewis, is to instantiate an interpretive theory in virtue of the causal roles occupied by a structurally isomorphic network of inner states. As with sensations,

²⁵ Lewis (1966), p. 19; (1994), p. 307. Lewis identifies mental state M with the *occupant* of the M-role, rather than the common second-order state, on causal grounds: given that M has a causal role, it must be identified with the occupant of the M-role, because only the first-order state is efficacious.

²⁶ Lewis (1994), p. 324.

²⁷ For details see Lewis (1974); for an overview see Lewis (1994).

intentional states are identical to physical states, but they have the mental properties in virtue of which they *count* as intentional states solely in virtue of the roles they occupy.

It is often said, but seldom elaborated upon, that Lewis is a type physicalist about mental state M, but a functionalist about *having* M. We are now in a position to understand precisely what this claim amounts to. Having M, for Lewis, is being in some state that occupies the M-role, which amounts to being in some state with the mental properties that are definitive of M *qua* mental. Now for me to fail to be *having* experience E requires that I am not in *any* state that occupies the E-role. It follows right away that my experiences do not rule out alternative possible realizations of the actual FS. Our counterparts at worlds with the same FS as actuality differ as to the fundamental role occupants, but not as to which roles are occupied, so they do not differ as to which experiences they are *having*. My evidence also fails to rule out sceptical scenarios—for instance, worlds in which my closest counterpart is an envatted brain functionally isomorphic to my own, and which is therefore having all the experiences I am having. This is of course just what Lewis thinks: it is not that our evidence rules out such scenarios, but that certain contexts entitle us to ignore them.

It is no accident that Lewis claims that a possibility *w* is eliminated by *S*'s perceptual experience *E* iff *w* is a possibility in which *S* is not *having* *E*. To my mind, this is a deliberate reference to the functional property of being in some state that occupies the E-role. Lewisian evidence does not depend on how the folk-psychological roles are occupied, but only on the roles occupied, for it is these that determine which experiences our counterparts are having at alternative possibilities, and hence which of those possibilities our own experiences rule out. Mental *properties* determine evidential significance, and these properties supervene on FS.

I shall now offer a metaphysical argument that proceeds directly from the structural supervenience of mental properties to humility. As in §1, I leave open that the FS includes primitive first-order qualities, spatiotemporal relations *and* second-order lawmaking relations. As before, I grant the conjunction of combinatorialism and quidditism, which entails that there are distinct possible worlds with the same FS as actuality. The metaphysical argument is a *reductio*: I assume we know an answer-proposition, and show that this violates structural supervenience. First, note that on any account of knowledge-*wh*, knowing which x is the F is having propositional knowledge. For instance, on the reductive account of knowledge-*wh*, knowing which x is the F is knowing *that* p , where p is the true answer-proposition to the question ‘which x is the F?’ Similarly, on Schaffer’s non-reductive contextualism, knowing which x is the F is knowing *that* **a** rather than any of a set of contextually determined alternatives **b**, **c**, **d**...is the F.²⁸ Either way, having knowledge-*wh* is being in a psychological state with intentional content.

Now consider a world w_p in which mass and charge are permuted relative to actuality *salva* FS. Structural supervenience entails that the contents of our mental states would have been the same had w_p been actual. Now suppose for *reductio* that we know an answer-proposition—for simplicity, let it be the proposition that q_1 occupies the charge role. At w_p , q_1 and q_2 (the actual occupant of the mass role) are permuted. It follows that we know something we would not have known had w_p been actual, because that in that case, q_2 would have occupied the charge role, and the proposition that q_1 occupies the charge role would have been *false*. Now this does not in itself violate the structural supervenience of mental properties. Perhaps the knowledge-ascriptions differ in truth-value solely because of extra-mental differences between actuality

²⁸ Schaffer (2007).

and w_p .²⁹ I need to show that there is a difference in *content* between actuality and w_p , not just in which knowledge-ascriptions are true.

At this point I make a crucial assumption: there is nothing *epistemically privileged* about the way the FS of the actual world is realized compared with w_p . Given the nature of Humean fundamental properties, our stipulative knowledge of the answer-proposition just *cannot* depend on the actual FS being realized by such properties in a specific way.³⁰ We can state this no-privilege thesis as follows: necessarily, for any world w^* with the same FS as actuality, we come to know the true answer propositions iff we would have come to know the answer-propositions true at w^* had w^* been actual. Note that this is *not* simply a way of stating the sceptical premise that evidence supervenes on FS. For present purposes, I am neutral as to the nature of evidence and happy to allow that it can differ between actuality and w_p . The point here is simply that whatever the method by which we come to know that q_1 occupies the charge role, our counterparts at w_p can employ the *same* method (*modulo* any fundamental physical differences in the method that result from the permutation of q_1 and q_2) to come to know that q_2 occupies the charge role. Assuming we come to know which properties realize the actual FS, it follows that had w_p been actual, we would have come to know which properties realized *its* FS in the same way. Had w_p been actual, however, the content of our knowledge would have been different. Because actuality and w_p are isomorphic as to FS, this is contrary to structural supervenience. Hence, we cannot know which properties realize the FS of our world.

²⁹ Williamson (2000) argues against analysing knowledge in terms of truth on the grounds that knowledge is a mental state, whereas truth is *mind-independent*.

³⁰ 'Just about all there is to a Humean fundamental quality is its identity with itself and its distinctness from other qualities. A Humean fundamental quality is intrinsically inert and self-contained,' Black (2000), p. 91.

4. Why Be Humble?

My aim in this section is to assess the relative merits of the sceptical, semantic and metaphysical arguments, but I shall first consider an objection that will help to clarify the relationships between them.³¹ I claim that knowledge *of the answer propositions*—and any others that fail to supervene on FS—is ruled out by the structural supervenience of mental properties. Structural supervenience does not rule out purely descriptive knowledge of fundamental properties, and is consistent with our knowing propositions of the form ‘the unique actual occupant of role *R* occupies *R*’. Suppose it is objected, however, that such knowledge is all it takes to know which property occupies a given role. If we can know the role occupancy facts by knowing propositional contents that supervene on FS, the metaphysical argument does not work. Call this the *descriptivist strategy*.

If knowing that the unique actual *F* is the *F* is all it takes to know which *x* is the *F*, then such knowledge, as Locke points out, seems too cheap.³² Intuitively, if someone tells me that exactly one person invented the zip, they do not thereby tell me *who it was*. Still, let us suppose that the descriptivist strategy is correct. In Lewis’ view, we know that *T* is uniquely realized, which is of course a central component of the Ramsey-Lewis semantics for theoretical terms. It follows that we only need to know which roles are occupied in order to know which properties occupy them. According to the descriptivist strategy, *knowing the answer-proposition is not necessary* for knowing which property occupies role *R*. We do not need to be in a psychological state whose content is the singular proposition that *q*₁ occupies the charge role in order to know which property occupies that role, so the metaphysical argument fails.

³¹ I thank an anonymous referee for the objection in question.

³² Locke (2009).

The descriptivist strategy targets a presupposition of the metaphysical argument: that to know which property occupies role *R*, we need to know the relevant answer-proposition. However, crucially for present purposes, the sceptical and semantic arguments *share* this presupposition. Tellingly, Whittle objects to Lewis' arguments by embracing the descriptivist strategy.³³ Whittle's target is the sceptical argument presented in §1, and she argues against it that provided we know that *T* is uniquely realized, then we know enough to know which properties occupy which roles. Whittle does not deny that evidence supervenes on FS, or that evidence is required for knowledge. In her view, we can know which property occupies the charge role without knowing the singular proposition that *q*₁ occupies it, so it does not matter that we could never gain any evidence for that proposition. The descriptivist strategy tells equally against the semantic argument. In two dimensional terms, it amounts to the claim that knowing the primary intension of 'the unique actual occupant of role *R* occupies *R*' is sufficient for knowing which property occupies *R*. This, however, is just to say that knowing the secondary intension is not necessary, in which case neither is grasping it.

Even if the descriptivist strategy is right about knowledge-*wh*, we can still run a *non*-Lewisian sceptical argument for humility. The descriptivist holds that we know which properties occupy which roles provided we know that the actual FS is uniquely realized, because uniquely identifying descriptive knowledge is sufficient for knowledge-*wh*. An evidential sceptic might then argue for humility by claiming that such knowledge is *necessary* for knowing the role occupancy facts, but that we can never have evidence for unique realization.³⁴ Unlike Lewis' sceptical argument, there is no question that this one is vulnerable to an abductionist reply, for the hypothesis that there are multiple occupants of any given role is clearly more complex than

³³ Whittle (2006).

³⁴ C.f. Bird (2007), pp.77-9.

the rival hypothesis that there is just one occupant per role. Worlds with the same FS, by contrast, do not differ in complexity *at all*.³⁵ Whatever its merits, this kind of sceptical argument is the only remaining route to humility once we deny that knowing the answer-proposition is necessary for knowing the answer to a ‘wh...?’ question.

Three arguments for humility can be distilled from Lewis’ (2009) remarks, given the context provided by his broader philosophical system: sceptical, semantic and metaphysical. Each one proceeds from combinatorialism and quidditism to the conclusion that there are multiple possible realizations of the actual FS, and singular propositions expressing which one is actual. Each then presupposes that knowledge of these answer-propositions is necessary for knowledge of the role occupancy facts. Thereafter they diverge, each offering alternative grounds for thinking we *cannot* know the answer-propositions. It is the additional premises employed at this point that set the arguments apart:

- | | |
|----------------------|--|
| <i>Sceptical:</i> | (i) Evidence supervenes on FS; (ii) Knowing that <i>p</i> requires evidence for <i>p</i> . |
| <i>Semantic:</i> | (i) We cannot grasp the answer-propositions; (ii) Knowing that <i>p</i> requires grasping or entertaining <i>p</i> . |
| <i>Metaphysical:</i> | (i) Mental properties supervene on FS; (ii) Actuality is not epistemically privileged in relation to possible worlds with the same FS but different realizers thereof. |

Assuming the truth of the common premises, I shall now assess the relative merits of these three arguments by considering the entailment relations between their additional premises.³⁶ If the negation of a premise of argument A entails, on reasonable assumption, the negation of at least one premise of argument B, then at least some reasons to reject A are also reasons to reject

³⁵ Similar points are made in Locke (2009).

³⁶ I shall use *Sceptical*(i) and *Sceptical*(ii) to refer to the individual premises listed under *Sceptical*, etc.

B. Conversely, if no such entailment relations hold, then a proponent of B need not worry about whether A's premises are true. I assume throughout that *Metaphysical(ii)* is both true and independent of all the other premises under consideration. Whatever we say about evidence, knowledge, grasp, or the metaphysics of mind, whether or not we have a method of coming to know the answer-propositions cannot depend on *which* n-tuple realizes the actual FS. I will therefore assess only entailment relations between *Sceptical*, *Semantic* and *Metaphysical(i)*. I shall begin by considering whether the negation of any premises of *Sceptical* entail the negation of *Metaphysical(i)*, then the converse; I shall then do the same with *Semantic* and *Metaphysical(i)*. Because there are (at least) three ways for *Metaphysical(i)* to fail, and these are common to both comparisons, I shall outline these ways before proceeding.

How could permutation of fundamental properties *salva* FS make a difference to our mental properties? Firstly, if at least some fundamental properties are essentially mental, as panpsychists suppose, then structural supervenience fails. On a version of panpsychism that Chalmers refers to as *constitutive Russellian panpsychism*, the fundamental physical roles are occupied by microphenomenal properties that also constitute the *macrophenomenal* properties of conscious experience.³⁷ Because microphenomenal properties are categorical—they do not have *nomic* essences—we can permute them *salva* FS; but because they do have *phenomenal* essences, permutation does not preserve the distribution of macrophenomenal properties.

Secondly, one might endorse a form of content externalism according to which the primitive identities of fundamental properties partially determine the intentional contents of mental states concerning them. According to Putnam's natural kind externalism, for example, the contents of our thoughts about water are partially determined by its chemical structure, whether or not

³⁷ Chalmers (2013).

forthcoming in *Philosophical Studies* (see published version:
<http://dx.doi.org/10.1007/s11098-017-0877-6>)

we know that water is H₂O.³⁸ Now suppose we endorse a similar theory of content concerning fundamental properties—never mind how plausible it is. On such a theory, when we think that charge occupies the charge role, the content of our thought is that a specific fundamental property q_1 occupies the charge role. Fundamental property externalism is implausible if it requires the kind of direct perceptual contact with fundamental properties that Putnam supposes we have with water, but perhaps we can avoid this. After all, H₂O is composed of atoms with fundamental properties, so why should content externalism *about water* stop at the level of chemical structure? Perhaps our perceptual contact with natural kinds such as water facilitates directly referring singular concepts for the fundamental properties that they instantiate.

Thirdly, one might embrace a very strong form of type identity theory about mental properties. As we have seen, in Lewis' identity theory, the mental state *human pain* is identical to a physical state, but it counts as pain in virtue of the functionally defined property of occupying the pain role. By contrast, if we identify the mental *property* of feeling painful—the property Lewis refers to as *having pain*—with a structural compound of fundamental physical properties, then phenomenal pain will not supervene on FS. Such a theory might identify the property of being in pain with a complex neural process involving, *inter alia*, the rapid diffusion of charged ions across a membrane. This specific property will not be instantiated at worlds where the charge role is differently occupied. Let us proceed to assess the arguments.

1. Metaphysical vs. Sceptical. If we reject *Sceptical(i)*, then which n-tuple of fundamental properties realizes the FS of *S*'s world makes a difference to *S*'s evidence. One way to make sense of this possibility is direct realism: the evidence depends on the fundamental role

³⁸ Putnam (1975). Putnam's arguments were intended to establish semantic externalism, but extend in a natural way to mental content; see McGinn (1977). I trust the details are sufficiently familiar and omit them for brevity.

occupants because we can directly perceive them via our interactions with the world.³⁹ Suppose *S* has some evidence *E* that fails to supervene on FS. If we suppose *S*'s having *E* is *S*'s being in a mental state with a certain phenomenal character or intentional content—for example: its seeming to *S* to be the case that *p*; *S*'s perceiving, seeing or knowing that *p*; *S*'s remembering that *p*—then it follows right away that mental properties fail to supervene on FS, and *Metaphysical*(i) is false. On any psychological theory of evidence, the negation of *Sceptical*(i) entails the negation of *Metaphysical*(i). Lewis' commitment to *Metaphysical*(i), as I suggested in §3, is arguably behind his commitment to *Sceptical*(i).

What about *Sceptical*(ii)? Unless supplemented with additional premises, Lewis' sceptical argument for humility can be blocked by rejecting evidentialism. There are complications: for instance, abductionists deny that evidence is necessary for knowledge via the claim that in cases where there is one, we can infer to the *best* explanation even when our evidence fails to decide between alternative hypotheses. It is not clear that this helps to avoid humility, because according to Lewis, we cannot even state, let alone compare, the alternative hypotheses in question. The key advantage of the metaphysical argument here is that it is *no* kind of sceptical argument, and depends on no particular theory of knowledge, so we can avoid the controversy over whether this or that anti-sceptical strategy defeats it.⁴⁰

Conclusion 1(a): *Metaphysical*(i) and *Sceptical*(i) plausibly stand or fall together, but rebuttals of *Sceptical* that work by rejecting *Sceptical*(ii) have no impact on *Metaphysical*.

³⁹ Schaffer (2005, pp. 21-2) discusses direct realism as a response to Lewis' sceptical argument.

⁴⁰ Schaffer (2005) argues that traditional anti-sceptical strategies, including abductionism, can be marshalled against Lewis. Locke (2009) responds that Lewis' argument differs from traditional sceptical arguments in ways that block abductionism and other traditional anti-sceptical strategies—although in some cases, Locke appeals to Lewis' semantic principles to make his case.

Let us now consider whether the negation of *Metaphysical(i)* tells against *Sceptical*. It is clear that the negation of *Metaphysical(i)* does not entail the falsity of *Sceptical(ii)*—if mental properties fail to supervene on FS, whatever the reason, that gives us no reason whatever to reject evidentialism. Does the negation of *Metaphysical(i)* entail that *Sceptical(i)* is false? Let us address, in turn, the three ways in which *Metaphysical(i)* may fail. First, panpsychism. On panpsychism, the phenomenal character of our experiences depends not only on which roles are occupied, but also on which fundamental properties occupy them, and so fails to supervene on FS. That being so, there is no obvious reason to deny that we lack evidence for the answer-propositions. Panpsychism does not afford direct access to the natures of fundamental properties, but as long as the identities of the microphenomenal properties make a difference to the character of the macrophenomenal properties they constitute, this is sufficient for us to have evidence that discriminates between alternative possible worlds with the same FS. If *Metaphysical(i)* fails due to panpsychism, so does *Sceptical(i)*.

What if *Metaphysical(i)* fails due to content externalism? Suppose perceptual experiences have intentional contents. If those contents differ between worlds with the same FS, then so, it seems, does the content of our perceptual evidence. Given that intentional contents do not supervene on FS, a proponent of *Sceptical* would need to commit to an additional principle ruling out these content differences as epistemically irrelevant. On a narrowly phenomenal conception of evidence, one could hold both that the intentional contents of mental states depend on the identities of the fundamental properties, but that intentional contents are not part of our evidence. That conception of evidence, however, is deservedly unpopular.⁴¹ On any reasonable conception of evidence, if *Metaphysical(i)* fails due to content externalism, so does *Sceptical(i)*.

⁴¹ Williamson (2000) argues at length against the phenomenal conception of evidence.

What if *Metaphysical*(i) fails due to the type identity theory sketched above? For illustration, consider Lewis' conception of evidence—remembering that this is *not* a Lewisian identity theory—according to which *S*'s evidence *E* rules out possibility *w* iff *w* is a possibility in which *S* is not having *E*. According to our current identity theory, the mental properties of *E*—its phenomenal character and intentional content—are identical to structural properties at least some of whose components are fundamental properties. On this theory, having *E* is not, as on Lewis' theory, being in *some* state that occupies the *E*-role; rather, it is being in a type identical *physical* state, all the way down to fundamental physics. Hence, *all* our experiences are experiences we are not having in worlds with alternative fundamental role occupants. More generally, and once again supposing that having evidence is being in a mental state with a certain phenomenal character or intentional content, if mental properties are identical to structural properties involving the fundamental role occupants, there is *no way* for us to have the *same* evidence at worlds where the fundamental roles have different occupants. If *Metaphysical*(ii) fails due to type identity theory, then so does *Sceptical*(i).

Conclusion 1(b): There are potential rebuttals of *Sceptical* that do not tell against *Metaphysical*, but no rebuttals of *Metaphysical* that are not also rebuttals of *Sceptical*.

2. *Metaphysical* vs. *Semantic*. Both Leuenberger and Kelly attribute to Lewis strong constraints on grasping a proposition, constraints that entail *Semantic*(i). Leuenberger, as we have seen, argues that grasp is not limited to propositions that supervene on FS, from which it follows that the non-supervenience of answer-propositions on FS gives us no reason to endorse *Semantic*(i). To provide reason to reject *Semantic*(i), we could argue for a weaker theory of grasp than the one suggested by Kelly based on Lewisian identification of the referent.

Does the negation of *Semantic(i)* entail the falsity of *Metaphysical(i)*? *No*. The supervenience of mental properties on FS requires only that whatever we grasp around here, we grasp at all worlds with the same FS as actuality. This is perfectly consistent with our grasping all the true answer-propositions, and even all the false ones; and it is also consistent with our not grasping any of these. *Metaphysical(i)* rules out only that we *know* the answer-propositions, because knowledge of those propositions entails differences in mental content between worlds with the same FS. Proponents of the metaphysical argument need not deny that we can grasp the answer-propositions, or embrace any specific account of grasp beyond the general constraint that the graspable contents cannot differ between worlds with the same FS.

Let us turn to *Semantic(ii)*. Someone who places very high standards on grasp, such as Kelly's or Leuenberger's Lewis, might well deny that grasp is necessary for knowledge. Ironically, the chief proponent of an epistemology that suggests the possibility of knowing that *p* without grasping *p* is Lewis himself, whose relevant alternatives theory allows for knowledge without *belief*.⁴² According to Lewis, '*S* knows that *p*' is true in a context *C* iff *S*'s evidence eliminates every possibility in which not-*p*, except for those possibilities *S* can properly ignore in *C*.⁴³ As we have already seen, Lewis holds that *S*'s experiential evidence *E* eliminates a possibility *w* iff *w* is a possibility in which *S* is not having *E*. Provided *E* eliminates a set of contextually determined alternatives to *p*, we know that *p*, regardless of whether we also meet whatever conditions on grasping *p* our theory of grasp imposes. If Lewis' contextualism is inconsistent with *Semantic(ii)*, so much the worse for semantic arguments as interpretations of Lewis.⁴⁴

⁴² Lewis (1996), p. 556.

⁴³ Lewis (1996). I have adopted a metalinguistic formulation to highlight Lewis' contextualism.

⁴⁴ Recognising this difficulty, Kelly (2013, p. 714) argues that Lewis' epistemology is incomplete unless supplemented with a belief clause.

We need not be drawn into such interpretive issues here, for the negation of *Semantic(ii)* does not entail the falsity of *Metaphysical(i)*. The mere possibility of knowledge without grasp has no bearing on the structural supervenience of mental properties. True, Lewis' epistemology allows for low-standards contexts in which it is proper to ignore the alternative possible realizations of the actual FS, and it may be that in such contexts, we know which property occupies which role, because our evidence rules out all the contextually relevant alternatives.⁴⁵ However, such low-standards knowledge explicitly does not require that our evidence rules out alternative possible realizations of the actual FS, and so does not require that our evidence would have been different had some other realization been actual. Lewis' epistemology also seems to allow for knowing that *p* without representing the content that *p* at all, and if we can know which property occupies which role without representing the answer-propositions, then of course we cannot run the metaphysical argument for humility. However, there is plenty of room to deny that *grasping p*—at least insofar as this requires meeting conditions on grasp such as those suggested on Lewis' behalf by Leuenberger and Kelly—is necessary for knowing that *p*, without also denying that *representing p* is necessary.

Conclusion 2(a): *Metaphysical* does not depend on the truth of *Semantic(i)* or *Semantic(ii)*, so there are several ways of rebutting *Semantic* that do not tell against *Metaphysical*.

Finally, let us consider whether the negation of *Metaphysical(i)* entails the falsity of the premises of *Semantic*. The failure of mental properties to supervene on FS has no direct bearing on whether propositional grasp is necessary for knowledge, so I shall focus mainly on the question of whether failures of supervenience entail that we can grasp the answer-propositions:

⁴⁵ See Langton (2004), Schaffer (2005) and Locke (2009) for discussion of contextualist replies to humility.

does the negation of *Metaphysical(i)* entail the falsity of *Semantic(i)*? I shall once more address the three ways in which *Metaphysical(i)* might fail, in turn, beginning with panpsychism.

As we have seen, given panpsychism, it is plausible that we have at least some epistemic access to the identities of the microphenomenal properties that constitute our conscious experiences. Lewis acknowledges that panpsychism refutes humility by making room for the identification of fundamental properties.⁴⁶ Assuming his argument for humility to be valid, Lewis then suggests that panpsychists must reject one of two components of the Ramsey-Lewis semantics: either (1) the O-language names the fundamental properties only as the occupants of roles, or (2) the O-language suffices to express all possible observations.⁴⁷ We could reject (1) by holding that our access to the essences of microphenomenal properties permits us to identify and name them as the very properties they are; or we could reject (2), and deny that the our language has the resources to express the relevant experiences. Either way, there is no principled reason—setting aside the obvious practical difficulties of isolating the individual contributions of microphenomenal properties to conscious experience—why we should not be able to grasp the answer-propositions given the truth of panpsychism.

Let us turn now to the consequences of fundamental property externalism for *Semantic(i)*. The content of my belief that charge occupies the charge role, given this brand of externalism, is the singular proposition that q_1 occupies the charge role. On this view, the Ramsey-Lewis conception of our cognitive access to fundamental properties is false, and with it goes any motivation for denying that we can grasp the answer-propositions. Given externalism, when I think to myself that charge occupies the charge role, the proposition I thereby entertain is the

⁴⁶ Lewis (2009), pp. 217-8.

⁴⁷ Leuenberger (2010, p. 336) considers fundamental phenomenal properties as counterexamples to *Structuralism*, the claim that O-language expressible propositions supervene on FS.

forthcoming in *Philosophical Studies* (see published version:
<http://dx.doi.org/10.1007/s11098-017-0877-6>)

singular answer-proposition, not a descriptive intermediary. One might still endorse a theory of grasp according to which merely entertaining singular propositions in this way is not sufficient for grasping them, but the challenge then would be to explain why this kind of grasp is necessary for knowledge, when it is not necessary for belief. *Semantic(i)* is arguably falsified, and *Semantic* is in any case undermined, by fundamental property externalism.

Finally, the type identity theory. If *Metaphysical(i)* fails because mental properties are identical to structural compounds of fundamental properties, does that entail the falsity of *Semantic(i)*? In a word: *no*. This kind of failure of mental properties to supervene on FS has no bearing that I can discern on the issue of whether we are capable of grasping the answer-propositions. It entails that we would have different mental properties at worlds where the fundamental roles had different occupants, but so what? In the case of both panpsychism and externalism, there are positive reasons to suppose that the *kind* of mental differences between worlds with the same FS are sufficient for grasping the answer-propositions, but there is no corresponding reason to suppose so in the present case. Whether anyone could sensibly endorse an identity theory of this kind is another matter entirely.

Conclusion 2(b): there are several rebuttals of *Semantic* that do not tell against *Metaphysical*, but all rebuttals of *Metaphysical* also serve to rebut *Semantic*, with the exception of one based on an implausibly strong and widely discredited version of the identity theory.

The conclusion of our comparison is that one who rejects the structural supervenience of mental properties will struggle to find a plausible argument for humility. However, those who reject evidentialist theories of knowledge, or who prefer not to commit to theories of propositional grasp, can still run the metaphysical argument.

5. Conclusion

The humility thesis places *a priori* constraints on our knowledge—given the combination of combinatorialism and quidditism with some further principles, we cannot know which of a set of structurally isomorphic possible worlds we occupy. *Which* further principles are need to establish humility? I have granted that Lewis offers distinct sceptical and semantic arguments for humility based on the addition of principles concerning evidence and propositional grasp, respectively. I have suggested, however, that a hitherto unappreciated metaphysical argument can also be extracted from his writings. The additional principle in this argument is that mental properties supervene on fundamental structure. Given Humean fundamental properties, there can be nothing epistemically privileged about the way the fundamental structure of our world is realized; the supposition that we could come to know which properties actually realize it therefore entails differences in content across structurally isomorphic worlds. Assuming the structural supervenience of mental content, humility follows, whatever theory of knowledge we endorse, and whatever we think it takes to grasp a proposition.⁴⁸

⁴⁸ This work was funded by an FCT Investigator grant from the Fundação para a Ciência e a Tecnologia (IF/01736/2014), and is based in part on research carried out while I was a postdoctoral fellow at Oxford, funded by the European Research Council. I am grateful to David Chalmers, David Papineau, Jonathan Schaffer, Célia Teixeira, Jessica Wilson and several anonymous referees for very helpful discussion and criticism.

REFERENCES

- Armstrong, D. (1983). *What is a Law of Nature?* Cambridge: Cambridge University Press.
- Bird, A. (2007). *Nature's Metaphysics*. Oxford University Press.
- Black, R. (2000). 'Against Quidditism', *Australasian Journal of Philosophy* 78, pp. 87-104.
- Braddon-Mitchell, D. and Nola, R. (eds.) (2009). *Conceptual Analysis and Philosophical Naturalism*. Cambridge MA: M.I.T. Press.
- Chalmers, D. (2013). 'Panpsychism and Panprotopsyism', *The Amherst Lecture in Philosophy* 8.
- Kelly, A. (2013). 'Ramseyan Humility, Scepticism and Grasp', *Philosophical Studies* 164, pp. 705-26.
- Langton, R. (2004). 'Elusive Knowledge of Things in Themselves', *Australasian Journal of Philosophy* 82, pp. 129-136.
- Leuenberger, S. (2010). 'Humility and Constraints on O-Language', *Philosophical Studies* 149, pp. 327-54.
- Lewis, D. (1966). 'An Argument for the Identity Theory', *Journal of Philosophy* 63, pp. 17-25.
- (1970). 'How to Define Theoretical Terms', *The Journal of Philosophy* 67, pp. 427-66.
- (1974). 'Radical Interpretation', *Synthese* 23, pp. 331-344.
- (1980). 'Mad Pain and Martian Pain,' in N. Block (ed.) *Readings in Philosophy of Psychology* Vol. I, Harvard University Press (1980), pp. 216-32.
- (1981). 'What Puzzling Pierre Does Not Believe', *Australasian Journal of Philosophy* 59, pp. 283-89.
- (1994). 'Reduction of Mind', first published in S. Guttenplan (ed.) *A Companion to the Philosophy of Mind*, Oxford: Blackwell (1994); repr. in D. Lewis, *Papers in Metaphysics and Epistemology*, Cambridge University Press (1999), pp. 291-324.

forthcoming in *Philosophical Studies* (see published version:
<http://dx.doi.org/10.1007/s11098-017-0877-6>)

—(1995). ‘Should a Materialist Believe in Qualia?’, *Australasian Journal of Philosophy* 73,
pp. 140-4.

—(1996). ‘Elusive Knowledge’, *Australasian Journal of Philosophy* 74, pp. 549-67.

—(2009). ‘Ramseyan Humility’, in Braddon-Mitchell and Nola (eds.) (2009), pp. 203-22.

Locke, D. (2009). ‘A Partial Defence of Ramseyan Humility’, in Braddon-Mitchell and Nola
(eds.) (2009), pp. 223-42.

McGinn, C. (1977). ‘Charity, Interpretation and Belief’, *Journal of Philosophy* 74, pp. 521-35.

Putnam, H. (1975). ‘The Meaning of ‘Meaning’’, *Minnesota Studies in the Philosophy of
Science* 7, pp. 131-93.

Schaffer, J. (2005). ‘Quiddistic knowledge’, *Philosophical Studies* 123, pp. 1-32.

—(2007). ‘Knowing the Answer’, *Philosophy and Phenomenological Research* 75, pp. 383-
403.

Shoemaker, D. (1980). ‘Causality and Properties’, in P. van Inwagen (ed.) *Time and Cause*,
Dordrecht: Reidel (1980), pp. 109-35.

Whittle, A. (2006). ‘On an Argument for Humility’, *Philosophical Studies* 130, pp. 461-97.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.