

Semantic Sovereignty

Stephen Kearns and Ofra Magidor

(Penultimate draft of paper to appear in *Philosophy and Phenomenological Research*

Please cite the final published version)

§1. Introduction

The question we are concerned with in this paper is this: *do semantic facts supervene on use facts?* According to the thesis we label *Semantic Supervenience* (SSUP), semantic facts do supervene on use facts. According to the thesis we label *Semantic Sovereignty* (SSOV), semantic facts do not supervene on use facts.

Semantic Supervenience is widely (indeed almost universally) assumed to be correct. Of course, many philosophers accept the thesis only if ‘use facts’ is given a sufficiently broad interpretation: perhaps one that includes facts about the behaviour of the linguistic community as a whole, or about the background physical environment, or about the relative naturalness of properties.¹ But the commonly accepted assumption is that on some such sufficiently wide interpretation, Semantic Supervenience is correct. As Jerry Fodor puts it:

“I suppose that sooner or later the physicists will complete the catalogue they've been compiling of the ultimate and irreducible properties of things. When they do, the likes of *spin*, *charm*, and *charge* will perhaps appear on their list. But

¹ One might accept that according to most philosophers semantic facts supervene on one such wider set of facts, but object that it is odd to classify such wider factors as ‘use facts’. But for the purpose of this paper, one can simply take the term ‘use facts’ to be a place holder for whatever facts semantic facts are generally taken to supervene on. (Cf. our discussion about how to interpret ‘use facts’ below, and Williamson (1994), §7.5 who employs the term ‘use facts’ in a somewhat similar way).

aboutness surely won't; intentionality simply doesn't go that deep...If the semantic and the intentional are real properties of things, it must be in virtue of their identity with (or maybe of their supervenience on?) properties that are themselves *neither* intentional *nor* semantic. If aboutness is real, it must be really something else."²

In this paper we argue that (even given a wide interpretation of 'use facts') Semantic Supervenience should be rejected in favour of Semantic Sovereignty: semantic properties really do go that deep.

A few preliminary clarifications are in order. First, *what do we mean by 'semantic facts'?* By semantic facts we mean facts like these³: 'Water' means water; 'water' in English means the same as 'eau' in French; 'Jill' refers to Jill; 'Vulcan' does not refer; 'The queen of England' as uttered in February 2009 refers to Elizabeth II; Jill's current thought is about coffee;

The latter example calls for further comment. One might worry that it is a case of an *intentional* fact, rather than a *semantic* one. Semantic facts seem intricately related to intentional ones. It is not uncommon to assume that words have the meaning or reference they have at least partially in virtue of facts concerning speakers' intentions when using these words. Conversely, it is not uncommon to assume that speakers can be in certain

² Fodor (1987), p. 97.

³ If one is worried that any of these are not facts (e.g. because one thinks that 'Vulcan' does refer, following a certain theory of fictional names), then one can substitute any favoured similar-looking semantic claims instead. The paper does assume, though, that global scepticism about semantics is false.

intentional mental states such as belief in virtue of accepting certain sentences (either in their spoken language or in the language of thought). In such cases, their intentional properties seem to very much depend on the semantic properties of the sentences in question.⁴ The relation between the semantic and the intentional is thus highly complex and clarifying the precise nature of this relation is far beyond the scope of this paper. For the purposes of the paper we will thus count intentional facts as a kind of semantic facts.⁵

Next, *what do we mean by ‘use facts’?* We start by giving the term a very wide interpretation, assuming that the use facts are the totality of physical facts (though we will discuss an even wider interpretation in §3). This seems like a good starting point. It is at least clear that whatever we mean by ‘use facts’, the term cannot be interpreted so widely as to include semantic facts such as the fact that speakers use ‘Jill’ to refer to Jill, or else it would be trivial that semantic facts supervene on use facts. Moreover, it seems that the commonly accepted picture is precisely that semantic facts supervene on physical facts: speakers utter certain sounds, while being in certain brain states, and so on, and it these facts that are sufficient to determine the semantic facts.

The traditional version of this view holds an even stronger claim: roughly that semantic facts pertaining to a certain agent supervene on the physical facts intrinsic to that agent.

So, for example, if two agents are physical duplicates then the word ‘water’ in the mouth

⁴ This latter direction of the relation between semantic and intentional facts is especially important to the issue in question. If an account of the intentional facts must already implicitly incorporate the full semantic theory, then adding the intentional facts to the supervenience base would render SSUP entirely trivial: the supervenience base will already include all the subvening facts.

⁵ If one is unhappy to count intentional facts as semantic ones, one can rephrase the central question of this paper to be whether semantic and intentional facts supervene on use facts, and take our paper to be arguing that they do not. Note that either way, the thesis we are defending goes against the view represented by Fodor’s quote above.

of one of them must have the same semantic properties as the word ‘water’ in the mouth of the other. Famously, semantic externalists reject this claim. According to semantic externalists, it is possible that two agents are physical duplicates, yet the word ‘water’ in the mouth of one, can have the semantic property of referring to H₂O, while in the mouth of the other, have the semantic property of referring to XYZ.⁶ Still, even semantic externalists standardly accept that semantic facts globally supervene on the totality of physical facts (including, for example, the molecular structure of the lakes each agent encounters).⁷

Some may nevertheless object, and insist that physical facts are too narrow to form the relevant supervenience base. But as we shall show in §3, many of the considerations that tell against the version of SSUP that uses physical facts as its supervenience base, generalise to the case where the supervenience base is widened as much as possible. If nothing else, taking the supervenience base to consist of physical facts forms a good entry to the discussion, helping us focus on the problems that SSUP faces more generally.

A third point that needs clarification is *what do we mean by ‘supervenience’?* The literature on the topic provides a large number of non-equivalent notions of

⁶ For expositions of and arguments for semantic externalism, see Putnam 1975, Burge 1979, and Davidson 1987.

⁷ According to some philosophers, one factor that goes into determining semantic facts is the question of which properties are more natural or eligible. (See Lewis 1983). We will assume throughout the paper that naturalness properties are either themselves physical properties at least supervene on physical properties, so that the thesis that semantic facts supervene on physical facts can accommodate the role of naturalness in determining semantic facts. Note that if one holds, as Lewis does, that naturalness facts are necessary, then this assumption is satisfied trivially. Also note that nothing of substance rides on this assumption: if one does not accept this assumption then the relevant arguments we present in §2 can instead be understood as attacking the more general thesis that semantic facts supervene on non-semantic facts (see §3).

supervenience.⁸ But at least an initially plausible candidate for what supporters of SSUP mean by ‘supervenience’ is this:

Strong-Global Supervenience: A-properties strongly-globally supervene on B-properties if and only if, for any two possible worlds w_1 and w_2 , every isomorphism between the objects of those worlds that preserves B-properties, is also an isomorphism that preserves A-properties.

Weak global supervenience seems too weak to capture the spirit of SSUP.⁹ Weak-global supervenience allows, for example, for a world w_1 , where half the speakers in w_1 use ‘water’ to refer to H_2O and the others use ‘water’ to refer to XYZ and where there is a physical duplicate world w_2 of w_1 , where all the speakers uttering ‘water’ in w_1 to refer to H_2O , refer to XYZ in w_2 and vice versa. Strong-individual supervenience, on the other hand, might be too strong to capture the spirit of SSUP.¹⁰ Suppose for example that we do not want count as legitimate physical properties such gerrymandered relational properties as ‘lives near lakes that contain H_2O ’. Semantic externalists will then want to reject the claim that semantic facts strongly-individually supervene on physical

⁸ See, for example, Shagrir 2002, Paull & Sider 1992, Bennett & McLaughlin 2008, and Leuenberger 2008a.

⁹ A-properties weakly-globally supervene on B-properties if and only if, for any two possible worlds w_1 and w_2 , if there is an isomorphism between the objects of those worlds that preserves B-properties, then there is also an isomorphism that preserves A-properties.

¹⁰ A-properties strongly-individually supervene on B-properties, if and only if, for any for any two possible worlds w_1 and w_2 , and any individual x in w_1 and y in w_2 , if x in w_1 is B-indiscernible from y in w_2 then they are also A-indiscernible.

properties, because two physically-indiscernible individuals can differ in the semantic fact of what the word ‘water’ in their mouth refers to. Yet semantic externalists can still accept that semantic facts strongly-globally supervene on physical facts.¹¹ Strong-global supervenience thus seems like a good middle ground, and we shall use it as a starting point for what supporters of SSUP mean by supervenience (though we will consider some alternatives in §4).

The structure of the paper is as follows. In §2 we present our main arguments in favour of SSOV. In §3, we will discuss whether SSUP can be rescued by taking ‘use facts’ to mean ‘non-semantic’ facts rather than ‘physical facts’, and argue that it cannot. In §4, we will discuss whether SSUP can be rescued by interpreting ‘supervenience’ in a more restricted sense than strong-global supervenience, and we will argue that it cannot. In §5 we conclude that that SSUP should be rejected, and highlight the importance of this conclusion.

§2 The case in favour of Semantic Sovereignty

For the purposes of this section, we will take SSUP to be the thesis that semantic facts strongly-globally supervene on physical facts, and SSOV to be the denial of this thesis.

We will present a series of arguments in favour of SSOV.

¹¹ Because global supervenience is formulated only for monadic properties, it is likely that an externalist would need to appeal to certain relational properties of a speaker in order to capture the supervenience of the semantic properties on the physical. If one does not wish to count relational properties as part of the supervenience base, one may prefer a formulation of global supervenience that allows relations and not just properties to be part of the subvening (and indeed the supervening) level. See Leuenberger 2008a for a discussion of the idea of extending formulations of global supervenience to cover relations.

The arguments in §2.1 mimic, to a certain extent, anti-supervenience arguments that occur in other domains (most notably the supervenience of the mental on the physical), but we show here that these arguments have many interesting variants that apply to case of semantic supervenience. This point is important, because while the supervenience of the mental on the physical received much attention in the literature and is considered a highly controversial, semantic supervenience seems to be widely taken for granted, with very little discussion or justification. It is thus well worth seeing that many of the issues that make trouble for supervenience of the mental on the physical, also make trouble for semantic supervenience.

In the remaining two subsections we focus arguments that seem unique to the case of semantic supervenience: in §2.2 we present an argument pertaining to haecceitistic semantic properties, and in §2.3 arguments pertaining to semantic scepticism.

§2.1 Possessors, blockers, and shifters

In this section, we present a range of problems associated with SSUP that have to do with the metaphysical possibility of non-physical entities, for example non-physical ghosts.¹² Such entities can cause trouble for SSUP in three ways. They can add *possessors* of semantic properties, i.e. new entities that possess semantic properties in a physically duplicate world; They can act as *blockers* of semantic properties, i.e. they can prevent entities that possess certain semantic properties from possessing any relevantly similar

¹² We do not discuss here the question of what is to be a physical entity, but we take it as fairly intuitive that there are some possible entities which are ghost-like, and do not count as physical.

property in a physically duplicate world; And finally, they can act as *shifters*, i.e. they can ensure that entities that possess a certain semantic property, possess an alternative similar semantic property in another physically duplicate world.¹³ To be more concrete, we present a series of examples where ghosts act as possessors, blockers, and shifters.

Possessors:

1. Ghost Languages: Consider a possible world w that is a physical duplicate of the actual world, but which also contains a finite number of ghosts. The ghosts in w communicate with each other using a language called ‘Ghostalish’. In Ghostalish, ‘bla’ means peanut. Now consider an isomorphism between w and the actual world that maps each physical object to its duplicate, and maps each of the ghosts to some mundane abstract object in the actual world.¹⁴ This isomorphism preserves physical properties, but not semantic properties. For example, one of the ghosts has the semantic property that she uses ‘bla’ to mean peanut, and that property is not preserved under the isomorphism, so SSUP fails.

¹³ The precise distinction between the three categories is not sharp, but the distinction is also not terribly important for our purposes. It just helps to keep track of what features the different types of examples play on.

¹⁴ We assume that there are infinitely many necessarily existing abstract objects (for example infinitely many numbers) and thus that we can map the ghosts to those abstract objects, while still maintaining a one-to-one correspondence between the domains of the two worlds. If the reader does not accept this assumption, rephrase the example by comparing w to another world w^* , also a physical duplicate of the actual world but one that contains some mundane ghosts who do not speak a language (the same number of ghosts as there are in w).

Shifters:

2. New Referents: Consider a possible world w that is a physical duplicate of the actual world, but which also contains ghosts. Suppose that in the actual world a certain human speaker, John, utters the phrase 'The object nearest to me'. In the actual world, the object nearest to John is a pencil, whereas in w , there is a ghost standing between John and the nearby pencil. So in the actual world and in w , John's utterance has different semantic properties: in the actual world it has the property of referring to the pencil, while in w the property of referring to the ghost, thus violating SSUP.

3. Disjunctively introduced words: Let us define the word 'Parbon' to mean Paris if there are no ghosts, and Lisbon if there are ghosts. In the actual world (we assume), there are no ghosts, so 'Parbon' means Paris. Yet consider a physically duplicate world w , in which there are ghosts. In w , we (or our physical duplicates) have just made the same disjunctive stipulation as above, but following their stipulation, 'Parbon' in w means Lisbon. Thus our utterances using 'Parbon' in the actual world have different semantic properties from those of the duplicate utterances in w , in violation of SSUP.

4. Drowning use-facts: Let us open with a thought experiment that does not involve ghosts. Suppose that the community of English speakers consists of one million speakers, and all but one of them use the word 'flower' in a way that, other things being equal, refers to flowers. Yet one deviant speaker uses the word 'flower' in a way that, other things being equal, refers to weeds. Semantic externalist would naturally allow that when the deviant speaker utters the sentence 'Flowers are nice', other things are *not equal*: the

utterance means that flowers are nice, not that weeds are nice, despite the deviant speaker's idiosyncratic use-facts. (Of course, the mere difference in number might not quite be sufficient here. We may, for example need to assume that the deviant speaker isn't the world's greatest botanic expert. The point is that the relevant details can be filled so as to make the deviant speaker's use patterns be drowned by those of his community).

With this thought in mind, consider a physical duplicate world of the actual world w , which contains zillions of ghosts. Now suppose that in w all the ghosts speak a language that is very much like English, and consider themselves as one with the community of the duplicate English-speakers.¹⁵ With enough of the details filled in as needed, we can assume that ghosts speak the same language as that of the duplicate English-speakers: call that language 'Ghostalish'.¹⁶

Suppose however, that when the ghosts use 'flower' they use it in a way, that other things being equal, refers to weeds. Since the ghosts form the vast majority of the community of

¹⁵ By 'the duplicate English-speakers' we simply the duplicates in w of those people who are English speakers in the actual world. We need not determine whether the language they speak in w is strictly identical to English.

¹⁶ To give a feel of how the details can be filled to satisfaction, we note that we have already assumed that the Ghosts view the human duplicate English-speakers as members of their community. Moreover, it is not stretch to assume that the humans count the ghosts as members of their community: even if the humans don't believe in ghosts, they might well be happy to accept that if there were ghosts in their vicinity that spoke a language very much like English, the ghosts would be part of the same linguistic community. Nor are we prohibited from assuming that the humans casually influence the ghosts, or even that ghosts casually influence the humans (though in a way that is over-determined by physical causation).

Ghostlish speakers (and, if you want, they are also much better botanic experts than all humans, and so forth), the word ‘flower’ in Ghostlish means weed rather than flower. It follows that the duplicate English speakers in w, who – we have assumed - are speakers of Ghostlish, use ‘flower’ to mean weed. Just as the one deviant speakers’ use in our initial thought experiment was drowned by the use of the 999.999,999 other English speakers, the use of the millions of duplicate English speakers in w is drowned by the use of the zillions of ghosts. In the actual world (where there are no ghosts) the word ‘flower’ as uttered by a human English speaker means flower, but as uttered by their duplicate in w it means weed, so SSUP fails.

5. New Reference Magnets: Here’s a rough version of a theory that supporters of SSUP often subscribe to. Consider a predicate like ‘rabbit’. How is the reference to it determined? First, there are the (narrowly construed) use facts of the predicate: that is they include the pattern of all cases where speakers utter the word ‘rabbit’, including speakers brain-states, environment, behaviour, and so forth. These facts, however, may not be sufficient to determine the reference of ‘rabbit’. For example, all the use facts might be (in some sense) equally compatible with the hypothesis that speakers use ‘x is a rabbit’ to mean ‘x is a rabbit and self-identical’. The thought, though, is that there is a further relevant fact to consider: the property of being a rabbit is a more natural property (whatever that means) than the property of being a rabbit and self-identical. Moreover, the property of being a rabbit is the most natural property that is also a good fit for the overall pattern of use for ‘rabbit’, and thus serves as a *reference magnet*.¹⁷

¹⁷ For more on this picture of how words come to refer, see Lewis 1983, Hawthorne 2006 and Sider 2009.

Now with this picture in mind, consider a physically duplicate world w , which contains a ghostly substance, call it 'ghost-water'. Suppose that ghost-water is always located wherever water is located. Suppose further that for some reason, the property of being ghost-water is more natural than the property of being water.¹⁸ In the actual world, where there is no ghost-water, the most natural property that is a good fit for English speakers' use of 'is water' is the property of being water. However, it may be that in w the most natural property that is a good fit for the use facts is the property of being ghost-water: ghost-water provides a more appealing reference magnet than water, thus overriding the latter's magnetic appeal.

Some might object by arguing that a non-physical substance such as ghost-water cannot count as more natural than water; others might worry that ghost-water cannot provide a good enough fit for the use facts for 'water' in w . In response, let us present another variant of the example that is less likely to raise such worries. Consider the meaning of the word 'ghost' as used by the English speakers in the actual world, and by their duplicates in w . In the actual world, 'ghost' describes a pretty general type of non-physical entity. In w , however, 'ghost' might mean something slightly more specific. Suppose, for example, that all the ghosts in w belong to some natural subspecies of ghosts: ghost-sub-1. Since w has some ghosts, and all of them are ghost-sub-1, being a ghost-sub-1 might serve as a much better reference magnet for the word 'ghost' than the

¹⁸ Why this might be the case depends on how one construes the notion of naturalness. On at least some formulations of naturalness, though, this should be possible: for example ghost-water might be simpler, more homogenous in its nature than water, and so forth.

more general property of being a ghost. (To see this, compare for example the actual meaning of ‘unicorn’, to the meaning of the word in a world very much like ours, but where unbeknownst to us there is exactly one instantiated natural kind which are horse-like and have single horn. It’s plausible that in that world, that specific natural kind would serve as the reference magnet for ‘unicorn’, while in our world ‘unicorn’ means something very general – something like ‘horse-like creature with one horn’). In each of our scenarios, the semantic properties concerning the meaning of ‘water’ and ‘ghost’ respectively differ between physically duplicate worlds, so SSUP fails.

Blockers:

6. Too many reference magnets: In the examples we presented in 5 above, new reference magnets served as shifters of semantic properties. But they can also serve as blockers. Suppose, for example that the property of being ghost-water described above is *just as natural* and just as good a fit for our use of ‘water’, as the property of being water, and that there is no other property that would serve a better reference magnet than ‘water’. (For example, we may assume that the disjunctive property of being either ghost-water or water is highly unnatural). We take it that on the standard reference-magnet picture, the conclusion to draw in this case is that in the ghost-water world described, ‘water’ would simply fail to refer.¹⁹ If this is so, then in such a case, ghost-water would serve as a blocker and again SSUP would fail.

¹⁹ Though see §2.3 for further discussion of this inference.

The idea of using possessors, blockers, and shifters to show failures of supervenience on the physical is familiar from the literature on the supervenience of the mental on the physical.²⁰ Indeed, one can extend some of the original examples in that debate to our current purposes. Consider for example a possible world that is a physical duplicate of the actual world, but where Jill's duplicate is a zombie – a creature with no qualia. When Jill utters in the actual world the phrase 'what it's like for me to see red things', she refers to a certain phenomenal property. However, when her zombie-duplicate utters the same phrase, she fails to refer. Similarly, consider a world where Jill's duplicate is an invert: inverts have qualia, but what it's like for Jill's invert to see red is what it's like for Jill to see green and vice versa. So when Jill and her invert duplicate utter 'what it's like for me to see red things' they refer to different phenomenal properties. Zombies can act as blockers for certain semantic properties, inverts as shifters. Either way, at least on our current construal of SSUP, the thesis fails.

§2.2 Haecceitistic semantic properties

In this section we present a different reason to doubt SSUP. Consider a physical duplicate w of the actual world (one with no ghosts or other supernatural entities), in which Jack's duplicate is not Jack himself, but some other physically indiscernible person – call him 'Jack*'. Now in the actual world, Jill has the semantic property of referring to Jack by 'Jack'. Yet in w , Jill's duplicate has the semantic property of referring to Jack* by 'Jack'. So the two worlds differ in their semantic properties, without differing in their physical properties. So SSUP fails.

²⁰ See Hawthorne 2002, Leuenberger 2008b.

What lies in the heart of this example, is that (some) semantic properties are *haecceitistic properties*, while physical properties are not haecceitistic: the actual world and the world *w* described above can count as physically indiscernible despite having the haecceitistically distinct physical objects Jack and Jack*, while they do not count as semantic duplicates.

Further argument is in order. First, why do we say that physical properties are not haecceitistic? To be clear, it is not because we think that at a fundamental level there are no haecceitistic facts about physical objects. We are willing to concede that there is a clear difference between a possible world which contains a certain particular electron *e* and one that contains a duplicate but distinct electron, *e**, and we are willing to concede that the electron has the haecceitistic property of being identical to *e*.²¹ The point we do want to make, however, is that such haecceitistic facts and properties are not the kinds of facts that people standardly refer to when they claim that a certain realm of facts or properties supervene on the physical ones. That is to say, in the sense of ‘physical properties’ that is used for the purposes of SSUP and similar supervenience theses, haecceitistic properties of physical objects do not count as physical properties.²² Consider for example supporters of the claim that phenomenal properties supervene on physical properties. Suppose it turns out that there is possible world *w*, which is exactly like the actual world, except for one electron *e* in Jack’s actual brain, which is replaced by

²¹ See Chisholm 1986 for arguments against this position.

²² If the reader objects to disqualifying such haecceitistic properties from counting as ‘physical’, they can rephrase SSUP to be the thesis that semantic properties supervene on physical* properties, where physical* properties are those physical properties that are not haecceitistic. We argue below that the new formulation captures what supporters of SSUP have intended all along.

duplicate electron e^* in w . Suppose that Jack and his duplicate have radically different phenomenal properties. Surely, if this possibility is genuine it constitutes a counterexample to the supervenience of the phenomenal on the physical. It would seem extremely odd for a defender of such a supervenience claim to argue that the two worlds are not really physical duplicates because of the haecceitistic difference between e and e^* . Or consider the supervenience of the moral on the physical. We can imagine a odd position that allowed for merely haecceitistic differences to make a moral difference: perhaps God has chosen to bestow Jill with the moral permission to kill whoever she wants, but has chosen not to bestow this permission on her duplicate Jill*. Thus Jill and Jill* can differ in the moral property of whether they ought not to kill Jack. Yet we take that if such odd position were true, it would amount to denying the supervenience of the moral on the physical. We thus conclude that the physical properties that are relevant for purposes of supervenience theses such as SSUP do not include haecceitistic properties.

Next, why do we insist that (some) semantic properties *are* haecceitistic? This question is more complex than it might initially appear. The mere fact that we can point out to the semantic fact that 'Jack' refers to Jack (rather than to Jack*) is not sufficient. After all, we conceded that there is a non-semantic fact that (the physical object) Jack is identical to the (the physical object) Jack, rather than to Jack*, yet we do not want to count that fact as part of our physical supervenience base. We also concede that there is a moral fact that Jill ought not to kill Jack, rather than Jack* (perhaps there is no moral prohibition on killing merely possible objects). Yet we take it that the existence of such facts is compatible with the supervenience of the moral on the physical. This implies that this

kind of haecceitistic moral-related fact does not belong to the set of supervening moral facts we want to account for.

A plausible story for what goes on in the moral case is this. There are some more general moral properties, such as the property of being such that one ought not to shoot other innocent persons. Now suppose such property holds of Jill in the actual world. It's certainly plausible that this property does not hold of Jill necessarily: after all, in possible worlds in which shot-wounds are highly pleasurable and heal themselves immediately, Jill lacks this moral property. However, if the supervenience of the moral on the physical is correct, then such property does hold of Jill (and her duplicates) in every physically duplicate world. (This is entirely plausible, because in all such worlds, shot wounds are just as painful and dangerous as they are in the actual world). The final piece in the line of reasoning the supporter of supervenience of the moral would take here is this: the moral fact that Jill ought not to shoot Jack isn't really a fundamental moral fact. It is derivative on the fundamental fact that Jill ought not to shoot innocent people, and the fact that Jack is an innocent person.

It is, however, far from clear that an analogous story would work for the case of the semantic. Consider the semantic property of Jill that she uses 'Jack' to refer to Jack. We can perhaps derive that property from a more general complex property such as Jill using 'Jack' to refer to whatever object satisfies the description of being a male person, who Jill often utters the sounds 'Jack!' in their vicinity, and so forth (call this complex general property 'P'). As opposed to the moral case, P does not seem to be the kind of property

that semantics is fundamentally concerned with. Semanticists have theories about the meaning and reference of words such as 'Jack', not about how those referents are determined relative to the complex behaviours of speakers. (The latter seems, if anything, the job of *meta-semantics*). Nor is this merely a matter of how we label different fields of linguistic theorising. Semantic properties enter into compositional semantic analyses. When we consider uses of 'Jack' in sentences such as 'Jill said that Jack is coming' or 'It is necessary that Jack is coming', it is clear that the meta-semantic property P will have little role to play.

Here are two further ways to see that the semantic properties that the defenders of SSUP allude to are *not* general meta-semantic properties such as P above. First, if properties like P were the semantic properties in question, it is hard to see why semantic externalism is widely taken to have taught us that semantic properties don't (strongly-individually) supervene on the intrinsic physical properties: after all, if the relevant "semantic" property concerning 'water' in Jill's mouth is that it refers to whatever clear liquid that is in her vicinity, one that she likes to utter the sounds 'water!' next to, and so forth, then it is very hard to see why an intrinsic physical duplicate of Jill will not have exactly the same property. It is precisely because externalists focus on specific properties of 'water' such as referring to H₂O, that they maintain that semantic properties do not (individually) supervene on intrinsic physical properties.

Second, consider the thesis of *arbitrary reference*.²³ Suppose we are presented with two qualitatively identical, spatially coincident spheres. According to the thesis of arbitrary reference, we can use stipulations such as ‘Let ‘A’ arbitrarily refer to one of the two spheres’, with the result that ‘A’ determinately refers to a particular one of the two spheres. Moreover, it is a random or arbitrary matter which sphere one ends up referring to following the stipulation. Consequently, there can be two physically duplicate worlds, w_1 and w_2 containing exactly the same pair of spheres, where the above stipulation results in ‘A’ referring to one sphere in w_1 , but to the other sphere in w_2 .

Now the thesis of arbitrary reference is certainly a radical one, and we do not wish to defend it here. Rather, we want to point out that the thesis of arbitrary reference is taken to be radical, precisely because it is assumed to violate the prevalent thesis of semantic supervenience. But if SSUP were interpreted so that semantic facts are general rather than haecceitistic, then arbitrary reference would be perfectly consistent with SSUP! After all, any purely qualitative claim one can make about the reference of ‘A’ following the stipulation (e.g. that it refers to sphere of this and that size, in such and such location) would not distinguish between w_1 and w_2 . Since we take that arbitrary reference would constitute an obvious violation semantic supervenience, we conclude that the latter is understood so as to include haecceitistic semantic facts.²⁴

²³ The thesis is defended and formulated in Breckenridge & Magidor (*forthcoming*)

²⁴ Robbie Williams suggested to us a different interpretation of supervenience, one according to which A-properties supervene on B-properties just in case both (i) non-haecceitistic A-properties supervene on non-haecceitistic B-properties; and (ii) haecceitistic A-properties supervene on haecceitistic B-properties. The idea is that this notion of supervenience both captures the spirit of the relevant theses (for example, the

The claims we have made in this section are hard to defend, because they are partially a question of terminology. We have tried to argue, however, that SSUP is best interpreted as the thesis that semantic properties (*including haecceitistic ones*) supervene on physical properties (*not including haecceitistic ones*). But on this interpretation, we have shown, SSUP is simply false.²⁵

§2.3 Underdetermination and semantic scepticism

There are numerous arguments in the philosophical literature that purport to show that use facts are insufficient to determine semantic facts – either in general, or in some specific domain of language.

electron swap example suggested above will be inconsistent with supervenience in this sense, because it would violate clause (i), and arbitrary reference will be inconsistent with supervenience because of clause (ii). Yet this notion of supervenience is consistent with the existence of haecceitistic semantic facts (such facts will supervene on the haecceitistic physical facts).

Our objection to this proposal is that we think it still fails to capture the spirit of SSUP. Consider a variant of the arbitrary reference scenario above: suppose w_1 contains the duplicate spheres X and Y, while w_2 contains the spheres X and Y*, where Y* is qualitatively identical but numerically distinct from Y. Now suppose following a stipulation such as above ‘A’ ends up referring to X in w_1 , and to Y* in w_2 . This scenario would not violate supervenience in the sense suggested by X (clause (i) would not be violated because the qualitative semantic properties are preserved, and clause (ii) would not be violated because the physical haecceitistic properties are *not* preserved). But this scenario still seems to us to intuitively violate SSUP.

²⁵ For a more detailed argument for the claim that if physical facts are non-haecceitistic and semantic facts are not, then supervenience fails, see Hofweber 2005. It should be noted, however, that Hofweber largely assumes, and thus does not argue in detail for, the antecedent of the conditional.

Consider Quine's arguments for the inscrutability of reference: nothing in speakers' use-facts is sufficient to determine whether by 'rabbit' they mean rabbit, or undetached rabbit parts, or rabbit time-slices.²⁶ Consider Kripkenstein's rule-following arguments: nothing in speakers' use facts is sufficient to determine whether by 'plus' they mean addition or quaddition.²⁷ Consider Putnam's model theoretic arguments: whatever (first-order) constraints the use-facts place on the meanings of our words, there will be different conflicting models satisfying these constraints, so the use facts are insufficient to determine which of these models is the intended one.²⁸ Consider Robbie Williams's new model theoretic argument: Williams argues that the standard attempts to address the model theoretic arguments by preferring more natural interpretations over less natural ones fail, by presenting a new version of the argument in which, he claims, the distinct models will be equally natural and thus we will not be able to use naturalness as a criterion to choose between them.²⁹ All of the above are fairly general arguments for the claim that use facts do not determine semantic facts.

There are also arguments that apply in somewhat more specific domains. Consider the problem of vagueness. Even if we are willing to accept that our use facts somehow determine that by 'plus' we mean addition rather than quaddition, many find it much harder to accept that our use-facts are sufficient to determine whether by 'tall' we mean taller than 1.8m or 1.80001m. Consider the problem of reference to indiscernibles:

²⁶ See Quine 1960

²⁷ See Kripke 1982, Wittgenstein 1953

²⁸ See Putnam 1977

²⁹ See Williams 2007

suppose there are two indiscernible co-located bosons. Can we successfully name one of them 'A'? If so, how can our use succeed in picking out one rather than other? Similarly, given the automorphism of the complex numbers onto themselves, how can we succeed in referring to one square root of -1 by 'i' and the other by '-i'?

The first thing to note about these arguments is that they all provide prima facie reasons to doubt SSUP. If use facts are insufficient to determine semantic facts (whether globally or in a specific domain) then SSUP fails. Interestingly, though, in the vast literature on these problems, few seem to take this conclusion at face value.³⁰ Instead, there are broadly two kinds of responses that are given to the sceptical arguments.³¹ The first is to concede the conclusion of the argument, and thus to accept some form of *semantic scepticism*: if use facts are insufficient to determine semantic facts in general, then perhaps there are no semantic facts of the matter. Or if use facts are insufficient to determine semantic facts in some specific domain (for example, in the domain of borderline cases of vague expressions), then perhaps there are simply no semantics facts of the matter in that domain (e.g. there is no fact of the matter concerning the precise boundary of 'tall').

³⁰ One rare exception is Boghossian 1989, which concludes that the so-called 'Kripkenstein puzzle' ought to be solved by accepting a non-reductive view of semantic facts. At least on one natural interpretation, Boghossian means to count supervenience accounts as reductive in this context.

³¹ For some (but not all) of the arguments there are also responses that aim to show that the argument is fundamentally flawed or doesn't even get off the ground. We do not consider such responses here.

The second type of response tries to avoid the arguments' conclusions by accepting certain forms of *brute supervenience*: perhaps use facts merely seem to be insufficient to determine the relevant semantic facts because we have implicitly placed too many constraints on the way in which semantic facts are determined by use facts. For example, we may have assumed that the function that takes us from use facts to meaning facts is in some way fairly simple, or systematic or (at least roughly) epistemically accessible to us. But supervenience does not have to take this form: semantic facts might supervene on use facts in a brute fashion that is not very systematic, informative, or explanatory. Responses to the arguments that appeal to naturalness or eligibility are one instance of this strategy. (If we had a good grasp of what 'naturalness' amounted to perhaps they would not count as postulating 'brute supervenience'. But it is notoriously difficult to come up with a notion of naturalness that would not be completely brute and yet will do the job of telling apart different semantic hypotheses).³² Another instance of this strategy is Williamson's view of vagueness.³³ According to Williamson use facts do determine semantic facts, but we have little to say about how precisely the supervenience function works, and have no hope of ever discovering it.

Given the prominence of the problems of semantic underdetermination, it is surprising that a third line of response to these problems has been almost completely neglected.³⁴

³² See Cameron (2010) for a view that explicitly endorses the bruteness of naturalness properties.

³³ See Williamson 1994.

³⁴ Again, Boghossian 1989 is one rare exception. It should be mentioned, however, that Boghossian only briefly defends a non-reductive account of meaning in the light of Kripkenstein's worries about rule following. It is unclear both whether Boghossian takes a non-reductive account of meaning to entail the

This is the solution that takes the phenomenon of underdetermination at face value, and simply rejects SSUP. Once SSUP is rejected, we can concede that the underdetermination arguments indeed succeed in showing that use facts are insufficient to determine semantic facts. However, we can avoid the sceptical conclusion that these arguments are standardly taken to suggest: the underdetermination of semantic facts by use facts *does not* entail that there are no semantic facts, because if semantic facts do not supervene on use facts, there is no puzzle about why they fail to be determined by use facts. This in turn paves the way to new and powerful solutions both to the general sceptical worries and to the specific ones such as vagueness and reference to indiscernibles. To be clear: we are not arguing that a mere rejection of SSUP solves the relevant problems once and for all. One would need to work out specific theories that address each of the problems in turn. Our argument is simply that rejecting SSUP opens the road to a wide range of new possible solutions, one that have been largely neglected.³⁵

One might worry that solutions to underdetermination problems that appeal to SSOV are no improvement on those that appeal to brute supervenience: what is the advantage of postulating brute semantic facts, rather than brute supervenience facts? One reason to think solutions involving SSOV fair better is simply that it is not clear that the solutions involving brute supervenience can be made to work. We mentioned briefly Williams's

non-supervenience of the semantic on the use facts and whether he would be happy to extend his idea to account for the other forms of semantic scepticism we mention.

³⁵ One way to develop these solutions is to use SSOV to defend the thesis of Arbitrary Reference alluded to above, and to then apply arbitrary reference to a range of these problems. This line of thought is suggested in Breckenrdige & Magidor (*forthcoming*)

argument that purports to show that eligibility is insufficient to solve the Putnam model theoretic argument, and in any case, it is very unlikely that eligibility response would suffice to address the problem of vagueness (the problem seems especially difficult precisely because none of the potential precisifications for ‘tall’ seems more natural than the others). And we have argued elsewhere that Williamson’s theory of vagueness has a hard time accounting for why we cannot come to know the sharp boundaries of vague expressions.³⁶

But more importantly solutions that appeal to SSOV have the following advantage. One way to understand the picture underlying SSOV is that (certain) semantic facts are (in some sense) fundamental. The thought is roughly this: if the semantic fails to supervene on the physical, then it does not *derive* (or derives only partially) from the physical (or non-semantic). And if the semantic is not derivative, it is fundamental. Thus failure of supervenience at least goes naturally with a picture of the semantic as a fundamental aspect of the world, and we should not find bruteness at the fundamental level particularly offensive (after all, some facts must be brute if we want the domain of facts to be well-founded). Brute supervenience, on the other hand, is dialectically puzzling: on the one hand, it insists that semantic facts are *not* fundamental and that it is important that we reduce them to more fundamental use-facts. On the other hand, they are in a certain sense nevertheless brute (because of the bruteness of supervenience). While there is no logical contradiction in this combination, it leaves us in a dialectically unhappy position, one that SSOV avoids.

³⁶ Kearns & Magidor (2008)

In this sub-section we have argued that the range of arguments for underdetermination of the semantic, provide us with at least *prima facie* reasons to think that SSUP fails. Moreover, if instead of taking appearances as misleading we simply endorse these cases as genuine failures of SSUP, the road is paved for new and powerful solutions to the puzzles, ones that have been by and large ignored up to now. If there are indeed successful solutions to the puzzles which rely on SSOV, then inference to the best explanation would provide us with a strong reason to accept the thesis.

§3 A different interpretation of use-facts

In §2, we argued that, assuming use facts are to be understood as physical facts, SSUP should be rejected in favour of SSOV. But perhaps this merely shows that our initial interpretation of ‘use facts’ was too narrow. An obvious alternative is to take use facts to be *non-semantic* facts. Understood in this way, SSUP claims that the semantic facts supervene on the non-semantic facts. (It is important to remember that the semantic facts include the intentional facts, so the non-semantic facts do not include any intentional facts.).³⁷ In this section, we will examine whether understanding SSUP in this manner can salvage the thesis, and we will argue that it cannot. In §3.1-33, we will show that although the new formulation affects some of our original arguments in §2 as we presented them, the essence of our arguments still stands, and consequently SSUP ought to be rejected.

³⁷ It’s also worth keeping in mind that the defender of SSUP is likely to accept that one does not need the full set of non-semantic facts or properties in one’s supervenience base, and would accept that semantic facts supervenes on some subset of relevant non-semantic properties.

§3.1 Ghosts and the non-semantic

Our first series of objections to SSUP (§2.1) were based on the possibility of non-physical objects such as ghosts (or non-physical properties, such as qualia). For example, we noted that there are worlds that are physical duplicates of ours but contain non-physical ghosts. The existence of these ghosts makes a difference to what semantic properties and relations are instantiated. Thus SSUP is false, given that we understand the use facts to be the physical facts. Does understanding the use facts to be the *non-semantic* facts help with this problem?

In certain ways it does. Consider for instance, the idea that these non-physical ghosts can act as new referents for terms such as ‘the object nearest to me’. When spoken by Jack in the actual world such a term refers to a pencil, but in a physically identical but haunted world, the term refers to a non-physical ghost. This shows that physical indiscernability does not suffice for semantic indiscernability. However, there is a clear *non-semantic* difference between the two worlds. The fact that, in one world, Jack is nearest a pencil and in the other, John is nearest a ghost is a non-semantic difference. A defender of SSUP may thus argue that the semantic difference rests exactly on this non-semantic one.

However, there are two kinds of variants of the cases we presented in §2.1 that make trouble even for supervenience on the non-semantic. The first involves purely semantic (and in particular, purely intentional) entities, and the second, purely semantic (or purely intentional) properties.

First, consider the case of purely intentional ghosts. Purely intentional ghosts have various mental states (beliefs, desires, etc.) but lack any (interesting) non-semantic properties. Such ghosts seem readily conceivable. Indeed, they seem to be what are normally called Cartesian minds. The essence of a Cartesian mind is thought. Insofar as thought is intentional, Cartesian minds are purely intentional, and thus (as we defined the term), purely semantic.

However, if purely semantic ghosts are possible, then their possible existence suffices to prove SSUP false.³⁸ To see how this causes a problem for SSUP, consider a possible world *w* which just like the actual world, except that it also contains some purely intentional ghosts. This world is a non-semantic duplicate of the actual world (the purely intentional ghosts do not make a difference to the non-semantic properties), but need not be a semantic duplicate: the purely intentional ghosts can act as possessors (a purely intentional ghost can have the property of thinking about philosophy), as shifters (we can still introduce the word ‘Parbon’ disjunctively as suggested in example 4 above, or let ghosts act as reference magnets for ‘ghost’, as suggested in example 5 above), and as blockers (purely intentional ghosts can still make for too many reference magnets as discussed in example 6 above).

³⁸ Defenders of SSUP might wish to argue that purely intentional ghosts are not possible. But this is a controversial metaphysical claim, that will require further defense (we remind the reader that SSUP is taken to be, by and large, entirely uncontroversial). Moreover, below we present other arguments, which do not rely on the possibility of purely intentional ghosts.

Similar problems arise when one considers purely intentional properties. According to at least one plausible hypothesis, qualia are purely intentional (and thus purely semantic) properties. Now consider a possible world w , which is just like the actual world except that Jill's duplicate in w is a zombie or an invert. If qualia are purely intentional properties, w only differs from the actual world in intentional respects, and thus counts as a non-semantic duplicate of the actual world. But as we have seen, such a world makes trouble for SSUP: Jill's utterance of 'What's it like for me to see red' would differ in its semantic properties between the actual world and its duplicate w , thus violating SSUP.

Counterexamples to SSUP, on its new formulation, need not be restricted to the purely intentional. Consider the case of words – semantic objects, which bear semantic properties. Let w_1 and w_2 be two worlds containing magic, in the style of Harry Potter.³⁹ The worlds differ in only one respect: w_1 contains a magic word - call it 'Abra' - that has the magical power to turn someone into a unicorn when uttered. In w_2 , Abra has no such magical powers. Other than this, w_1 and w_2 do not differ at all. In particular, Abra is never uttered in either world.

We claim that in this scenario w_1 and w_2 are non-semantic duplicates, and that the only differences between them are semantic. Why so? First, it seems that differences between the properties of Abra itself are semantic. One reason for this is that the magical powers of the word in w_1 may plausibly be seen as part of the word's semantics. Think of the analogy with performatives such as 'promise'. Just as one promises through appropriate utterances of sentences containing the word 'promise', and this fact is due to the role the word plays in the utterance, one may (in w_1) turn someone into a unicorn via an utterance

³⁹ Thanks to Vann McGee for suggesting this example.

of the word *Abra*, and this fact is due to the role this very word plays in the utterance.⁴⁰ A second, somewhat more complex consideration, is that there are good reasons to classify the property of ‘magically turning people into unicorns when uttered’ as a semantic property, independently of whether the property is a constitutive part of the semantic value of *Abra*. To see why, note for a start that the property of *being uttered* should plausibly be classified as a semantic property: necessarily, only semantic objects can be uttered, and saying of a non-semantic object that it is uttered results in a kind of category mistake. (Recall that as we are thinking of it, merely producing some relevant sounds is not sufficient for uttering.)⁴¹ Furthermore, note that given the spirit of SSUP, mixed properties (namely ones that have both a semantic and a non-semantic component) should be classified on the side of the subvening class, rather than on that of the supervenience base. It follows then that the mixed property of having the power to turn someone into a unicorn when uttered is a semantic difference.

A similar observation holds with respect to differences between the worlds which concern properties of objects other than *Abra* itself. The worlds differ, for example, in that Harry Potter has in one world but lacks in the other, the property of being such that, if he were to utter *Abra*, he would turn someone into a unicorn. But this again, is a ‘mixed property’, and should thus be classified on the side of the semantic.

⁴⁰ It is important to realize that in order for one to produce the magic in w_1 one has to utter the very word *Abra*, and not, for example, some homonym of it. Thus this is really a fact about the word, and not about, e.g., the sounds by which it is pronounced.

⁴¹ If one is still not convinced that being uttered is a semantic property, one can change the example where the magical laws of the w_1 are such that some magic is preformed whenever someone refers to cats using the word *cat* (and there are otherwise no differences between the worlds). On the amended example the magical powers of the word *cat* may not be part of its semantic value, but the differences between the worlds would still be semantic because of the second consideration discussed above.

These considerations are perhaps not clear cut (they depend on how one understands the highly vague distinction between the semantic and the non-semantic). But the driving intuition is this: if one were asked to give a description of w_1 and w_2 in purely non-semantic terms, this description could not include facts about words, what happens when they are uttered, and so forth, and thus would not be able to capture the difference between the two worlds.

A final objection to the case might be raised: if Abra really does differ in its magical powers between w_1 and w_2 , then even accepting that such differences are semantic, there must also be a non-semantic difference in the non-semantic properties of the non-semantic objects upon which this semantic difference depends. According to this objection, it is simply question-begging to assume that there are worlds which differ purely in the respects we have suggested. In reply, we would like to point out that there is nothing clearly incoherent about the scenario we have presented. It would be highly surprising if such a scenario were indeed impossible. Insofar as our example entails the falsity of SSUP, it is clear that a proponent of SSUP must either give up her position or claim that the example as we have described it is impossible. But this does not make our example question-begging. Our case presents a challenge to the proponent of SSUP precisely because the scenario we describe seems perfectly possible. Of course we have not *proven* that the scenario is possible, but we have done enough to cast serious doubt on SSUP (and certainly on the common assumption that SSUP is obviously true).

We conclude that the reinterpretation of ‘use-facts’ does not threaten the essence of our argument from §2.1. The reinterpretation merely requires us to choose the examples in our arguments a little more carefully.

§ 3.2 Haecceitistic properties

We argued in §2.2 that the semantic properties include certain haecceitistic properties while the physical properties do not contain any such properties. As such, we were able to construct cases in which there is a haecceitistic difference between two worlds without there being a physical difference between them. This haecceitistic difference is then able to ground a semantic difference between the two worlds, once again showing that the semantic does not supervene on the physical.

Does understanding the supervenience base to be non-semantic properties solve the problem? Consider for example the possible world w , which is a physical duplicate of the actual world, but where instead of Jack, w contains a physical duplicate of his - Jack*. As we argued in §2.2, the actual world and w are nevertheless physical duplicates. Isn't the difference between the actual world and w , though, at least a *non-semantic* difference? We have no problem in characterising the difference between the actual world (containing Jack) and w (containing Jack*) as in some sense non-semantic. But the crucial point is that property of being Jack (rather than Jack*) ought not to belong to the set of *relevant non-semantic properties* that form our new supervenience base. The reason for this is precisely the one we gave for not counting such haecceitistic properties as physical in §2.2. Recall that our reasoning there did not rely on some intuitive

understanding of the term ‘physical property’. Rather, the reasoning showed that defenders of supervenience theses such as SSUP should not, on reflection, allow such haecceitistic properties into their supervenience base. For one thing, if such properties are allowed, then purely haecceitistic swaps that make for a large difference in the supervening properties would be consistent with supervenience.⁴² The bottom line is that in so far as our argument for not taking *physical* properties to be haecceitistic was successful, the argument would equally show that the *relevant non-semantic* properties ought not to be haecceitistic either, and thus our original argument stands.

As a second line of thought, we want to point out that even if one were to allow non-semantic haecceitistic properties to form part of the supervenience base, it is not clear that this will suffice to account for all *semantic* haecceitistic properties. That is, we believe it is plausible that there are certain semantic haecceitistic properties that are instantiated in one world while not being instantiated in a non-semantic duplicate (right down to the non-semantic haecceitistic properties) of this world. Consider, for example, the semantic haecceitistic property of being the word ‘bus’. Insofar as it is possible for any two cross-world objects to be qualitatively identical yet distinct objects (which most philosophers allow), there seems no bar on two words being qualitatively identical yet distinct. Thus in the actual world, there is a word that has the semantic haecceitistic property of being the word ‘bus’. There is another world which is a non-semantic duplicate of this world (including the non-semantic haecceitistic properties) that contains

⁴² And even focusing only on qualitatively identical worlds, allowing haecceitistic properties to be part of the supervenience base would allow for scenarios such as those suggested in footnote 24 above.

a qualitatively identical but distinct word. This word thus has a different semantic haecceitistic property. There is, then, a semantic difference between the two worlds without a non-semantic difference. Similarly, consider haecceitistic properties concerning the purely intentional ghosts, discussing in §3.1 above. Suppose some ghost-name, ‘George’, has the property of referring to one such purely intentional ghost George. A non-semantic duplicate of that world might have a duplicate ghost, George* instead of George, and thus differ in this semantic property. That we insist on the exact same duplicate *non-semantic* properties does not settle the issue. We conclude, therefore, that even when phrased in terms of non-semantic properties, SSUP fails.

§ 3.3 Semantic scepticism

In §2.3, we claimed that the many arguments for the idea that the physical facts underdetermine the semantic facts provide good evidence that the semantic facts do not supervene on the physical facts. Thus the phenomenon of vagueness plausibly arises due to the fact that the physical facts do not completely determine what property is picked out by a vague predicate. It is a simple step from this view, plus a general commitment to bivalence, to the idea that the semantic facts ‘go beyond’ the physical facts and thus fail to supervene on them. Similar stories can be told of many other interesting semantic phenomena.

We wish simply to point out that it is obvious that all of these worries apply with equal force (however forceful that is) to the claim that the semantic supervenes on the non-semantic. It is just as plausible that the non-semantic facts underdetermine the semantic

facts as it is that the physical facts do. Thus the non-semantic facts do not completely determine what property is picked out by a vague predicate. The same point applies to all the other types of underdetermination we looked at above. We conclude, therefore, that understanding the use facts as the non-semantic facts leaves the points we make in §2.3 completely untouched.

§4 A different interpretation of supervenience

As noted above, some of the concerns we raised against SSUP in §2 (particularly those involving possessors, shifters, and blockers) mimic concerns that have been raised in the literature about the supervenience of the mental on the physical. For example, the possibility of non-physical ghosts which have mental properties seems to cause an obvious problem for physicalism about the mental. One response to such worries that physicalists have opted for is to deny the possibility of non-physical entities with mental properties. Many, however, consider this move implausible or poorly motivated. A different kind of response concedes the possibility of such non-physical entities, but argues that the supervenience of the mental on the physical should be understood in a way that is not inconsistent with this possibility. The rough line of thought behind this second kind of response is this: the supervenience of the mental on the physical is not a necessary truth, but a contingent one. Crudely put, the thesis is that in the actual world what brings about or is responsible for the mental properties are the physical properties. This is nevertheless consistent with the claim that in other worlds, other kinds of properties might be responsible for the mental properties. Following this line of thought

there have been various formulations in the literature trying to capture this kind of contingent supervenience.⁴³

Could SSUP be defended against our previous arguments by interpreting supervenience in terms of such contingent notions? In this section we explore this idea. In §4.1 we focus on one such prominent interpretation of supervenience and raise several reasons to think that this is not the correct way to understand SSUP. In §4.2, we show that even if SSUP is understood this way, most of our previous objections to the thesis still stand.

§4.1 Concerns with interpreting SSUP in terms of J-Supervenience.

Let us focus on one of the prominent formulation of contingent supervenience, that presented by Frank Jackson.^{44, 45} According to Jackson's formulation, A-properties supervene on physical properties (relative to a world w) if any world which is a *minimal physical duplicate* of w , is an A-duplicate of w . A 'minimal physical duplicate' (MPD) of a world w_1 , is a world w_2 that is a physical duplicate of w_1 and contains no additional 'non-physical' objects and properties, except for those which any such world *must* have. This formulation is obviously vague, but the rough idea is fairly clear: a minimal physical

⁴³ See, for instance, Chalmers 1996, Jackson 1994, Lewis 1983

⁴⁴ Jackson 1994. Most of our remarks will hold equally if one were instead to appeal to Chalmers's or Lewis's formulations (*mutatis mutandis*).

⁴⁵ Jackson formulates the supervenience thesis that *everything* supervenes on the physical. We here amend the formulation to express the more restricted claim that a particular set of properties supervene on the physical.

duplicate of the actual world would not contain ghosts (because they are neither physical, nor a necessary addition to the world) but might contain numbers (because they are necessary additions to the world). Moreover, if (and only if) physicalism in philosophy of mind is correct, then a minimal physical duplicate of the actual world would contain the mental properties of human beings, because they ‘must be added’, given the physical properties of human beings.

Combining Jackson’s formulations with our previous understanding of duplication, we get the following thesis:

J-supervenience: A-properties J-supervene on physical properties (relative to a world w) if and only if for any MPD w' of w , any isomorphism between w and w' which preserves physical properties, preserves A-properties.

Is it then plausible to interpret SSUP as the thesis that semantic properties J-supervene on physical properties? It is not clear that it is. First, it should come as no surprise by now that if SSUP is interpreted in this way, then it is at best only contingently true. But it seems to us that many defenders of SSUP would insist that the thesis is a necessary one: semantic facts are just not the kind of facts that can be fundamental in any world, and they must be derivative on use facts.⁴⁶

⁴⁶ This seems to us in contrast to the case of the supervenience of the mental on the physical, where some at least seem to be motivated to physicalism by contingent or a posteriori considerations (e.g. the empirical success of physics and biology at explaining various aspects of our world, and in particular various aspects of our mental lives).

There are other more specific problems with the proposal. Following an objection of Hawthorne's to Jackson⁴⁷, we note that SSUP as formulated in terms of J-supervenience is consistent not only with the possibility of possessors but also with the possibility of blockers and shifters: worlds which contain ghosts (whether acting as possessors, blockers, or shifters) are not minimal physical duplicates of the actual world, and thus cannot make problems for J-supervenience. Now, the fact that SSUP as formulated in terms of J-supervenience is consistent with the possibility of possessors should not count against it: after all, the purpose of reformulating supervenience in terms of J-supervenience was precisely the thought that other worlds, ones in which new semantic facts are brought about in other ways than in the actual world, need not trouble us. However, the consistency of the thesis with blockers and shifters is a lot more troubling. Consider for example a haunted world *w*, as described in §2.1, example 5. The world *w*, is a physical duplicate (though not a minimal one) of the actual world in which facts about ghosts entail that the word 'water' in Jack's mouth refer to ghost-water, in contrast to the actual world, where it refers to water. Thus there is a physical-duplicate of our world in which some positive semantic fact of the actual world fails to obtain, which suggests that the physical properties in our world are in some sense not even sufficient to determine the semantic facts of our world. A formulation in terms of J-supervenience is thus not only weaker than the claim that *new* semantic facts in other worlds might be determined by facts other than the physical use facts (this was intended); but it is also too weak to capture the claim that the *old* semantic facts of our world are completely determined by the physical use facts.

⁴⁷ Hawthorne 2002.

A different problem for the formulation in terms of J-supervenience is this: a world might have no MPDs. Suppose that any physical duplicate of a certain world *w* must have *some* non-physical object, but different duplicates can have different non-physical objects. For example, suppose that God ensures that every world has at least one non-physical soul in it, but creates different souls in different worlds. Or suppose it is indispensable to the physics of *w* that there are countably many abstract mathematical objects, but different worlds can have different sets of those objects (one can have all the even numbers, while another might have all the odd numbers). Any such case would entail that *w* has no MPDs and thus that *any* set of properties trivially J-supervenes on physical properties, which seems absurd. Nor can the problem be easily fixed by insisting that J-supervenience of A-properties on the physical only holds in a world *w*, if *w* has at least one MPD: after all, some worlds which don't have any MPDs for reasons noted above are ones in which intuitively supervenience does hold. Consider for example the number case above. It would be absurd to suppose that this fact about numbers means that in *w* mental facts cannot supervene on physical facts or that semantic facts cannot supervene on use facts.

We conclude, then, that the formulation of SSUP in terms of J-supervenience is not a promising one.

§4.2 The objections to SSUP revisited

Would formulating SSUP in terms of J-supervenience help with the objections we raised in §2? It would certainly help with many of the problems concerning blockers,

possessors, and shifters raised in §2.1. Assuming there are no ghosts in the actual world, no MPD of the actual world would contain ghosts, so none of the haunted worlds we described can make trouble for SSUP, thus understood. Interestingly, though, some of our objections still stand. For example, consider the world *w* which is a physical duplicate of the actual world, contains no additional non-physical stuff (except what is necessary), but where Jill's duplicate is a zombie. As we have seen, such a world causes problems for SSUP, and since it is a MPD of the actual world the move to J-supervenience does not help avoid the problem. Similar remarks hold with respect to the possibility of inverts. Since some shifting and blocking of semantic properties is possible without adding new non-physical objects to the world, the core of our objections still stand.⁴⁸

It is also obvious the move to J-supervenience will not help with the problem concerning haecceitistic properties (clearly those problems can be raised by MPDs of the actual world), or to the issues concerning indeterminacy and scepticism (the problems there had nothing to do with extra non-physical entities).

Thus formulating SSUP in terms of a J-supervenience is neither a promising interpretation, nor is it sufficient to block our argument.

⁴⁸ Of course, supporters of SSUP might wish to deny the possibility of zombies and inverts. But this is a controversial debate in metaphysics, and it does not seem that defenders of the possibility of zombies are typically also defenders of SSOV. It is also worth noting in this connection that the combination of views which holds SSUP while allowing for the possibility of zombies is not a blatant contradiction: it is at least *prima facie* open to those philosophers who deny that phenomenal facts are intentional (see Block 1990).

§5 Conclusion

In this paper we have argued that on any plausible formulation, SSUP cannot be correct. The endless formulations and re-formulations of the term ‘supervenience’ one finds in the literature, might eventually lead to a suspicion that there is no interesting relation expressed by the term. But if no interesting relation is expressed by ‘supervenience’, then no interesting fact is expressed by SSUP, and the thesis should be abandoned. We take it that such an abandonment would amount to accepting SSOV in spirit, even if not in letter.

Among those who have not yet given up on the notion of supervenience, many have argued that all modal formulations of the type we have explored in this paper are too weak to capture what we intuitively mean by ‘supervenience’. Our arguments in this paper, however, have attempted to show that even on such weak modal understanding SSUP fails, so if our arguments are successful, it is highly doubtful that SSUP would be true on the apparently stronger notions of dependence.

The importance of our conclusion should by now be clear. First, SSOV goes naturally with the view according to which semantic facts are fundamental rather than derivative.⁴⁹ Second, as we have argued in §2.3, adopting SSOV opens the way to new solutions for a wide range of deep problems in the philosophies of mind and language. If semantic facts do not supervene on use facts, we can avoid one of the main sources of global scepticism in the domain of semantics. Moreover, we can propose new solutions to more specific

⁴⁹ To clarify: we are not assuming that fundamentality entails failure of supervenience. But it is at least plausible to hold that failure of supervenience implies fundamentality.

difficult puzzles such as vagueness, reference to indiscernibles, Benacerraf's problem, and arbitrary reference, to name just a few. This paper, we hope, is the first step in reinstating The Semantic to its rightful crown.⁵⁰

References

Bennett, K. and McLaughlin, B. (2008): 'Supervenience', *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/supervenience/>.

Breckenridge, W, and Magidor, O. (*forthcoming*), 'Arbitrary Reference', forthcoming in *Philosophical Studies*.

Boghossian, P. (1989): 'The Rule Following Considerations', *Mind*, Vol. 98, No. 392, pp. 507-549

Burge, T. (1979): 'Individualism and the Mental', in French, Uehling, and Wettstein (eds.) *Midwest Studies in Philosophy IV*, Minneapolis: University of Minnesota Press, pp. 73-121.

⁵⁰ We are grateful to audiences at the University of Texas Austin, Logos Barcelona, Boston University, Leeds University, Manchester University, MIT, University of Oxford, and the Oxford-Paris Workshop on Language and Ontology for helpful discussion. Special thanks to Cian Dorr, John Hawthorne, Vann McGee, Josh Schechter, Robbie Williams, and an anonymous reviewer for their valuable comments and suggestions.

Cameron, R. (2010): 'Vagueness and Naturalness', *Erkenntnis* 72, pp. 281-293.

Chalmers, D. (1996): *The Conscious Mind*, Oxford University Press: Oxford and New York

Chisholm, R. (1986): 'Possibility without Haecceity', in *Midwest Studies in Philosophy XI: Studies in Essentialism*, ed. Peter A. French et al., University of Minnesota Press: Minneapolis

Davidson, D. (1987): 'Knowing One's Own Mind', *Proceedings of the American Philosophical Association*.

Fodor, J. (1986): *Psychosemantics*, MIT Press.

Hawthorne, J. (2002): 'Blocking Definitions of Materialism', *Philosophical Studies*, 110(2): 103–113

Hawthorne, J. (2006) 'Epistemicism and Semantic Plasticity', *Oxford Studies in Metaphysics Volume 2*, Oxford University Press: New York

Hofweber, T. (2005): 'Supervenience and Object-Dependent Properties', *Journal of Philosophy* 102 (1):5-32.

Jackson, F. (1994): 'Armchair Metaphysics', in O'Leary-Hawthorne & Michael (eds.) *Philosophy in Mind*. Kluwer.

Kearns, S., and Magidor, O. (2008): 'Epistemicism about vagueness and meta-linguistic safety', *Philosophical Perspectives* 22, pp. 277-304

Leuenberger, (2008a): 'Supervenience in Metaphysics', *Philosophy Compass* 3/4 (2008): 749–762,

Leuenberger, S. (2008b): 'Ceteris Absentibus Physicalism', in D. Zimmerman (ed.), *Oxford Studies in Metaphysics*, Oxford: Oxford University Press, pp. 145–170.

Lewis, D (1983): 'New Work for a Theory of Universals', *Australasian Journal of Philosophy*, 61:343-377

Kripke, S. (1982): *Wittgenstein on Rules and Private Language: an Elementary Exposition*. Harvard University Press, Cambridge, Mass.

Paull, C.P. and Sider, T.R. (1992): 'In Defense of Global Supervenience', *Philosophy and Phenomenological Research*, 32: 830-45.

Putnam, H. (1975): 'The Meaning of Meaning', *Philosophical Papers, Vol. II : Mind, Language, and Reality*, Cambridge: Cambridge University Press.

Putnam, H. (1977): 'Realism and Reason', *Proceedings of the American Philosophical Association* 50, pp. 483-498, reprinted in Putnam (1978), *Meaning and the Moral Sciences*, Routledge and Kegan Paul.

Quine, W. V. O. (1960): *Word and Object*, MIT Press

Shagrir, O. (2002): 'Global Supervenience, Coincident Entities and Anti-individualism', *Philosophical Studies* 109: 171-196.

Sider, T. (2009): 'Ontological Realism', in David J. Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.

Williams, R. (2007): 'Eligibility and Inscrutibility', *Philosophical Review* 116(3):361-399

Williamson, T. (1994): *Vagueness*, Routledge: London.

Wittgenstein, L. (1953): *Philosophical Investigations*, G.E.M. Anscombe and R. Rhees (eds.), G.E.M. Anscombe (trans.), Oxford: Blackwell