

Hyperintensional Propositions

MARK JAGO

Forthcoming in *Synthese*, special issue on hyperintensionality.

Abstract: Propositions play a central role in contemporary semantics. On the Russellian account, propositions are structured entities containing particulars, properties and relations. This contrasts sharply with the sets-of-possible-worlds view of propositions. I'll discuss how to extend the sets-of-worlds view to accommodate fine-grained hyperintensional contents. When this is done in a satisfactory way, I'll argue, it makes heavy use of entities very much like Russellian tuples. The two notions of proposition become inter-definable and inter-substitutable: they are not genuinely distinct accounts of how propositions represent what they represent. Semantic theorists may move freely between the two conceptions of what propositions are. Nevertheless, the two approaches give different accounts of the metaphysical nature of propositions. I argue that the sets-of-worlds view provides an adequate account of the nature of propositions, whereas the Russellian view cannot.

Keywords: Propositions, Hyperintensionality, Fine-grained content, Impossible worlds, Semantics

I Introduction

In contemporary semantics, approaches to propositions (and to content more generally) divide into two main camps. In the first camp, we find *structured entities* playing the role of content; in the second camp, we find sets of worlds playing that role. A prime example of the former approach is the *Russellian* notion of a proposition, on which the proposition *that Anna loves Bec* is a structured entity containing Anna herself, the relation *loving* and Bec herself, in that order (King 1995; 1996; Salmon 1986; 2005; Soames 1987; 2008). Since the structure is important to such entities, we might identify them with ordered tuples. Russellian propositions have a lot going for them. They do a great job of capturing what a sentence is *about*. The tuple (Anna, *loving*, Bec) captures the fact that 'Anna loves Bec' is about Anna, Bec and the *loving* relationship between them. That's one important role that propositions are meant to play.

The main opposition to this view comes from the sets-of-worlds approach to content (Lewis 1986; Stalnaker 1976a;b; 1984). Typically, on this view, the proposition *that Anna loves Bec* is the set of all possible worlds according to which Anna loves Bec. This entity is unstructured in the sense that it's just a set of worlds, with no order placed on them. This view, too, has lots going for it. The proposition *that A* entails the proposition *that B*, on this view, when the former is a subset of the latter. A proposition is possible when it is non-empty, and necessary when it is the set of all possible worlds. Moreover, truth-at-a-world reduces to set-membership: *that A* is true at world *w* iff *that A* contains *w*. Since entailment, possibility, necessity and truth (at a world) are properties (and plausibly, the

essential properties) of propositions, there is much to like about the sets-of-worlds account of propositions.

It seems clear that the two approaches are in conflict. Soames (1985; 1987; 2008) argues that the sets-of-worlds approach is deeply flawed (in ways that the Russellian approach is not), whereas Ripley (2012) argues for precisely the opposite position. Contrary to both Soames and Ripley, I want to argue that the two approaches are not so distinct. In particular, when the sets-of-worlds account is extended to accommodate fine-grained contents, it should make use of entities very much like Russellian tuples. The representational properties of sets-of-worlds propositions are given by entities very much like Russellian tuples, I'll claim. If this is correct, then the two views about propositions are inter-definable and inter-substitutable. Semantic theorists may move freely between the two conceptions of what propositions are. In short, the Russellian and the sets-of-worlds approaches are not genuinely distinct accounts of propositional representation.

If correct, this is a very surprising result. I will not draw the conclusion that sets-of-worlds theories of content are defunct, however. Sets of worlds play a useful role in semantics, even if much of the representational action happens at the level of Russellian tuples (or similar entities). Indeed, I'll argue that, on purely metaphysical grounds, sets of worlds provide us with a better account of the *nature* of propositions than the Russellian account can.

I shall argue for the main result of the paper by claiming that, in order to accommodate fine-grained contents, the sets-of-worlds view must work with *impossible* as well as possible worlds, thought of in a certain way. It is a key premise of this argument that fine-grained contents are desirable in a semantic theory. As I'll use the term, a theory of content is fine-grained when it allows for *hyperintensional* contents: that is, contents that are distinct but logically equivalent. The argument for fine-grained content, in outline, is simple. First, there exist genuinely hyperintensional operators, such as 'believes that', 'knows that', 'is informed that', 'it is cognitively significant that' and 'it is trivial that'. Second, the theory of content is compositional, so that the semantic value assigned to 'OA' is a function of the semantic values assigned individually to operator 'O' and to sentence 'A'. Together, these points entail that logically equivalent sentences may express distinct contents, and hence that we require a fine-grained theory of content. For the remainder of the paper, I will take it for granted that fine-grained hyperintensional contents are desirable.

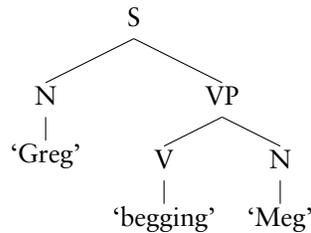
After briefly introducing the Russellian and sets-of-worlds approaches (§2 and §3), my argument will proceed by considering how the sets-of-worlds view can best accommodate fine-grained content. I will argue that it must use *impossible* as well as possible worlds. I'll consider what these worlds are and how they represent in §3 and §4, and argue that their answers push the sets-of-worlds theorist into the conclusion advertised above (§5). I'll then argue that sets of worlds still have a role to play in semantics (§6), and indeed, that they provide a better account of the metaphysical nature of propositions than the Russellian account does (§7). §8 is a brief conclusion.

2 The Russellian Account

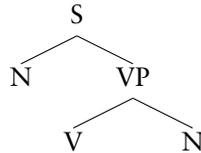
In this section, I'll briefly review what I take to be the best approach to structured propositions, which comes from King (1995; 1996) and Salmon (2005). King is concerned with the question of *propositional structure*, that is, the way in which the constituents of a proposition are structured. Given that propositions are structured entities of some kind, just what kind of structure do they have? His answer in King 1996 is that propositional structure is identical to *syntactic structure*, understood at the level of syntactic analysis of sentences at which semantic interpretation begins. (King (1995) calls this level *semantic input*, or SI). To illustrate, take a sentence

(1) Greg begged Meg

which has the syntactic structure:



On the view from King 1996, the relevant structure here is



and so the proposition expressed by (1) has the structure:



where the leaf nodes are the particulars Greg and Meg, and the relation *begging*. Here, the relation which structures lexical items in the sentence ('Greg', 'Meg' and 'begged') is *identical* to the relation which structures the semantic components of the proposition (Greg, Meg and the *begging* relation).

King (1995) presents a slightly different idea, on which propositional structure contains syntactic structure as a component. Propositional structure is the composition of syntactic structure with the semantic relations which hold between

the relevant lexical items and their semantic values. King doesn't say why he changed his view from the simple to the more complex account, and I can't see that the change is required. After all, the relation between lexical items (in context, after disambiguation) and semantic values should be one-to-one, and so the addition of semantic relations to the propositional-structuring relation do not add any relevant additional structure. (King (1995, 520–1) notes this point.)

Following King's lead (but omitting some of the details of his approaches), I'll think of structured propositions as ordered tuples (i.e., set-theoretic entities) that are structurally isomorphic to the syntax tree of the sentence in question. On this view, (1) is associated with the Russellian tuple:

(Greg, (*begging*, Meg))

Such contents are fine-grained because syntactic structures often differ between logically equivalent sentences. Let's take it for granted that this view (or something near it) is the best way to pursue the Russellian strategy. I shall now turn my attention to the alternative conception of propositions as sets of worlds.

3 The Sets-of-Worlds Account

On the sets-of-possible-worlds account of propositions, the proposition *that A* is the set of all possible worlds according to which it is the case that *A* (Stalnaker 1976b). A consequence of this account is that necessarily equivalent propositions are identical (Stalnaker 1976a, 9). Sets-of-possible-worlds are intensional, not hyperintensional, entities. The simplest way to generalise this approach to include hyperintensional contents is to generalise the nature of the worlds involved. We do this by including impossible as well as possible worlds. Impossible worlds are worlds according to which something is the case which could not possibly be the case. What an impossible world represents, in totality, could not be the case.

There are many questions surrounding impossible worlds. There are questions about their *nature*: what kind of entity are they? (These questions apply equally in the case of possible worlds.) There are distinct questions about their *granularity*: if an impossible world represents that *A*, what else must it thereby represent, or thereby not represent? Are such representations governed by some non-classical (but non-trivial) logic?

The answer to the granularity question depends on what we want worlds (and sets of worlds) to do in our theory. We might want to analyse *what is said*, or information content, or the content of epistemic and doxastic states, or a notion of *aboutness*, using sets of worlds. It might be that different kinds of world, with differing granularities, are required for these applications. A full semantic theory will need to analyse all of these notions and more. So we should adopt the principle that the application demanding the finest grain should determinate the granularity of worlds in general.

Take the case of doxastic content. One's beliefs are not closed under classical consequence, for we all have inconsistent beliefs and yet there are some things we do not believe. Even if we purged our beliefs of inconsistency, our beliefs would still

not be classically closed, for we cannot follow through all the consequences of what we hold true. (It is sometimes said that we ‘implicitly’ believe all consequences of what we believe. But unless there is some guarantee that we always avoid inconsistent beliefs, this notion is trivial. We do have non-explicit beliefs, but these do not include all consequences of one’s explicit beliefs.) So there is at least one standard inference rule with respect to which a given agent’s beliefs are not closed. There is some instance of that rule such that she believes the premises, but not the conclusion. Given that this is the case for *some* standard inference rule, is there any reason to think that beliefs in general must be closed under *other* standard inference rules? (They are trivially closed under the identity rule, $A \vdash A$, of course. By ‘inference rule’ here, I am thinking of those rules typically used to set up a natural deduction system.)

If so, it could only be because, given the *meaning* of the associated connective, anyone who believes the premises of some instance of that rule thereby believes the conclusion. But the meaning of any of the standard connectives stands to the corresponding rules just as the meaning of any other connective stands to its inference rules. So if we have an argument from the meaning of *some* connective to closure of an agent’s beliefs under the associated inference rules, then we have an argument for closure for *all* the standard inference rules. But we have already rejected that conclusion, and so we must hold that there is no such argument. Beliefs need not be closed under any of the standard inference rules (or any other inference rule, other than identity). So, if we are to capture doxastic contents via sets of worlds, those worlds too must violate all of those inference rules. In other words, what worlds represent is not closed under any inference rule (other than identity).

This on its own does not entail that, for every instance of one of the violated inference rules, there is some world which represents the premises but not the conclusion as being the case. But it would be highly strange if, say, what worlds represent about cats must be closed under *modus ponens*, whereas what they represent about dogs need not be so closed. If the meanings of the connectives do not guarantee closure (of what worlds represent) *everywhere*, then they cannot guarantee any *particular* relationship between what worlds represent. There are no particular logical connections between representations *that A* and *that B* (other than identity) such that, if world w represents that A then w must thereby represent that B , or must fail to represent that B .

In sum, worlds in general must be very fine-grained representations. What they represent is not, in general, closed under any nontrivial logical rules. This allows us to say that, if it is (logically) impossible that A , then there is an impossible world which represents that A (Nolan 1997). That impossible world may represent that A without representing obvious consequences of A . Having discussed the granularity of impossible worlds (and hence of worlds in general), it is now time to discuss the *nature* of impossible worlds (and worlds in general). Just what kind of entity are impossible worlds? This is the topic of the next section.

4 The Nature of Impossible Worlds

In this section, I will argue that impossible worlds, and worlds in general, are best thought of as arbitrary sets of sentences of a suitable world-building language. I'll argue, first, that impossible worlds cannot be *genuine* worlds, which represent (*de dicto*) that such-and-such by being that way. Instead, they must be ersatz worlds: mere representations that such-and-such is the case. I then argue that ersatz worlds are best thought of as linguistic constructions, and that the language in question is best thought of as being constructed from the very particulars, properties and relations to which it refers.

The first claim, that the worlds in question are ersatz worlds, is the least controversial of these claims. Lewis (1986) argues (convincingly) that impossible worlds cannot be *his* kind of genuine world. Lewisian genuine worlds (in *de dicto* cases) represent that *A* by being that way. If there's a flying hedgehog, according to Lewisian world *w*, then that's so because *w* contains a flying hedgehog as a part. Hence, if some world represents that there's a flying hedgehog, then there really is a flying hedgehog, *simpliciter*. Flying hedgehogs are at least consistent, in contrast to hedgehogs that are simultaneously spiky and not spiky all over. We cannot countenance the existence of such hedgehogs whilst maintaining classical (or any non-paraconsistent) logic. But if there is a world which represents that *A*, for any logical impossibility *A*, then there is an impossible world according to which there exists just such a spiky-and-not-spiky hedgehog (§3). So that world cannot be of the Lewisian variety.

The only account of genuine impossible worlds I know of is that given by Yagisawa (2010). For Yagisawa, worlds are fundamentally like times and places. Entities exist and have properties at a world *w* by having *modal stages* at world *w* which have those properties intrinsically. Just as we might say that I was happy on Monday in virtue of having a happy monday-stage, Yagisawa says that Hattie is actually spiky in virtue of having a spiky actual-stage; she could have been spineless in virtue of having a (merely) possible spineless world-stage; and she is necessarily a physical being because all of her world-stages are physical beings. These world-stages are fully real entities. Moreover, they have properties like *being spiky* or *being spineless* intrinsically. Hattie's actual-world-stage is intrinsically spiky: is it not merely spiky-at-the-actual-world. Similarly for all her other modal stages: they are spiky or spineless *simpliciter*, and not merely spiky-at-*w* or spineless-at-*u*.

Given this, the possibility of Hattie's being spineless entails that there is an intrinsically spineless Hattie-stage, somewhere in modal space. So, we can infer from 'Hattie could have been spineless' to 'something is spineless' (namely, Hattie's spineless modal-stage). By the same token, the *impossibility* of Hattie's being spiky-and-not-spiky entails that there is an *intrinsically* spiky-and-not-spiky Hattie-stage, certainly not actually, but out there somewhere in impossible modal space. That Hattie-stage is (impossibly) both spiky and not spiky *simpliciter*. So we can infer from 'it is impossible that Hattie is *F*' to 'something is *F*' (namely, one of Hattie's impossible-world stages). Such modal stages are legitimate objects of quantification, according to Yagisawa. They are as real as possible modal stages, temporal stages and spatial parts. So Yagisawa is committed to the reality

of entities – modal stages – with contradictory properties. Such entities are not merely *F*-and-not-*F*-at-*w*; they are *F*-and-not-*F simpliciter*. But logic alone forbids such entities, for the existence of something that's *F*-and-not-*F* entails *A* for any *A* whatsoever, which is absurd. So impossible worlds cannot be of the Yagisawa variety ([reference omitted](#)).

Impossible worlds are not genuine worlds, and so they must be ersatz worlds. It remains an option to adopt a hybrid approach, on which possible worlds are genuine (perhaps, Lewisian) worlds, whereas the impossible worlds are ersatz worlds. That's the position [Berto \(2010\)](#) adopts. But notice that, even if one accepts the existence of genuine possible worlds, it does not follow that they play a role in sets-of-worlds propositions. If there exist impossible ersatz worlds, then there also exist ersatz possible worlds: these are the ersatz worlds which correspond to the genuine possible worlds. So, on the hybrid account, we have a choice: sets-of-worlds propositions may comprise genuine possible worlds plus ersatz impossible worlds, or they may comprise ersatz worlds (possible and impossible) only. Call the former *hybrid propositions*.

There's much to be gained, and nothing to be lost, by avoiding hybrid propositions, and instead constructing propositions purely from ersatz worlds. In so doing, we maintain a uniform account of how propositions represent, and of how propositions are constructed from the basic stuff of reality. That's so independently of whether there are genuine worlds other than the actual world. The question of multiple genuine worlds is really about the *reduction* of modality, as [Berto \(2010, 481\)](#) emphasises. The genuine worlds tell us *which* of the ersatz worlds are to be treated as possible worlds. If it's paramount to have a reductive account of modality, then accept a plurality of genuine worlds; if not, don't. Either way, given that the impossible worlds are ersatz worlds, it's better to do semantics with ersatz worlds across the board. That's the case in favour of treating sets-of-worlds propositions as sets of ersatz worlds.

Next, I'll argue that such worlds are best thought of as sets of sentences in some world-building language, as *linguistic ersatzism* maintains. The argument here is rather simple. Worlds must be very fine-grained (§3) and the best (perhaps the only) way to get such fine-grained worlds is by adopting the linguistic ersatz approach.

Suppose, with [Lewis \(1986\)](#), that ersatz worlds represent that such-and-such either by linguistically *saying* that such-and-such or by *picturing* that such-and-such. (I am ignoring the view that worlds represent what they represent 'magically', such that no explanation of how they represent is available.) One can represent a variety of impossible situations pictorially, as Escher and other ingenious artists have demonstrated. Nevertheless, pictorial representations are never as fine-grained as we require. It is not easy to see how one can pictorially represent a man without thereby representing him as being thin, or as being fat, or as being average build. (Perhaps a cluster of pictorial representations, all differing on the build of the man in question, will do the trick. Even if so, we cannot pictorially represent that there's a thin man without thereby representing that there's a man.) More generally, we cannot pictorially represent that $A \wedge B$ without thereby representing that *A*. Pictorial representations are closed under conjunction

elimination, whereas we require worlds which are not so closed (§3). The worlds we require do not represent pictorially.

The remaining option is that worlds represent linguistically. So long as the language in question allows us to distinguish between the sentence $A \wedge B$ and the pair of sentences A, B , we can build representations which are not closed under conjunction introduction or elimination (or any other inference rule, other than identity). Those representations are merely arbitrary sets of sentences, with no closure conditions specified. Linguistic ersatzism identifies worlds with such representations, which are as fine-grained as the syntax of the language in question. (I'll assume that *all* such representations count as worlds. Of course, many will not count as metaphysically, logically, or epistemically possible worlds, but that is a different issue.)

The third point I want to establish is that the language from which such worlds are built should be something like the *Lagadonian language* suggested by Carnap (1947) and Lewis (1986). In this language, particulars, properties and relations themselves serve as names and predicates. Linguistic ersatzism re-casts the question of how worlds represent as the question of how sentences of the world-building language represent. Since we ultimately want to give semantics for various natural languages in terms of sets-of-worlds propositions, we cannot interpret the world-building language via a translation into some natural language. (We will have to state the rules for interpreting the world-building language in some natural language, of course. What we cannot do is to stipulate that sentence S of the world-building language is to mean what the English sentence S' means.) Somehow, world-building sentences must make contact with non-linguistic reality. The simplest way to achieve this is to take those lexical items to be the very things with which we want our world-building language to make contact. In particular, we do this by taking names and predicates to be the relevant particulars, properties and relations.

We don't have to accept this move. We could choose other entities (certain sets, say) to be the lexical items of the world-building language. In interpreting those items, we must link them to non-linguistic reality, and in particular, to particulars, properties and relations. Whatever entities provide the semantic values for this interpretation, we will require a world-building lexical item for each semantic value. So the interpretation must be surjective. And we must not lose structure in moving from the world-building language to its interpretation, for it is pointless having very fine-grained propositions if what they represent is of a coarse-grained nature. So the interpretation must be injective. Combining these two points, the interpretation must be a bijection from lexical items to semantic values. Moreover, we are not one jot interested in the intrinsic properties of the lexical items: we care only about what they represent. So nothing is gained by choosing lexical items distinct from their semantic values, except an extra layer of complexity. That's a strong reason for anyone adopting linguistic ersatzism to adopt the Lagadonian world-building language, or something very like it.

Through the past two sections, I've been arguing for a certain view of what impossible worlds (and worlds in general) are. In the next section, I'll consider the consequences of this view for the debate between the Russellian and the

sets-of-worlds views of propositions.

5 The Inter-Definability Argument

The direction of my argument should by now be clear. The kind of sets-of-worlds account of propositions I have been recommending analyses worlds as sets of sentences. Those sentences are taken from a language whose words are particulars, properties, relations and, more generally, the semantic values of words of the target natural language (§4). Such sentences are already looking very much like Russellian tuples (§2).

More precisely, the lexical items of the ersatz-world-building language are precisely the lexical items from which Russellian tuples are built (and those items fall into the same grammatical categories in each language). To show that the languages are in fact identical, therefore, we would need to show that they have the same syntactic formation rules. The syntactic structure of Russellian tuples is given by the syntax of the language in question, understood at the level of analysis appropriate to semantic input (§2). There's no reason for the ersatz-world-building language to adopt precisely this syntax, and so I do not claim that ersatz-world-building sentences *must* be identical to Russellian tuples. Nevertheless, there is good reason to think that, if Russellian tuples and ersatz-world-building sentences do differ, then the respects in which they differ cannot be important for semantic theory.

First, note that the logical structure of the two languages must be the same. They must contain the same quantifiers, variables, connectives and so on. Any differences in the languages must be traced to the differing syntax of their logically atomic sentences. Moreover, the atomic ersatz-world-building sentences should be isomorphic to the atomic Russellian tuples. Suppose, contrary to this, that atomic Russellian tuples draw distinctions where the atomic ersatz-world-building sentences do not. This could not be down to the constituents of those tuples, since those constituents are shared by the atomic ersatz-world-building sentences. It must be because the syntax of atomic Russellian tuples, and hence the syntax of the natural languages in question, is more fine-grained than the syntax of atomic ersatz-world-building sentences. As a consequence, natural languages will draw *syntactic* distinctions where sets-of-worlds contents do not. But each such syntactic distinction in natural language marks a distinction in some possible agent's belief state, which will not be respected by sets of worlds. To avoid this unwanted result, the atomic ersatz-world-building sentences should be at least as syntactically fine-grained as the atomic Russellian tuples.

Suppose, on the other hand, that such sentences are syntactically more fine-grained than atomic Russellian tuples. The worry here is that this extra syntactic complexity brings no benefits with it. By design, Russellian tuples capture all the syntactic distinctions required in a theory of natural language. (They probably capture many more distinctions than are necessary, since any generative theory of syntax is likely to over-generate syntactic forms.) There's nothing to be gained by adopting a finer-grained syntax for atomic ersatz-world-building sentences. So

it's highly plausible that, on the best theory of ersatz worlds, the atomic ersatz-world-building sentences will be syntactically isomorphic to the atomic Russellian tuples.

Let's grant that there is no theoretically important syntactic difference between atomic Russellian tuples and atomic ersatz-world-building sentences. Then, there is no theoretically important difference between Russellian tuples and ersatz-world-building sentences, *simpliciter*. For simplicity, we may assume that ersatz-world-building sentences are identical to Russellian tuples. (To reiterate: if they are not, the differences between them are 'mere' syntactic differences, and not of theoretical import.) This granted, both views of propositions identify worlds with sets of Russellian tuples and hence analyse *what* a world represents in terms of what the corresponding Russellian tuples represent. Moreover, sets-of-worlds propositions represent what they represent in virtue of what is represented by the worlds they contain. So, on either view of how propositions represent, Russellian tuples play the central role. Already, it seems that we do not have two genuinely distinct approaches to propositional representation. Rather, it seems that we have one basic theory, which can be cashed out in one of two ways.

There's more. The worlds we require are extremely fine-grained entities (§3): we count every set of ersatz-world-building sentences as a world. So every set of Russellian tuples is a world. What a set of worlds represents is determined by the representational commonalities between the worlds it contains. Set-of-worlds W represents that A iff its intersection $\cap W$ (a set of tuples) contains a tuple T which represents that A . Let's use the notation ' T_A ' for this tuple. Notice that $\cap W$ may be empty, in which case, W represents nothing. (If we restricted to possible worlds, $\cap W$ could never be empty, for it would always contain all necessary truths. Not so with impossible worlds in play.) If non-empty, $\cap W$ is itself a world (a set of Russellian tuples), and W represents whatever world $\cap W$ represents. Note the general difference between representing both *that A* and *that B*, on the one hand, and representing *that A and B*, on the other. This difference shows up in the context of explicit propositional attitudes, for example. (See the note on 'implicit' beliefs in §3 above.) We previously insisted (§3) that impossible worlds must respect this difference. Now, we find that we have propositions that simultaneously represent two states of affairs without representing their conjunction, as well as propositions which represent nothing.

If there is to be a unique sets-of-worlds proposition *that A* (and that is surely what we want), then we will require that set to contain *all* worlds which represent that A . There will be many sets of worlds X for which $\cap X = \{T_A\}$, and they cannot all be the proposition *that A*. The only non-ad-hoc way to single out a unique proposition *that A* is to identify it with the set of *all* worlds which represent that A . In other words, the sets-of-worlds proposition *that A* is to be defined as the maximal set of worlds for which $\cap X = \{T_A\}$, i.e., $\{w \mid T_A \in w\}$. But then, since T_A is the Russellian proposition *that A*, the Russellian and the sets-of-worlds propositions *that A* are completely inter-definable. As semantic theorists, we can move freely between the Russellian and the sets-of-worlds notions.

I've argued that the Russellian and the sets-of-worlds approaches are inter-definable and essentially equivalent to one another. Now it's time to look at

the consequences of this fact. The most obvious is that (with some specific exceptions) the problems for one of the approaches will also be problems for the other approach, and likewise, the solutions available to one approach are equally available to the other.

By way of example, it is sometimes claimed that the sets-of-worlds approach can provide a solution to Frege's problem in a way not available to the Russellian approach (Ripley 2012). The problem for the Russellian is clear. If both 'Hesperus' and 'Phosphorus' have Venus as their semantic value, then 'Hesperus is *F*' and 'Phosphorus is *F*' will express the very same proposition. This makes it hard for the Russellian to explain how 'Anna believes that Hesperus is *F*' and 'Anna believes that Phosphorus is *F*' can take opposite truth-values (as they surely can). It might appear that the sets-of-worlds approach (with impossible worlds) has an easy solution, as follows. Since it is impossible that Hesperus is not Phosphorus, there is (supposedly) an impossible world according to which Hesperus is not Phosphorus. This allows for impossible worlds which represent Hesperus but not Phosphorus, and that in turn allows for Hesperus-contents to be distinct from Phosphorus-contents (qua sets of worlds).

Things are not nearly so simple for the sets-of-worlds approach, however. *How* does an impossible world represent Hesperus without thereby representing Phosphorus? We require an explanation of how worlds represent what they represent, and the general explanation we've adopted is in terms of the interpretation of the world-building language. So, if Hesperus-representations are to be distinct from Phosphorus-representations, the world-building language must have distinct names corresponding to 'Hesperus' and 'Phosphorus'. But we also assumed that, in general, world-building names are identical to their semantic values. If we accept all that, then Hesperus-representations cannot be distinct from Phosphorus-representations. The Russellian's problem with co-referring names is equally a problem for the sets-of-worlds approach.

There are responses to this problem available, of course. Perhaps we should think of the semantic value of a name as a cluster of properties, rather than as the bearer of the name. (This is particularly plausible if we think of names as general terms.) If so, then we can account for the differing semantic values of 'Hesperus' and 'Phosphorus' in terms of different clusters of properties. The crucial point here is that, if the sets-of-worlds theory can use this (or some other) approach to differentiate between Hesperus-representations and Phosphorus-representations, then so can the Russellian. (Perhaps the resulting view would not then be a *Russellian* view. The important point concerns the view of propositions-qua-tuples, rather than the view of what we take the semantic values to be.)

In this section, I first argued that ersatz-world-building sentences are identical to, or at least very similar to, Russellian tuples. If they are not identical, then there is no theoretically important difference between Russellian tuples and ersatz-world-building sentences. As a consequence, on either view of how propositions represent, it is Russellian tuples that play the central role. Moreover, the two notions of propositions are inter-definable and inter-substitutable. As semantic theorists, we can move freely between the two conceptions of what propositions are. The Russellian and the sets-of-worlds approaches are not genuinely distinct

accounts of propositional representation.

6 Doing Without Sets of Worlds?

In the previous section, I claimed that the Russellian and the sets-of-worlds approaches are inter-definable and so each shares many of its advantages and disadvantages with the other. Each account involves assigning semantic values to the lexical items of the target language and forming syntactically structured entities from these semantic values. Let's introduce the neutral term *S-sentence* ('S' for semantic) for those structured entities. The Russellian approach identifies propositions directly with S-sentences, whereas the sets-of-worlds view adds further structure, by identifying propositions with sets of sets of S-sentences. It is tempting, therefore, to think that the sets-of-worlds view adds nothing to the Russellian view except extra complexity, which doesn't bring with it any benefit. Following this line of thought, one might be tempted to take the discussion above as an argument for preferring the Russellian approach to the sets-of-worlds one.

That's not how I interpret the conclusions reached above. Given that there exist S-sentences, there certainly exist sets of sets of S-sentences, i.e., sets of worlds. The question is, whether those sets of worlds have a useful role to play in semantics that would not be played equally well by the S-sentences themselves. As I see it, there are such roles for sets of worlds. They play a role in semantics, not because they do things that cannot be defined in terms of S-sentences, but because phenomena emerge at the level of sets of worlds which are not readily apparent at the level of S-sentences. This happens when we restrict the notion of *world* in play.

Here is one case in which content aligns very neatly with sets of worlds. Suppose, as seems right, that knowing that *A* amounts to ruling out all possibilities to the contrary. And suppose that the domain of 'all' here varies with context, so that we get the kind of contextualism about knowledge of Lewis 1996. On that view, the content of one's epistemic state varies with context. This contextual variation is most simply accounted for by identifying that content with sets of worlds, as follows. The reason the content of one's epistemic state varies as the domain of 'all' varies, on this account, is that the epistemic state in question is a subset of the domain over which 'all' ranges, i.e., it is a set of worlds. As worlds are excluded from the domain of 'all', they may be excluded from the agent's epistemic state, too, resulting in variation of the content of that state. The connection between contextual variation and variation in epistemic content is most simply explained by identifying that content with a set of worlds.

Here is another example of sets of worlds proving useful in semantics. (reference omitted), following a suggestion from Chalmers (2010), takes the epistemically possible worlds to be those worlds that are not obviously *a priori* impossible. Such worlds may be inconsistent, but not obviously so. Just which worlds count as the epistemically possible ones is a vague matter. Epistemic contents are then sets of epistemically possible worlds. We can use these epistemic contents in analysing epistemic attitudes, including knowledge, belief and states of cognitive information. One important consequence of this approach is that the content of an epistemic

state may itself be indeterminate. Our agent knows that A and, since she is rational, she knows what trivially follows from A . But she does not know all consequences of what she knows, for she is not logically omniscient. Precisely *what* she knows is indeterminate. We capture this indeterminacy best in terms of sets-of-worlds epistemic contents, rather than in terms of S-sentences. This is so even though the epistemic contents in question are defined in terms of epistemically possible worlds, themselves defined in terms of S-sentences (and rational relationships between them).

A third example of a phenomenon best analysed in terms of sets of worlds is the kind of information provided by non-trivial deductions (reference omitted). Suppose we take it for granted that valid deductions can be informative but that many, such as the trivial move from A to $A \vee B$, are not. *What* the content of an informative deduction is, is best viewed in terms of sets of epistemically possible worlds, on the model of (reference omitted). On that model, some inferences are literally uninformative: they have no information content. Other inferences genuinely have informative content: the set of possibilities associated with the inference is non-empty. And for some inferences, it is indeterminate whether they have informative content or not. These are precisely the results we should expect from a theory of informative inference. It is not clear how to get those results when working directly with S-sentences, rather than with sets of worlds.

I've given three cases in which sets of worlds are of value in semantics, even if the worlds in question are themselves defined in terms of S-sentences. So I do not think we should take the argument from §5 as reasons for aligning all notions of content with S-sentences, as some Russellians think we should (Soames 1987; 2008). *Some* notions of content are best analysed as sets of worlds. But, for all I've said so far, it is an open question whether we should analyse *propositions* as S-sentences or as sets of worlds. Either way, our theory will assign a prominent role to S-sentences, worlds (qua sets of S-sentences) and sets of worlds. So, purely from the perspective of how well our semantic theory works, it matters little which of these entities we label as the propositions. That's not to say that there are *no* reasons to plump for one option rather than the other. In the next section, I'll discuss one such reason.

7 The Nature of Propositions

Given what I've said so far, it might seem that I'm suggesting agnosticism about the nature of propositions (or perhaps that we shouldn't care too much about giving an answer). But I want to distance myself from this agnosticism. From a purely semantic point of view, there is no over-riding reason to prefer one approach over the other. Nevertheless, in this section, I will argue that there are metaphysical reasons for preferring the sets-of-worlds analysis of what propositions are.

I take it to be *of the nature* of propositions to be the kind of things that bear truth-values, that are possible or necessary, and between which entailments hold. Let's take it for granted that propositions are the entities that play these theoretical roles (or the entities that come closest, if nothing plays all those roles). How are

we to analyse claims about an entity's *nature*? In general, the question is hard to resolve. But let us not forget that, on either account of propositions under consideration, propositions are *sets*. And there is a simple and compelling account of the nature of sets: their nature is given entirely by their membership.

On this view, if $X = \{x, y, z\}$, then it is of X 's nature to be the set containing x, y, z and nothing further, for this is what fixes X 's identity. Thus, it is of X 's nature that $x \in X$, and that $\{x, y\} \subseteq X$. But it is not of x 's nature to be a member of X ; nor is it of $\{x, y\}$'s nature to be a subset of X . This asymmetry ties in with what Fine (1994) says about an object's *essence*:

It is no part of the essence of Socrates to belong to [$\{\text{Socrates}\}$]. Strange as the literature on personal identity may be, it has never been suggested that in order to understand the nature of a person one must know to which sets he belongs. There is nothing in the nature of a person, if I may put it this way, which demands that he belongs to this or that set or which even demands that there be any sets. (Fine 1994, 4–5)

Socrates fixes the identity of $\{\text{Socrates}\}$, but the converse does not hold. The crucial point is *not* merely that Socrates is a person, and the nature of persons does not make reference to sets. Fine's point is more general than that: we have such an asymmetry in the nature of objects X and Y whenever X fixes the identity of Y , but not *vice versa*. Even when X and Y are both sets, it can be that X fixes the identity of Y but not *vice versa*. Suppose we adopt the usual iterative conception of pure sets as constructions, beginning with the empty set \emptyset . Then \emptyset 's nature is given *wholly* by being *the set with no members*. The singleton $\{\emptyset\}$, by contrast, has its nature given by being the set containing only \emptyset . So, \emptyset is part of the nature of $\{\emptyset\}$, but not *vice versa*. These points support the view that, whenever $X \in Y$ or $X \subset Y$, it is of the nature of Y to have X as a member or proper subset, but not of the nature of X to be a member or proper subset of Y . Call this the *nature-of-sets thesis*.

Now we can apply this thesis to our two conceptions of propositions. On the Russellian view, propositions are identified with S-sentences and worlds are sets of propositions. Truth-at-a-world is merely set-membership: proposition $\langle A \rangle$ is true-at-world- w just in case $\langle A \rangle \in w$. So, on the Russellian view, it is of world w 's nature that it contains such-and-such propositions as members, but not of the nature of those propositions to be members of world w , and hence not of the nature of those propositions to be true-at-world- w . So the Russellian view cannot accommodate the truism that propositions, by their nature, are the things that are true or false at worlds.

The same goes for possibility and necessity. A proposition is possible when it is true-at-a-world for some possible world, and necessary when it is true-at-a-world for all possible worlds. So, for the Russellian, proposition $\langle A \rangle$ is possible when $\langle A \rangle \in w$ for some possible world w , and necessary when $\langle A \rangle \in w$ for all possible worlds w . But whether or not this is so, it is not of $\langle A \rangle$'s nature to be a member of w , and hence not of $\langle A \rangle$'s nature to be possible or necessary. On the Russellian view, it is not of the nature of propositions to be the things that are possible or necessary, contrary to what we want.

Finally, there can be no essential connection between premises and the conclusions they entail, on the Russellian view of propositions. On this view, proposition $\langle A \rangle$ is a member of $\langle A \vee B \rangle$, and so it is of $\langle A \vee B \rangle$'s nature to have $\langle A \rangle$ as a part. But this connection is not *entailment* (from $\langle A \rangle$ to $\langle A \vee B \rangle$), for the same connection holds between $\langle A \rangle$ and $\langle A \wedge B \rangle$, and the former does not entail the latter. So, it is not of the nature of S-sentences to entail one another, whereas it is of the nature of propositions to entail one another. These cases provide a strong metaphysical argument against identifying propositions with S-sentences, as the Russellian does.

The sets-of-worlds view fares much better on this score. On that view, proposition $\langle A \rangle$ is true-at- w when $w \in \langle A \rangle$. So, given the nature-of-sets thesis, it is part of $\langle A \rangle$'s nature to be true-at- w when $w \in \langle A \rangle$. Similarly, $\langle A \rangle$ is possible when it contains some possible world and necessary when it contains all possible worlds. If $\langle A \rangle$ contains a possible world, then it is of $\langle A \rangle$'s nature to contain that world; and if it contains all possible worlds, then it is of $\langle A \rangle$'s nature to contain all those worlds. So, on the sets-of-worlds view, it is of a proposition's nature to be possible or necessary.

One may object to the latter argument as follows. It may be of $\langle A \rangle$'s nature to contain a world w , and w may be a possible world, but it does not follow that it is of $\langle A \rangle$'s nature to contain some possible world. Similarly, it may be of $\langle A \rangle$'s nature to contain such-and-such worlds, which are all the possible worlds, but it does not follow that it is of $\langle A \rangle$'s nature to contain all the possible worlds. Whether some worlds that are contained by $\langle A \rangle$ are possible or impossible goes beyond what makes $\langle A \rangle$ the thing it is, and hence cannot be a part of $\langle A \rangle$'s nature. This objection has force. I will return to it below and offer a response. Before that, I want to consider what the sets-of-worlds view says about the entailment-role of propositions.

Suppose that propositions are sets of *possible* worlds only (a view I have already rejected). Then, proposition $\langle A \rangle$ entails proposition $\langle B \rangle$ just in case $\langle A \rangle \subseteq \langle B \rangle$ (and more generally, some set of propositions Γ entail $\langle B \rangle$ just in case $\cap \Gamma \subseteq \langle B \rangle$). When this is the case, given the nature-of-sets thesis, it is of $\langle B \rangle$'s nature that $\langle A \rangle \subseteq \langle B \rangle$ and hence of $\langle B \rangle$'s nature to be entailed by $\langle A \rangle$. And more generally, when $\langle B \rangle$ is entailed by Γ , it will be of $\langle B \rangle$'s nature to be entailed by $\cap \Gamma$. The sets-of-possible-worlds analysis of propositions maintains the correct connection between the nature of propositions and entailment. When we introduce impossible worlds into the account, we lose this simple account. Nevertheless, there is a story available for the generalised sets-of-worlds account.

This story begins by noting that what one means by *possible world* is subject to qualification. We often distinguish between nomic, metaphysical, logical, conceptual and epistemic possibility. These senses of possibility can be refined still further. Let's focus on the notion of logical possibility. Which logic do we have in mind? Intuitionistic logic seems to carve out a space of possibilities just as much as classical logic does, and similarly for many other well-entrenched logical systems. For each such system, there is an associated space of worlds, which we consider possible relative to that logical system. Where will this process of refinement end? Should we say that only finitely axiomatisable logics, or only those of interest to a

sufficient number of logicians, come with a corresponding notion of possibility? That treats our interests as being more important than they are. Instead, we can say that, for *any* logic – and there are infinitely many – there is a corresponding notion of logical possibility.

The picture that is emerging is that we should be as flexible as possible with notions of possibility. The most flexible approach is this: for any set of worlds X , there is a corresponding notion of possibility, $possibility_X$, which takes all and only the worlds in X to be possible. (Of course, for most sets X , $possibility_X$ will be of little or no interest to us.) Suppose we adopt this view. Then, a sets-of-worlds proposition $\langle A \rangle$ is $possible_X$ just in case $\langle A \rangle \cap X$ is non-empty, and $necessary_X$ just in case $X \subseteq \langle A \rangle$. Similarly, $\langle A \rangle$ entails $_X$ $\langle B \rangle$ just in case $\langle A \rangle \cap X \subseteq \langle B \rangle$. Given this and the nature-of-sets thesis, it follows that if $\langle A \rangle$ entails $_X$ $\langle B \rangle$, then it is of the nature of $\langle B \rangle$ that this is so. For $\langle A \rangle$ entails $_X$ $\langle B \rangle$ just in case $\langle A \rangle \cap X \subseteq B$, and when this is so, it is of $\langle B \rangle$'s nature for it to be so. Similarly, if $\langle A \rangle$ is $necessary_X$, then it is of $\langle A \rangle$'s nature to be $necessary_X$. For $\langle A \rangle$ is $necessary_X$ just in case $X \subseteq \langle A \rangle$, and when this is so, it is of $\langle A \rangle$'s nature for it to be so.

Interestingly, it does not quite follow that it is of a $possible_X$ proposition $\langle A \rangle$'s nature to be $possible_X$. $\langle A \rangle$ is $possible_X$ when $X \cap \langle A \rangle$ is nonempty, and this is not a matter of $\langle A \rangle$'s nature alone. Rather, it is of the nature of $\langle A \rangle$ and X , jointly, that $X \cap \langle A \rangle$ is non-empty. (Note that we cannot in general define ' $\langle A \rangle$ is $possible_X$ ' as ' $\langle \neg A \rangle$ is not $necessary_X$ ', as we usually do in modal logic, since we have accepted that incomplete worlds, which represent neither that A nor that $\neg A$, are $possible_X$ for some X .) Nevertheless, we can capture the original intuition that, by their very nature, propositions are the kinds of entity which are possible or impossible as follows. To say that a proposition $\langle A \rangle$ is the kind of entity which is possible is to say that there is some set of worlds X for which $\langle A \rangle$ is $possible_X$, which is so just in case $X \cap \langle A \rangle$ is nonempty, for some X . This is the case whenever there is some world $w \in \langle A \rangle$. In other words, it is of the nature of $\langle A \rangle$ to be the kind of thing which is possible just in case it is of $\langle A \rangle$'s nature to contain worlds. And this is precisely what the sets-of-worlds account says!

In sum, the sets-of-worlds approach explains how propositions are, by their very nature, the kinds of things that are true or false, possible or necessary, and which entail one another, whereas the Russellian approach does not. That's a strong reason for thinking that our concept of *proposition* picks out sets of worlds, rather than Russellian tuples.

8 Conclusion

Fine-grained sets-of-worlds propositions make use of impossible as well as possible worlds (§3). The best theory of such worlds analyses them as sets of Russellian tuples, or entities very similar to Russellian tuples (§4). The representational properties of such worlds, and hence of sets-of-worlds propositions, ultimately derive from those Russellian tuples (or from very similar entities). On the sets-of-worlds view, as well as on the Russellian view, Russellian tuples play the central representational role. The two notions of propositions are inter-definable and

inter-substitutable (§5). As semantic theorists, we can move freely between the two conceptions of what propositions are.

Nevertheless, some notions of content are best analysed as sets of worlds (§6). Moreover, the sets-of-worlds approach explains how propositions are, by their very nature, the kinds of things that are true or false, possible or necessary, and which entail one another, whereas the Russellian approach cannot do this (§7). That's a strong reason for thinking that our concept of *proposition* picks out sets of worlds, rather than Russellian propositions.

References

- Berto, F. (2010). Impossible worlds and propositions: Against the parity thesis, *The Philosophical Quarterly* 60(240): 471–486.
- Carnap, R. (1947). *Meaning and Necessity*, University of Chicago Press.
- Chalmers, D. (2010). The nature of epistemic space, in A. Egan and B. Weatherson (eds), *Epistemic Modality*, Oxford University Press.
- Fine, K. (1994). Essence and modality: The second philosophical perspectives lecture, *Philosophical perspectives* 8: 1–16.
- King, J. (1995). Structured propositions and complex predicates, *Noûs* 29(4): 516–535.
- King, J. (1996). Structured propositions and sentence structure, *Journal of Philosophical Logic* 25(5): 495–521.
- Lewis, D. (1986). *On the Plurality of Worlds*, Blackwell, Oxford.
- Lewis, D. (1996). Elusive knowledge, *Australasian Journal of Philosophy* 74(4): 549–567.
- Nolan, D. (1997). Impossible worlds: A modest approach, *Notre Dame Journal of Formal Logic* 38(4): 535–572.
- Ripley, D. (2012). Structures and circumstances: Two ways to fine-grain propositions, *Synthese* DOI: 10.1007/s111229-012-0100-4.
- Salmon, N. (1986). *Frege's Puzzle*, MIT Press, Cambridge, MA.
- Salmon, N. (2005). *Metaphysics, mathematics, and meaning*, Oxford University Press, New York.
- Soames, S. (1985). Lost innocence, *Linguistics And Philosophy* 18: 59–72.
- Soames, S. (1987). Direct reference, propositional attitudes and semantic content, *Philosophical Topics* 15.
- Soames, S. (2008). Why propositions cannot be sets of truth-supporting circumstances, *Journal of Philosophical Logic* 37(3): 267–276.
- Stalnaker, R. (1976a). Possible worlds, *Noûs* 10(1): 65–75.
- Stalnaker, R. (1976b). Propositions, in A. MacKay and D. Merrill (eds), *Issues in the Philosophy of Language*, Yale University Press, New Haven, pp. 79–91.

Stalnaker, R. (1984). *Inquiry*, MIT Press, Cambridge, MA.

Yagisawa, T. (2010). *Worlds and individuals, possible and otherwise*, Oxford University Press.