A theory of metaphysical indeterminacy

Elizabeth Barnes and J. Robert G. Williams*

(February 8, 2010)

Contents

*We have presented and discussed material from this paper in various guises in many situations. We gratefully acknowledge the support of the British Academy for their sponsorship of the workshop *Metaphysical Indeterminacy: the state of the art*. We owe thanks to everyone —too many to name—who has given valuable comments, criticism and advice. We'd particularly like to thank Ross Cameron, Matti Eklund, Andrew McGonigal, Jason Turner, Rich Woodward, and Dean Zimmerman.

Introduction

In this paper, we aim to provide a theory of metaphysical indeterminacy. But there are many different projects such an ambition could encompass, so it is important to be clear about what we're trying to do. Leaving aside the particulars of *metaphysical* indeterminacy, the subject matter of a theory of *indefiniteness* more generally is a familiar phenomenon: the indefiniteness that attaches to the claim *that guy is bald* when we point at a borderline-bald man, and a host of related examples. Naturally, the project of providing a theory of this phenomenon is multi-faceted. The start of an open-ended list would include:

1. Describing what indefiniteness is: its nature or source.
2. Describing how to *reason* with the notion of indefiniteness, or in the presence of indefiniteness: what the logic of indefiniteness is.
3. Describing the cognitive role of indefiniteness: e.g. what impact knowledge that p is indefinite should have on one's opinion as to whether p.
4. Describing paradigm instances of indefiniteness; if there are multiple kinds of indefiniteness (individuated by nature, logic or cognitive role), describing which account is applicable to which case.
5. ...

We do not have the space to address all such issues here. But since our concern is to set out the foundations of a theory of metaphysical indeterminacy, (1) and (2) are the most crucial. That is, we think that to undertake a theory of metaphysical indeterminacy one must *at least* say what metaphysical indeterminacy consists in and describe its logic. These are our targets for the paper, and we take it that in meeting them we will, in a minimal but important sense, have provided a theory of metaphysical indeterminacy.

From varying answers to challenge (1), the familiar classifications of theories of indeterminacy or indefiniteness flow. If a theory says that indefiniteness is a certain kind of in-principle ignorance, we're likely to call the overall package an *epistemic* theory of indefiniteness. If a theory says that indeterminacy is sourced in the lack of semantic conventions governing certain terms, then we're apt to describe it as *semantic*.

Our project here is to describe and develop a theory of *metaphysical* indeterminacy. Accordingly, Part I addresses the first issue: what metaphysical indeterminacy is. We argue (section 1) for the legitimacy of a *primitivist* conception of indeterminacy, where indeterminacy itself is metaphysically fundamental, and thus not reducible to anything more basic. We also (section 2) briefly address some concerns about whether one can even grasp the concept of metaphysical indeterminacy, so characterized.

Thus our positive story about the nature of indeterminacy is short and sweet. But just in virtue of this, it is compatible with many different answers to the other facets of theory on our list. In particular, one might combine primitivism with either a classical or a non-classical logic (that indeterminacy is primitive does not itself weigh in either direction). In Parts II and III we describe a *fully classical and bivalent* logic and semantics for metaphysical indeterminacy. We proceed iteratively: building from an intuitive picture of how indeterminacy should interact with other theoretical commitments (part II), to a fully explicit model theory for a language including indeterminacy, and finally to a fully explicit model theory for a language including indeterminacy and modals (part III). In an appendix, we discuss the relationship between *de re* vagueness and metaphysical indeterminacy.

Part I
What is metaphysical indeterminacy?

1 The nature of metaphysical indeterminacy

We propose to provide a theory of metaphysical indeterminacy. However, we will not be attempting to offer any kind of *reduction* or *analysis* of metaphysical indeterminacy. Indeed, we are sceptical as to whether metaphysical indeterminacy admits of reduction. To appreciate the significance of this, let us begin by a brief sketch of some philosophical background.

First, terminology. We will use 'indefiniteness' to express a generic, pretheoretic notion. That it is indefinite whether our friend is bald, for example, is a datum for philosophical theorizing.

One aspect of giving a theory of indefiniteness is to say *what indefiniteness is*. Some famous examples:

> (SI) It is indefinite whether P iff (i) P is true on some classical interpretations of our language fitting with verdicts about 'clear cases' and 'conceptual truths'; and (ii) P is false on other such interpretations.

> (EI) It is indefinite whether P iff it is unknowable whether P, as knowledge whether P would violate metalinguistic safety principles.

The first is an approximation of the proposal that Fine (1975) develops. The second is an approximation of the proposal advocated by Williamson (1994).

Such specific accounts of the nature of indefiniteness can then be sorted into kinds according to salient features that feature in the account. Thus, *epistemicists* about indefiniteness will give an account of the nature of indefiniteness making central appeal to epistemic notions such as knowledge—Williamson being a paradigmatic example. *Semantic* theorists of indefiniteness make central appeal to semantic and/or metasemantic notions: interpretations, meaning-fixing facts etc.[1]

'Indefiniteness' is the generic term for the phenomenon. Indefiniteness which is underwritten by an epistemicist theory (such as Williamson's) we call e-indefiniteness. We use 'indeterminacy' as a label for non-epistemic kinds of indefiniteness. Indeterminacy that is underwritten by a semantic account (such as Fine's) we call s-indeterminacy.[2]

This paper investigates the claim that (some cases) of indefiniteness are (at least in part) cases of *metaphysical* (m-) indeterminacy. But what do we mean by this?[3]

[1] Does this mean that anyone agreeing with the 'minimal definition' of vagueness in Greenough (2003) counts as an epistemicist, given that this definition is given in epistemic terms? No: for even if one agrees that there are necessary and sufficient conditions for a term to be vague formulable in epistemic terms, one need not agree that those conditions tell us what the nature or source of the vagueness is.

[2] Our reason for separating indefiniteness from indeterminacy is purely terminological. Some philosophers are reluctant to admit that e-indefiniteness should be called 'indeterminacy'. So we use the more inclusive label 'indefiniteness' for the generic concept, and label some subspecies of that concept as 'indeterminacy'. But if you're happy to call what the epistemicist is talking about 'indeterminacy', then it's fine to label the generic concept 'indeterminacy'. No matter what term is applied—'indefiniteness' or 'indeterminacy'—we think there's a single, unifying generic concept

[3] We do not wish to be 'imperialistic' and claim that every instance of indefiniteness is a case of metaphysical indeterminacy: we are quite happy to think that there will also be semantic and perhaps even epistemic indefiniteness also. Moreover, we do not want to be committed to the claim that a given case of indefiniteness has a *single*

Following the pattern above, it is natural to expect the friend of metaphysical indeterminacy to provide an analogous biconditional, setting out what *metaphysical* indeterminacy reduces to. But there is an alternative. She may simply deny that facts about metaphysical indeterminacy are reducible to any more basic facts—i.e., that facts about metaphysical indeterminacy (if there are any) are fundamental.[4] The view is analogous to positions familiar in the metaphysics of time and modality: views that take *contingency* or *tense* as fundamental aspects of reality, not to be reduced to putatively 'more basic' ingredients.

By contrast, the two more familiar options for characterizing indefiniteness build in from the start a certain kind of reductive ambition. In each case, there are more basic facts in play (what expressions mean, what is knowable) and the aim is to show how indefiniteness can be accounted for in terms of the respective reductive bases. On our favoured view, the putative subject matter of metaphysical indeterminacy is strikingly disanalogous, in this respect, to that of either epistemic or semantic indeterminacy. To parallel the earlier biconditionals, we would have to identify a reductive basis of recognizably *metaphysical* facts, and then try to characterize necessary and sufficient conditions for m-indeterminacy in terms of what such facts obtain. While we have a rough sense of what semantic facts are (facts about language-use, meaning, and so on) and what epistemic facts are (facts about knowledge, ignorance, justification, and so on), what are we to make of the idea of distinctively *metaphysical* facts?[5] Perhaps this is one source of the scepticism one finds towards the notion of metaphysical indeterminacy in much of the contemporary literature. If we conceive of the name as reflecting a reductive project parallel to those of epistemic or semantic theorists, it's just not clear what friends of metaphysical indeterminacy have in mind.[6]

To sum up. Every theory of indefiniteness needs to say what (on their view) the nature of indefiniteness is. One stands out as distinctive: the position that says that indefiniteness is metaphysically primitive. This, we think, is appropriately classified as a metaphysical account of indeterminacy, and it is this account that we will explore in what follows.

2 Conceptual matters

Just as some theorists take *modality* or *tense* to correspond to irreducible aspects of reality (resisting reductions to truth-at-worlds or truth-at-times) we take metaphysical indeterminacy—if there is any—to correspond to an irreducible aspect of the world. However, one might think that the *metaphysical* simplicity of our proposal gives rise to *conceptual* difficulties. One common complaint among those sceptical of metaphysical indeterminacy is that they cannot *understand* the notion, or that they suspect it *makes no sense*. And in the absence of anything illuminating to say by way of explaining what m-indeterminacy *consists in*, one might think they had a point.

source. Hence the hedges in our characterization of the position.

Furthermore, we are not committed to thinking that the trichotomy of epistemic, semantic and worldly sources for indefiniteness exhausts the theoretical options—though they are the ones we see as the current main contenders.

[4]Different conceptions of metaphysics can express this in various ways. In the generalization of Lewis (1983) favoured by Sider (2009), we might say that the operator *indeterminately* is *perfectly natural*, and it is this which 'm-indeterminate' expresses. In Dummettian terminology, we might say that sentences about metaphysical indeterminacy are barely true: true, but not true in virtue of anything else. And so forth.

[5]Suppose one thought that facts about knowledge or semantics were metaphysically basic: would that make facts about meaning metaphysical facts? It's not clear to us how to begin to address such questions.

[6]One candidate is the account defended by Akiba (2004). On his account, reality contains not only spatial and temporal dimensions, but also a 'precisificational' dimension. If P holds in one such dimension, but fails to hold in another, it is indeterminate. Just as a reduction of tensed facts to tenseless facts about the happenings in a temporal dimension to reality might be called a 'metaphysical reduction' of tense, this might deserve the name a 'metaphysical reduction' of indefiniteness.

In this section, we argue that our opponents *can* make sense of what we're saying.

One advantage of providing an constitutive account of a contentious notion in independently understood terms is that it provides one's audience with a unproblematic route to understanding what is being said. For example, if you have the concept of knowledge, and acquaintance with some of the semantic and modal notions in terms of which Williamson frames his characterization of e-indefiniteness, then there is no room for incomprehension; one can just follow the definition.

But given we think there is no reductive characterization in the offing for m-indeterminacy, we have no such simple response to one who claims not to understand the notion. We can't point to a reductive definition to *force* our audience to admit they understand our starting point as Williamson could.

However, we think that we have an adequate replacement. All parties should admit that they have a grasp on a generic notion of indefiniteness (and related notions) as deployed in ordinary speech, and used informally in philosophy.[7] This generic concept of indefiniteness is arguably all one needs to have a working understanding of our target notion.[8] In particular, using it we can formulate the following biconditional:

> it is metaphysically indeterminate whether p iff (1) it is indefinite whether p, and (2) the source of this indefiniteness is the non-representational world

We would reject any reading of this as a story about the *nature of m-indeterminacy*. However, we think that all parties can *understand* the terms on the right hand side, and thereby come to have a working understanding of the theoretical term m-indeterminacy. (If (2) is found problematic, we have other options—we could replace it by the claim that the indefiniteness corresponds to an irreducible aspect of reality).[9]

Opponents may try and reject the idea that they have a truly generic notion of indefiniteness. They may claim that while they understand what it would be for something to be indefinite in an epistemic or semantic sense, their concept of indefiniteness in general is just a disjunction of these particular cases. If 'it is indefinite whether p' just means 'it is either s-indeterminate or e-indefinite whether p', and it would follow immediately that there can be no m-indeterminacy.

It's not terribly plausible that our understanding of indefiniteness is merely disjunctive. Unlike the more particular concepts of e-indefiniteness, s-indeterminacy and m-indeterminacy, the generic concept is something we work with in everyday life. Young children (at least with a little assistance) can classify things as indefinite or otherwise: but its surely implausible to credit them with either an understanding of specific accounts of the source of indefiniteness, or even of the broad categorizations. The epistemic and semantic readings seem to us pieces of theory

[7] If conceptual groundclearing is needed in order to allow someone to get a grip on this, we could point to paradigm cases.

[8] We should, of course, allow for the possibility of 'mixed cases'—instances of indefiniteness which are, e.g., both semantic and metaphysical. The claim of metaphysical indefiniteness should be construed minimally: at least some of the indefiniteness in question has its source in the non-representational world.

[9] While we have some sympathy for scepticism over 'source'-talk, notice that in giving this we are at least appealing to a *general-purpose* metaphysical notion. Nor would we be alone in claiming that a generic notion can potentially have multiple sources. Compare, for example, views according to which 'absolute necessity' is a single, uniform concept, but can still have multiple sources. Compare Fine (2005).

One potential deflationary way of understanding this 'source'-talk in the present case is in terms of the truth of certain counterfactuals. The idea, roughly, is that a case of indefiniteness with respect to S is at least partially metaphysical iff: were all the representational content of S precisified, there might still be non-epistemic indefiniteness with respect to S (i.e., indefiniteness which is not reducible to facts about knowledge). See Barnes (forthcoming) for elaboration. In the same paper, Barnes proposes a second, less metaphysically neutral elaboration, in terms of the truth-makers of S.

suggested as an account of what underlies generic indefiniteness, not concepts out of which generic concept of indefiniteness is built.

Moreover, there are several models we could provide for the relationship between the generic concept of indefiniteness and the more specific e-indefiniteness, s-indeterminacy etc. Here is one we find congenial: we should characterize generic indefiniteness via its conceptual or functional role, consisting perhaps of the characteristic attitudes and hedged responses that indefiniteness prompts.[10] Anything that fulfills this conceptual role will (prima facie) count as a case of indefiniteness. But, like any role-functional concept, it's open that the concept can be multiply realized. So it could be that, for example, semantic indecision most often fills the conceptual role of indefiniteness. Indefiniteness with respect to "mass" or "bald" are realized by facts to do with the word-usage of natural language speakers. Yet there could still be cases where the same generic conceptual role of indefiniteness is realized by the non-representational world. The role-functional conception of generic indefiniteness fits nicely with our recipe for latching onto m-indeterminacy.

Of course, to fully work out a story about the concept of indefiniteness along these lines, one would need to spell out the required functional role to outline the respects of similarity binding together different forms of indefiniteness. But so long as *something* along these lines ultimately works, our use of the generic notion in characterizing m-indefiniteness will be legitimate—the details of how exactly the generic concept works will not matter for what follows, so we will not pursue this further.[11]

It is also significant to note what our characterization of m-indeterminacy is *silent* about. It says nothing about objects with indeterminate spatio-temporal extension, objects which indeterminately instantiate properties, indeterminate identity, *de re* indeterminacy, or things which are "neither true nor false". All of the above have previously been taken as definitional of what metaphysical indeterminacy *is*.[12] Though we think some or all such cases are potential manifestations of metaphysical indeterminacy, none of them are built into the very understanding of

[10] It's an interesting question what this role is. Some characterizations put substantive constraints on the rest of the theory. A *good* way to argue against m-indeterminacy would be to (a) describe the theoretical role of indefiniteness; (b) argue that m-indeterminacy as we are conceiving it cannot satisfy that role. For suggestions about the characteristic conceptual role that might point in directions that are in tension with the view developed here, see Field (2000) and Wright (2001). (NB: Field's account is intended to characterize a the generic notion of *indeterminacy* rather than indefiniteness in general (i.e. it is supposed to distinguish this concept from Williamson's e-indefiniteness).)

The literature on 'defining vagueness' might be seen as having similar ambitions. However, often vagueness is analytically linked to sorites-like phenomena, and m-indeterminacy need not be. See, for example, Greenough (2003) and Smith (2005).

[11] The primitivist view of the nature of indeterminacy is also compatible with the view—defended by Barnett (2009) that indefiniteness (or vagueness) is itself a primitive, irreducible *concept*. But we are not committed to this kind of primitivism about the concept of indeterminacy. Just as a functionalist about the mind may think that the concept of pain is analyzable in functionalist terms, while picking out a sui generis phenomenon, the concept of indefiniteness may be analyzable while the aspect of the world it picks out is sui generis. (Barnett's speculation that there may be unanalyzable 'ur-vagueness' which is the source of everyday vagueness, may be closer to the primitivism we are interested in. It should be noted that he uses the term 'indeterminacy' in a way different from us.)

[12] Examples of these kind of definitions are often found in the introductory glosses of various papers concerned with 'metaphysical indeterminacy' or 'ontic vagueness'. Sometimes arguments described as attacking metaphysical indeterminacy or ontic vagueness really take aim at some much more specific thesis. For illustrative examples, see the discussion in: Merricks (2001), Eklund (2008), Rosen & Smith (2004), Heller (1996). For the identification of vagueness in reality with vagueness *de re*, see Williamson (2003). For a discussion of the terminological issues here, see Williams (2008a).

One way that substantial theses might be built into the notion of m-indeterminacy is if they were built into the generic notion of indefiniteness itself. What we wish to disclaim is the thesis that any of the above are *distinctive* commitments of the friend of m-indeterminacy.

the phenomenon.

At this stage, we think we've said enough about the concept of metaphysical indeterminacy to proceed with more theorizing about how it works. What about those skeptics who still maintain that metaphysical indeterminacy can't be properly understood or doesn't make sense?[13] Read in the most literal minded way, we think that such objections are just wrong—we have in this section argued that on minimal assumptions everyone can understand our notion of metaphysical indeterminacy (and of course, understanding the notion is compatible with thinking it has no application).

Perhaps a more charitable reading would represent the sceptics as asking for a story about what metaphysical indeterminacy *consists in* or *holds in virtue of*—a metaphysical reduction. To repeat: the view to be explored here is that nothing like this is available. If understanding metaphysical indeterminacy does in fact depend on being able to provide metaphysical (or conceptual) reduction then our project does not help allay any such skeptical worries. But we want to see the *argument* that reduction is necessary—without some further elaboration, it seems to us mere dogma.

[13]see especially Sainsbury (1994) and Lewis (1993)

Part II
A modal framework for metaphysical indeterminacy

To this point, we have addressed the first item on the agenda of a theory of metaphysical indeterminacy: the question of what indeterminacy is. On the view advocated above, m-indeterminacy itself is metaphysically fundamental.

Our view of the nature of m-indeterminacy is short and sweet: there is no 'more fundamental' story to be told about the nature of indeterminacy. But of course, in giving this story we have addressed only one facet of the overall theory of indeterminacy. In particular, we have said nothing as yet about what the *logic* of this notion should be.

Often, metaphysical indeterminacy is associated with revisionary logics, non-bivalent semantics and the like. We have no reason for thinking that primitivism itself rules out this sort of approach. For example, you could construe metaphysically fundamental indeterminacy as governed by the sort of nonclassical logic that Hartry Field has developed in recent work.[14] Given this, when p is indeterminate, excluded middle for p would fail.[15] Field's logic is only one among many non-classical options. The association of metaphysical indeterminacy with non-classical logic is a familiar one, and is defended by prominent advocates of metaphysical indeterminacy, *inter alia* Peter van Inwagen (1990) and Terence Parsons (2000) (these theorists differ among themselves as to how best to develop the non-classical logic and semantics).

But equally, there's nothing in the thesis that indeterminacy is primitive that *forces* non-classicism upon us. Describing the logic of metaphysical indeterminacy is independent, we think, of everything we have said so far.

So we are in the following situation. An overall package that deserves the name 'a theory of metaphysical indeterminacy' has several facets. One non-optional component is an account of the nature or source—and this we have done. Another non-optional component is an account of the logic and semantics of indeterminacy. And if our account of the latter is underdetermined by what we say about the former—as we think is the case here—then we need to add an explicit account of the logic/semantics into the package. The account to be developed here, in contrast to those mentioned previously, will be fully classical and bivalent.[16]

This is the project for the remainder of the paper. We kick off in Part II with a 'big picture' description of how primitivism can be developed and brought into relation with prior theory (in particular, an account of ersatz worlds) in a way *prima facie* consistent with classical logic and bivalence. We motivate a view of indeterminacy as a particular kind of *modality*. But while this material provides a way of conceiving and reasoning about metaphysical indeterminacy compatible with classicism, it does not give us an explicit specification of a logic for indeterminacy.[17] In Part III, therefore, we go on to work through an explicit model theory for a language

[14] See, for example, Field (2003).

[15] This follows in Field's logic from the following: $p \vee \neg p \models Dp \vee D\neg p$.

[16] We find classical approaches to indeterminacy attractive—at least as a starting point. Classical logic is simple and attractive, and given that many of our best extant theories presuppose classical logic, we have a lot of reconstruction to do if we give it up. For these reasons, an 'innocent until proven guilty' methodology is attractive (that is, you should keep classical logic unless you're forced not to). We're also motivated to explore a classical approach to indeterminacy for purely dialectical reasons. If we can develop a fully classical and bivalent model of metaphysical indeterminacy, then there is no argument from classical logic against metaphysical indeterminacy (contra what many assume).

[17] By 'logic' here we mean a specification of the consequence relation for the language (rather than, for example, a proof procedure). We shall use model-theoretic methods to specify the logic.

containing our primitive indeterminacy operator (and other things besides). By the end of part III we will be in a position to claim to have discharged two principle components of a theory of metaphysical indeterminacy: its nature and logic.

3 A modal framework for metaphysical indeterminacy

We think that metaphysical indeterminacy consists in a fundamental kind of unsettledness in the world. When *p* is metaphysically indeterminate, there are two possible (exhaustive, exclusive) states of affairs—the state of affairs that p and the state of affairs that not-p—and it is simply unsettled which in fact obtains. No further explication is possible or needed. A primitivist about indeterminacy might take a different line: claiming that there is a tripartite (exhaustive, exclusive) division amongst states of affairs: the state of affairs that p, the state of affairs that $\neg p$, and, incompatible with either, the state of affairs of p being indeterminate. This is not the conception to be pursued here, and leads to quite different pictures of primitive indeterminacy.[18]

To illustrate the claims, consider Fred. Fred is a foetus at an intermediate stage of develop—and let's suppose that it is (in our sense) indeterminate whether he has the cognitive capacities that would make him a person. What is the relationship between Fred and the property of *being a person*? We say: not that Fred fails to instantiate *being a person* and instead has a special *kind-of-a-person-but-kind-of-not* property. We say: not that Fred fails to instantiate *being a person* and instead bears a sui generis *kind-of-instantiating-and-kind-of-not* relationship to that property. Rather, the ontology just contains Fred and the property of *being a person*. Each thing either instantiates the latter or it doesn't. But in some cases—like Fred's—it is indeterminate which of the two polar options obtains.

We think that a very natural way of capturing (and, in the end, formalizing) this thought is via appeal to a broadly ersatzist theory of possible worlds. Familiarly, on this conception, possible worlds are abstract objects which represent (classically complete) ways the world might be. On standard ersatz theory, there is a single world which is *actualized*. The actualized world is the abstract object which represents things exactly as they are in the reality consisting of us and our surroundings.[19]

Once we have the set of all the abstract possible worlds in place, we can carve it up in various ways—we can look at those worlds which are metaphysically possible, the worlds which are nomologically possible, the worlds which are epistemically possible, etc. What we want to do here is introduce another division amongst the worlds that will single out a new modality—the worlds which are precisificationally possible.[20] Using our primitive notion of determinacy, we can characterize what a world has to be like to be precisificationally possible: it must be one that does not determinately misrepresent reality.[21]

Naturally, if the actual world doesn't contain any indeterminacy, then there will only be one world in the space of precisifications.[22] If everything is determinate then one (and only one)

[18]In particular, our view is compatible with the thesis that (at least for claims formulated in perfectly natural vocabulary) that the state of affairs that p obtains iff p. The defender of the tripartite classification would either have to give up this plausible claim, or give up on the law of excluded middle.

[19]Concrete reality is standardly dubbed 'the actual world'—thus the actualized world and the actual world are different things.

[20]Compare Akiba (2000), who also urges a view of indeterminacy as a kind of modality. Akiba's positive view of what indeterminacy is is different from ours, however.

[21]If one will accept propositional quantification, then w will be a precisificational possibility iff $\forall p((w \text{ represents that } p) \rightarrow \neg D\neg p)$. These are the worlds that in Williams (2008b) are called the 'actualities'. Notice that here we do not define truth as truth at every precisificational possibility, which would lead to truth-value gaps.

[22]For the sake of argument, we assume a standard ersatz theory rejection of distinct indiscernible worlds. Should

world is a candidate to be actualized, given the way things are—and that's just whichever world represents exactly how things are.

There will be more than one world in the space of precisifications just in case there is indeterminacy in reality. Again, consider indeterminacy with respect to p. If it is fundamentally unsettled whether p, there are two candidate representations for actualization—the abstract world which represents that p, and the abstract world that represents that $\neg p$. Neither of these are determinately correct, but neither is determinately incorrect, because in reality it's simply unsettled whether p or rather $\neg p$ obtains. If there is fundamental unsettledness in the actual world, then there will be no determinately correct way of representing how things are in reality. Given this basic idea, a lot of things fall out quite naturally.

Once we have a space of precisificational possibilities, we can invoke the tools familiar from other precisificational theories of vagueness—though our 'precisifications' will be worlds rather than interpretations of a language. As is standard, for any p, p will be determinate if all the worlds in the space of precisifications represent that p. Similarly, p is indeterminate if some worlds in the space of precisifications represent that p and some represent that $\neg p$. That is, p is indeterminate just in case the precisificationally possible worlds disagree over whether p is the case.

Importantly, given our picture of indeterminacy, all the worlds in the space of precisifications are themselves maximal and classical. For any p, each precisification will opt for one of p or $\neg p$, and thus, every precisification will represent as true instances of excluded middle, $p \vee \neg p$—and similarly for every classical tautology. It is only because *multiple* precisifications are involved that overall our model does not settle one way or another whether p—while all precisifications agree that, for example, $(p \vee \neg p)$ holds, they disagree over which disjunct makes it the case that the disjunction holds (thus the individual disjuncts are themselves still indeterminate).

In constructing the space of worlds, the initial thought was that there was a unique actualized world—a world which represented reality correctly. It is natural for us to follow standard erstazist treatments and think of what is True as corresponding to what holds at this unique actualized world. No matter which world is actualized, that world either represents that p, or it represents that $\neg p$—and by the characterization of Truth, either p will be True, or $\neg p$ will. Thus, not only do we have the law of excluded middle (along with every other classical tautology) we also have bivalence. When matters are metaphysically indeterminate, it is indeterminate which world is actualized—and hence it will not be settled which of p or $\neg p$ is True, though one or the other must be.[23] In a certain sense, then, we can agree that *there is a precise way that things are*—so long as by 'precise' one means that for every p, either p or $\neg p$, and either p is True or $\neg p$ is. We say that all this is compatible with it being primitively indeterminate which precise way things are.

The results we get—a fully non-revisionary, bivalent account of indeterminate language—may remind some of the kind of treatment of *semantic* indeterminacy favoured by McGee & McLaughlin (1994). They start from the idea that our thoughts and practices don't determine which among various equally good rivals is *the* interpretation of our language. They suggest that we can take there to be one correct interpretation, though metasemantics doesn't fix which

this assumption be rejected, the claim is simply that every world in the space of precisifications will represent exactly the same thing.

[23] If there are two precisifications of the actual world, w and w*, w says that w and not w* is actualized, whereas w* says that w* and not w is actualized. Each represents a single world as actualized - so determinately one world is actualized. But they disagree over which world is actualized - so it's indeterminate which world is actualized, though it's determinate that only one is. This follows on from the standard ersatzist idea that each world represents itself as being actualized, and allows us (via our precisificational framework) to build indeterminacy into the basic ersatz picture of actualization.

interpretation this is. Since any interpretation will be bivalent, they claim a non-revisionary account of semantic indeterminacy.[24]

A central challenge faced by such approaches is laid out in Williamson (1994, ch.5.). The notion of determinacy—of thoughts and practices *determining* one interpretation over others as intended—features apparently irreducibly in the characterization of this position. But it is then far from clear why it should be described as putting forward a *semantic* account of indeterminacy; we haven't got a clean reduction of determinacy to independently motivated semantic notions, and various attempts to replace appeal to it by appeal to supervenience or other general-purpose notions seem unpromising. Absent further explanation, one might interpret 'determinacy' just as the epistemicist's 'e-definiteness', and one would have a reading of these words compatible with there being facts of the matter about which interpretation is intended.

Our approach can meet the analogous challenge. Our framework does not have the ambition of reducing determinacy to independently understood semantic terms—so obviously the central charge cannot get going. But more positively, we are quite explicit that we regard *determinacy* as itself bedrock, and that we are *using* this fundamental operator to characterize the notion of precisificational possibility. Someone might try a hostile interpretation of the characterizations above—reading our 'indeterminacy' as the epistemicist's 'e-indefiniteness'—but in doing so they would be obviously ignoring the clear intent, for they would be treating as holding in virtue of facts about knowledge and ignorance what the position says is metaphysically brute—facts about indeterminacy.

If we're not offering a reduction though, what is the point of all this elaboration? Obviously we're not attempting to (conceptually or metaphysically) 'analyze' the notion of metaphysical indeterminacy in terms of the presence of more than one world in a modal space of precisifications—if nothing else, given the way we characterized precisificational possibility, this would be blatantly circular. Rather, we conceive of the above as offering a way to show someone how to *work with* the notion.[25] Given the above framework, we can generate biconditionals that take you from statements about metaphysical indeterminacy to statements in a more familiar area (modality) and back again. We can thus use the familiar framework of modality and possible worlds as an aid to theorizing about the somewhat more mysterious metaphysical indeterminacy.[26] This, we hope, enables us to offer a kind of operational understanding of metaphysical indeterminacy without the need to reduce indeterminacy to anything more basic. Just how illuminating the minimal characterization given so far will be must be measured by looking at the extent to which it can diagnose and remove confusion in application. Elsewhere similar ideas have been used to resolve puzzles about m-indeterminate survival, m-indeterminate parthood, m-indeterminate identity in general, m-indeterminacy in facts about the future (i.e. the

[24] Their view is quite subtle: they think that both revisionary and non-revisionary characterizations of 'truth' are possible, but useful for different purposes. So fully specified, they would claim that there is *a* disambiguation of key notions such as 'truth' and 'consequence' on which no revision of classical logic and semantics is enforced.

[25] As a consequence, we're fairly neutral, at least at this stage, about how seriously the metaphysical commitments of the framework need to be taken. You could be instrumentalist about the whole model, just thinking of it as a useful way of cashing out the bare-bones idea of unsettledness. Or you could take on board the abstract representations (or at least those in the space of precisifications) but think that they don't have any bearing on your broader theory of modality. Or you could dovetail theories of indeterminacy and modality, with determinacy just understood as a restricted form of necessity, etc. The choices here hang largely on considerations from elsewhere, particularly on what you think the metaphysics of modality ought to be.

[26] Compare those views on modality where translation of sentences involving modal operators to sentences involving first-order quantification over worlds aren't taken as giving reductions of modality. They are, rather, just theoretical aids to understanding modality. Our project is very much the same; we give translations (from sentences involving determinacy to sentences about worlds in the space of precisifications) not as a reduction, but as a theoretical aid.

‘open future’), and so on.[27] For the distinctions that need to be drawn in these cases, the basic ideas sketched above suffice.

[27] See Williams (2008b), Barnes (2009) Barnes & Cameron (2009), Barnes & Williams (2009).

Part III
The logic of metaphysical indeterminacy

The outline just given is enough for us to have the tools to go on and do some first-order theorizing about this or that putative example of metaphysical indeterminacy. But it would also be nice to do some more theoretical work. For example, when comparing a theory of some putatively indeterminate phenomenon to some rival account, it would be nice to have a systematic way of addressing the question of whether the theory is *consistent*. And when philosophers start to claim that unattractive results *follow from* what we say, we would like to be able to engage in critical evaluation. For both purposes, we would like a *logic* for our indeterminate language. And the best way to do that is to write down its model theory.

The ersatz-worlds framework is highly suggestive in this regard. It suggests that we might think of 'determinacy' as analogous to 'necessity', with the space of ontic precisifications playing the role of the space of possible worlds.

But being suggestive in this regard is one thing; working through the details is quite another. Good housekeeping demands that we show explicitly how the motivating ideas play out.

We now face a choice point: what kind of object-language should we develop a logic and semantics for? There are several natural choices:

1. A language suited to making claims about indeterminate subject-matters. For example, if one thinks that the whether one object is part of another can be metaphysically indeterminate, we might want to know about the logic and semantics of mereology in this setting.

2. A language suited to (1) above, but also containing the resources to make claims about indeterminacy itself: to state explicitly that it is indeterminate whether x is part of y.

3. A language suited to (1) and (2), but which is richer in expressive resources: in which we can express claims about it being *possible* for it to be indeterminate whether x is part of y, for example.

4. ...

To discharge (1), we need only consider standard first-order languages, and give a logic and semantics for these compatible with indeterminacy. To discharge (2), on the other hand, we need to include a language which itself contains an indeterminacy operator. To discharge (3), we need to include in addition standard modal operators. Clearly this list is open-ended: we might wish to add counterfactual conditionals, temporal operators, and so on and so forth.

Each step can be significant. We can illustrate this by analogy with a well-known semantic framework for indeterminacy: standard supervaluationism. Elaborated to suit *semantic* accounts of indeterminacy, supervaluationism says that there are multiple *classical interpretations* of our language that equally fit the meaning-fixing facts. Call such interpretations *sharpenings* of the language. For the 'standard' supervaluationist, a sentence is true iff it is true on all sharpenings.

Now, standard supervaluationism is not fully classical. It allows for meaningful sentences for which bivalence fails, for example: sentences which are true on one sharpening, and false on another. But as regards the *logic* of an object-language (what is a consequence of what), folklore has it that it is *fully classical* with respect to a language of type (1).[28] Not only are

[28] The folklore has recently been challenged by Graff (2003).

all classical logical truths still logical truths, but also classical metarules such as reductio ad absurdum and reasoning by cases are preserved. But (again according to the folklore[29]) this situation does not fully transfer to languages of type (2). Once we have a determinacy operator in the language, then we have the following pair:

- $p \wedge \neg Dp \models \bot$
- $\top \not\models \neg(p \wedge \neg Dp)$

For $p \wedge \neg Dp$ to be true on a supervaluational model, all the sharpenings would have to make the sentence true. But this cannot happen. Nevertheless, the negation of the sentence is never true on any model: it is at best neither true nor false. The pair together constitutes a counterexample to the metarule of reductio.

Standard supervaluationists do not usually explicitly address the question of how their indeterminacy operator plays with modality, and so with languages of type (3). But there is reason for thinking that the departures from classicism may be even more dramatic when we consider such contexts. Suppose a standard supervaluationist wished to preserve the following very natural thought: that *inconsistencies* are *impossible*. That is: whenever $p \models \bot$, we can conclude that $\neg\Diamond p$.[30] In the language of type (2), the supervaluationist is already committed to true disjunctions where each disjunct is an inconsistency. Given this, the inconsistency to impossibility principle, and the factivity of necessity, it will follow that we can have possible disjuncts, where each disjunct is impossible:

$$\Diamond(A \vee B) \wedge \neg\Diamond A \wedge \neg\Diamond B$$

But this is inconsistent in classical modal logic, and so this would mean that the supervaluationist is committed to *asserting* something that is classically inconsistent—something that is does not occur when we limit our attention to expressively poorer languages.[31]

So as we move from languages of type (1) to (3), each stage increases the revisionism: from no revisionism (type 1 languages), to revisionism over classical metarules (type 2 languages), and finally to asserting things that are classically inconsistent (type 3 languages).

To be clear: the point is not that the standard supervaluationist *must* elaborate their theory in this way.[32] It is that nothing in what they say about (1) and (2) debars them from this elaboration. Although the particular issue above won't arise for our theory, we need to recognize that to defend the non-revisionism of a framework, one needs to make sure that the non-revisionism is robust under such additions.

This challenge is open-ended: in principle modal, temporal, counterfactual operators etc could be the locus of novel revisionism. But since our guiding picture has indeterminacy itself presented as a kind of modality, it is natural to worry about interactions between determinacy

[29]For critical discussion, see Williams (2008c).

[30]One articulation of this thought would be to endorse the following rule (a modal weakening of reductio):

$$q \models \bot \Longrightarrow \bot \models \neg\Diamond q$$

[31]The specific case that leads to the above is the following. Take some p which is indeterminate. So $\neg Dp \wedge \neg D\neg p$ holds. Because of the non-revisionism over classical propositional tautologies, $p \vee \neg p$ holds. It follows that $(p \wedge \neg Dp) \vee (\neg p \wedge \neg D\neg p)$ will hold. Hence (assuming that truths are possible) $\Diamond[(p \wedge \neg Dp) \vee (\neg p \wedge \neg D\neg p)]$ holds. But each disjunct is a supervaluational contradiction, and hence by the above metarule, $\neg\Diamond[(p \wedge \neg Dp)]$ and also $\neg\Diamond[\neg p \wedge \neg D\neg p)]$. Conjoining these three, we have an instance of the required form.

[32]A natural alternative would be for them to refrain from endorsing the inconsistency-to-impossibility move, and instead endorse the rule: $q \models \bot \Longrightarrow \Diamond q \models \bot$. The question of which option to choose is beyond the scope of this essay.

and metaphysical modality. So we should *at least* develop a logic and model theory for a language rich enough to address such concerns.

We proceed in stages, and we shall see that there are various choice-points along the way. First, we look at a semantics adequate for a (modal) language with an indeterminate subject-matter, but which itself does not 'invoke' indeterminacy. It turns out that classical semantics can be imported wholesale to cover this case. Things start to get interesting when we add 'determinately' into the object-language we are studying. We start by tentatively extending the *D*-free language model to allow us object-language expression of claims that this or that is indeterminate. From a formal perspective, this initial suggestion seems a little unnatural. We go on to provide a much smoother and more general model, allowing full interaction between modality and indeterminacy (a language of type 3).

In the final section, therefore, we set out a key philosophical puzzle in the vicinity—that of higher-order vagueness—and identify two responses which motivate the different formal treatments we have developed.

4 Languages with an indeterminate subject-matter

What is to be done?

We are to proceed in stages. Our starting point is the familiar language of first order quantified modal logic. We shall develop a recursive characterization of truth at model, and use this to characterize consequence. All the technology is completely standard (and so it will serve as a good base-line for subsequent introduction of non-standard machinery). Nevertheless, it is something that the theorist of metaphysical indeterminacy, guided by the picture of Part II, can accept wholesale. We shall finish by explaining how indeterminacy fits in—specifically in the characterization of truth *simpliciter* (rather than truth *at* a world, variable assignment, or model).

A word about the ambitions of the project before we dive in. In this, and all the semantic theories that follow, even one completely sceptical of the coherence of the notion of m-indeterminacy will be able to follow much of the discussion. The notion of logical consequence, and of truth-on-a-model, never appeal to indeterminacy as such. The situation is comparable to that facing Quinean sceptics about modality confronting 'pure semantics' for modal logic. Even given their scepticism, the notion of truth on a model, and consequence in the specified language, are perfectly clear. They need not blanch until the topic turns to 'applied' semantics, where we need to look at what is true according to one very special model: featuring worlds that are genuinely the *possible* worlds. Likewise, in the present setting sceptics about indeterminacy will only run into trouble when faced with the characterization of the intended model, for it is there that we make essential appeal to (in)determinacy itself.

We are happy to offer sceptics as much as we can (including consequence and truth-on-a-model), but in the end, we have to point them back to the opening sections of the paper, where we argued that complaints of *unintelligibility* against our notion are ill-founded. Our intended audience at this point are those who are open to the prospect of indeterminacy of this kind, prepared to grant us *pro tem* that the notion is in good order, and who wish to see how it relates to truth and consequence.

The semantics

We start with a very familiar setting. Take as our target, initially, the language L of first order quantified modal logic. A model for this language is of the form $m = \langle W_m, O_m, a_m, \|_m\rangle$ where W_m is a set of worlds, $a_m \in W_m$, O_m a set of objects and $\|_m$ is an interpretation function defined on L, mapping each basic name and predicate of L to a function from worlds to objects or sets of objects respectively. A *simple m-variable assignment* is a mapping from variables in L to elements of O_m. Let the *denotation* of a singular term (name or variable) at a world w and variable assignment v be either: (i) the value at w of the intension that the interpretation function assigns to the name; or (ii) the object that the variable assignment gives to the variable.[33] Say:

1. Fn is true at w, v, m iff the denotation of n at w, v, m is a member of $|F|_m(w)$.

2. $\neg\phi$ is true at w, v, m iff ϕ is not true at w, v, m

3. $\phi \wedge \psi$ is true at w, v, m iff ϕ and ψ are each true at w, v, m.

4. $\exists x\phi$ is true at w, v, m iff ϕ is true at w, v', m, where v' is an m-variable assignment differing from v if at all only in what it assigns to x.

5. $\Diamond\phi$ is true at w, v, m iff ϕ is true at w', v, m, for some w' in W_m.[34]

The notion of 'true at world w, variable assignment v, and model m' is therefore defined for sentences of L. We then set:

- For $\phi \in L$, ϕ is true in m iff ϕ is true at a_m, v, m for all m-variable assignments v.

Given this, we can define logical consequence in L: we say that ϕ follows from Γ iff $\forall m$ if all of Γ are true-on-m, then ϕ is true-on-m too.

Notice that all this is perfectly orthodox. According to our guiding thought, metaphysically indeterminacy is manifest in indeterminacy concerning which world is actual. And since the orthodox definitions of truth-on-m and logical consequence need never appeal to the idea of the actual world (or the intended model) the presence of indeterminacy changes nothing.

This changes when we come to characterize the notion of truth simpliciter. To anticipate: a sentence will be true iff it is true on the intended model. And so, we need to start characterize which among the models is the intended one.

We can make the standard first moves: the intended model for the language will have W be the set of ersatz possible worlds, and O the set of all objects whatsoever.[35] All that remains, then, is to specify the intended interpretation of the language, $||$, and the designated actual world element a.

Two kinds of indefiniteness—semantic and metaphysical indeterminacy—surface exactly when we try to address these aspects of the intended model. The guiding idea of semantic indecision accounts of indefiniteness is that the sorts of factors that fix the meanings of our words—the intended interpretation—leave unsettled which one is designated. The guiding idea here is that there is no fact of the matter which one of the ersatz worlds of the intended model is in fact the actualized one.

[33]For simplicity, the semantics below will deal only with monadic predicates, but we can allow polyadic predicates in the obvious way: by letting the denotation of an n-adic predicate at w, v, m be an n-tuple of objects drawn from O_m. Mutatis mutandis for the other semantics to follow.

[34]This language is appropriate for S5 possibility. In the standard way, we could enhance our model with an accessibility relation on the worlds and so define modal operators of varying modalities. The definition of consequence has to be altered accordingly. See, e.g., Priest (2001) for details.

[35]We're ignoring for now set-theoretic paradoxes. See (Rayo & Uzquiano, 1999) for a way of doing semantics allowing for an unrestricted domain. We are also setting the model up using a single-domained setting, where the single domain contains only objects that actually exist. There are of course alternatives: e.g. a single-domained model where the domain contains mere possibilia (thus giving rise to metaphysical puzzles about the ontology of the intended model).

Our concern here is not with semantic indecision and its semantic representation, nor do we wish to engage with the variety of theories about what sorts of naturalistic facts fix which interpretation is intended (if indeed, any do). So we shall ignore such issues, and simply assume that the intended interpretation has been settled.[36] That leaves the specification of the actual world element of the intended model as the only remaining task.

Our guiding conception of metaphysical indeterminacy has it that one among the possible worlds 'gets matters right', but it was indeterminate which world that was. To this point we have specified three of the four elements of the intended model: the space of worlds, domain and interpretation. So the candidates to be the intended model are the various models containing those three elements together with some specified ersatz world. Some of these models will designate as actual worlds that are determinately nonactual. These models will be determinately unintended. Some models will designate worlds that are neither determinately actual nor determinately non-actual (what we called the ontic precisifications). These models will be neither determinately intended nor determinately unintended.

We can still speak of 'the intended model'. For according to our guiding conception, there's exactly one world that depicts reality correctly, and adding this world to our space, domain and interpretation will give a unique, intended model. But as it is indeterminate which world that is, it will be indeterminate which model is the intended one.

It therefore makes perfect sense to characterize truth *simpliciter* in the following, standard, way:

- ϕ is True iff ϕ is true on m for m the intended model

The semantics is fully bivalent; we can prove that $\mathrm{True}(\phi) \vee \mathrm{True}(\neg\phi)$ holds for any closed sentence of L. However, notice that due to indeterminacy as to which is the intended model, the truth-values assigned to sentences will be indeterminate. Suppose 'a' picks out o, and 'F' picks out the property P; that the ontic precisifications are w_1 and w_2 and w_1 represents o having P and w_2 represents the opposite. Then 'Fa' will be true at the model containing w_1, and false at the model containing w_2. Either way, 'Fa' is either true or false. But given the indeterminacy in which model is intended, it will be indeterminate whether it is true, and indeterminate whether it is false. This is a metalinguistic analogue of our motivating thought: that a disjunction can be determinate without either disjunct being so.

Notice that the machinery deployed and the definitions of truth and consequence are entirely standard. No appeal has been made *in the semantic machinery* of indeterminacy: it only features when we work out the *implications* of the definitions offered when put together with independent claims about how indeterminacy is manifested in designating one ersatz world as actualized.

5 Adding determinacy to the object-language

What is to be done?

The semantics just sketched shows that a language expressing metaphysically indeterminate states of affairs can function in a fully classical way. However, one limitation of the above approach should be clear. The language we're working with doesn't contain the resources to express its own indeterminacy: all talk of determinacy and indeterminacy occurs in the *metalanguage*. It is still philosophically significant. If accepted, it assures us, for example, that a

[36]One might wonder if we can really separate off these questions—whether ontic unsettledness in the meaning-fixing facts might not lead to semantic unsettledness. We don't want to prejudge this issue, but more progress will be made if we set it aside *pro tem*.

philosophical theory or theory of applied science that can be formulated in a non-indeterminacy-invoking way need undergo no logical or semantic revisions, upon noting that its *subject-matter* is itself metaphysically indeterminate on some occasions.

But as flagged earlier, we do want to consider theories that not only have an indeterminate subject-matter, but which explicitly deploy notions like determinacy. If *some* theory one accepts has an indeterminate subject-matter, one's *overall* philosophical theory better be of this sort. If the subject-matter of mereology or applied science is indeterminate, we want to be able to *say* that this is so, and we want to understand how the language in which we say this functions. The task of this section and the next is to sketch how this works out.

We shall presently set out a way of incorporating indeterminacy into the object-language that will allow all sorts of standard embeddings—truth-functional compoundings, quantification into the scope of the operators, modal and counterfactual embeddings, etc. But we shall build up to this slowly, because there are intermediate options of philosophical interest.

The restricted scope semantics

We begin by embedding a new element into our models. A model m will now be of the form: $\langle W_m, A_m, O_m, a_m, \|_m\rangle$. Here, W is the set of worlds, $A \subseteq W$ represents those worlds which are according to the model not determinately unactualized, $a \in A$ represents the actual world, O the domain and $\|$ the interpretation function.

Let the language L' be constructed from L by adding a new operator, D. We let a formula $D\phi$ be well formed in L', iff ϕ is well-formed in L. L' allows well-formed object-language statements whose initial element is a determinacy operator, but it doesn't allow embeddings of such sentences—not even conjunctions and negations.

What do we do to characterize the properties of L'? We already have a notion of true-in-m and Truth simpliciter for L. We can now extend this notion to the whole of L'. Suppose $\phi \in L' - L$, i.e. $\phi = D\psi$ for some $\psi \in L$. Say $D\psi$ is true-at-m, where $m = \langle W, A, O, a, \|\rangle$, iff ψ is true-at-m' for all m' where $m' = \langle W, A, O, w, \|\rangle$, for some $w \in A$. Having a notion of truth-on-a-model for such sentences, we can characterize notions of consequence and truth *simpliciter* just as before, as truth-preservation in all models, and as truth at the *intended* model respectively (in the latter case we anticipate that it may be indeterminate which model this is). The one extra wrinkle is that we are obliged to specify one extra element to go in the intended model: the set representing worlds that 'are not determinately unactualized' (the 'ontic precisifications'). So long as we think that there's a fact of the matter, of each world, whether or not it's determinately unactualized, then fixing on a unique such set in this way should raise no new problems, and the only locus of indeterminacy in which model is the intended one, as before, will concern which world is actualized.[37]

The essential thought behind the extension to L' is to insert some previously metalinguistic ways of talking into the object language.[38] So the notion of 'determinate truth' now gets expressed within the object language with the operator D. But it's natural to want more. Let L'' be the smallest language that contains L' and is closed under truth-functional connectives. The extension to L'' allows certain obvious combinations of sentences to be well-formed: for example, the sentence 'it is determinate that Harry is bald, but indeterminate whether Gary is

[37]If there is 'higher order indeterminacy'—indeterminacy over which worlds are determinately unactualized—then we can get indeterminacy in the specification of the model at this point, as well as in the specification of the actualized world. But for now, we're setting these complications aside.

[38]Compare Quine's 'second grade of modal involvement' (Quine, 1953). The previous section corresponds to the first grade, and the subsequent section corresponds to Quine's third grade of involvement, in which quantificational and other embeddings are allowed.

bald'.

The definition of truth on a model will proceed rather disjunctively. It is exactly as before for sentences of L. For sentences of L' we extend the notion of *truth-on-m*:

- $D\phi$ is true in m iff ϕ is true at v, w, m for all $w \in P_m$ and all m-variable assignments.

And again for sentences of $L'' - L'$:

- If at least one of $\phi, \psi \in L'' - L$, then $\phi \wedge \psi$ is true-on-m iff ϕ and ψ are true-on-m.
- If $\phi \in L'' - L$, then $\neg\phi$ is true-on-m iff ϕ is not true-on-m.

We have in this way a definition of true-on-m for D-involving sentences, which allows us to stick with our original, standard definition of truth and consequence for L''. In particular, the availability of a definition of *consequence* means that we have a characterization of the logic of a indeterminacy-invoking language (and one that is fully classical and whose semantics are bivalent). But this is a language of limited resources: quantificational and modal operators can't take wide-scope over D, and D itself cannot be iterated.

Let us step back to see what we've been doing. In effect, the truth-functional connectives are defined twice over. Once, for sentences of L, in terms of their interaction with truth at w,v,m. Second, we define them again for those truth-functional combinations of sentences which are D-involving, and this time we define them in terms of truth-on-a-model.

Looked at from a technical point of view, this seems unnatural. Why not just treat D as syntactically analogous to the modal operator $\Box$, and allow full interaction with quantifiers, modal operators and iterations of D itself? Semantically, this could be effected by giving conditions for when $D\phi$ is true at w, v, m for any formula ϕ in the extended language.

We do exactly this in the next section. Some new technical questions arise at this point to do with the interaction of D and modality, with which we haven't had to deal here, because we constructed L'' to make such combinations ill-formed. After we have both proposals on the table, we shall look at puzzles of higher-order indeterminacy, and explore some philosophical motivations for looking kindly on the approach outlined above.

6 Haloes and worlds

What is to be done?

The language to be studied in this section is one in which modality and indeterminacy (and quantifiers) are allowed to freely interact—including iterated determinacy operators and modal and quantificational embeddings of determinacy. This will be the richest language we will consider, and we shall again see that classicism can be maintained.

The unrestricted semantics

We shall now set out a semantics for a language that allows for full embeddings of expressions within others.

Let L^* be the extension to L including the D operator, and allowing the obvious embeddings with quantifiers. One might try simply to define D as a modal operator over the space of actual ontic precisifications. But something more subtle is needed—for $\Box$ and D must interact if we are to get sensible results. For example, it is possible for you to be a contender—and it is possible for you to *determinately* be a contender. To get these coming out true, we need to have some machinery that means that world-shifting modal operators also shift which worlds D ranges over. So the semantics for the modal operators and the semantics for D must interact—we can't just plug them in separately.[39]

What we propose to do is to alter the indices to which the truth-definitions are relativized prior to the definition of the key notion of truth-in-a-model. Up till now the 'ersatz worlds' of the model have been assumed to be classical complete representations of the way the world might have been. And we have been treating indeterminacy as a modal notion: worlds don't explicitly say what is indeterminate or determinate at them (just as they don't say what is necessary or contingent at them). In the case of unmodalized indeterminacy claims, we could deal with this simply by appealing, not only to the One True actualized world, but also to a 'halo' of worlds that surround it: those that are not determinately unactualized.

Now let us extend this thought to include modal facts. Absent indeterminacy, a single world gets picked out that represents reality correctly. To account for modal truths, this single actualized world is accompanied by a collection of ersatz worlds—and we focus on that collection that represents the facts about possibility and necessity correctly.

The motivating picture we have here says that all this can be preserved. There is the One True actualized world—it is just indeterminate which world this is. And there is the One True set of possible worlds—it is just indeterminate which set of worlds it is. There are not just ontic precisifications of actuality (those worlds that don't determinately misrepresent it), there are ontic precisifications of the whole of modal space (those worlds that don't determinately misrepresent the possibilities). Consider Fred the foetus. If it is actually indeterminate whether Fred is a person, then there are two ontic precisifications of actuality—one where he is a person, one where he is not a person. Similarly if it is *possibly* indeterminate whether Fred is a person, then these will be two ontic precisifications of that possibility.

So every possibility stands in two interesting relations to worlds. There are those worlds it thinks are possible, and worlds that it thinks don't determinately misrepresent the ways things are. The two relations are relatively independent. Consider an indeterminate non-contingency: for example, perhaps it is indeterminate whether the continuum hypothesis holds. The actual world says that either the continuum hypothesis holds or it doesn't: no matter which option is actualized, the other isn't determinately unactualized. But whichever option is actualized, it better be that the other is *impossible* (on pain of treating the continuum hypothesis as contingent).

If we are to allow for this, 'precisifications' must operate on modal space as a whole. Given two possibilities that differ in various ways, but which both have the continuum hypothesis being indeterminate, *if* the One True representation of the first possibility makes the continuum hypothesis true, the One True representation of the second possibility makes it true too. We precisify possibilities simultaneously, not individually and independently.

Our representation of possibilities will consist of the following pair: a set of possible 'haloes of indeterminacy', representing the worlds that are not determinately incorrect, by the lights of each possibility; and a 'modal selection function' that says of *every* such possible halo of indeterminacy, which ersatz representation is the One True representation of the possibility corresponding to the halo. This captures all the information described above. Just as we can think of the actualized world as 'selecting' the One True representation from the actual halo of indeterminacy, the intended selection function selects the One True representation for *each* possible halo of indeterminacy.

So much for the motivating picture. Now to capture this in our models. In general, a

[39] In general, in a multi-modal system, we have to watch out for interactions between the modalities when writing up the semantics. See for example Thomason (1984) on combinations of tense and modality for one well-studied instance.

model will have the form $\langle H_m, S_m, O_m, A_m, s_m, \|_m\rangle$. Here's how to think of what such a model represents. H_m is our stock of possible haloes of indeterminacy of the model (it contains a multitude of sets of ersatz worlds). S_m is model's stock of selection functions, precisifying the haloes so they represent complete possibilities. O_m is the domain of the model, and $\|_m$ the model's interpretation function. A_m is the model's actual halo of indeterminacy; and s_m is the model's actual selection function—something that precisifies each possible halo in the One True way (by the lights of the model). The actualized world of the model is fixed by these last two elements—as the One True representative of the actual halo of indeterminacy. So we don't need to list it separately. Two plausible interrelations between H_m and S_m are insisted upon: (i) that every $\sigma \in S_m$ is such that $\sigma(U) \in U$; and (ii) for every $U \in H_m$ and each $w \in U$, there is some $\sigma \in S_m$ such that $\sigma(U) = w$. Without this, the connection between $U \in H_m$ as 'haloes of indeterminacy' and $\sigma \in S_m$ as 'ways of precisifying the indeterminacy' would be lost.

Given this setup, we can give clauses for parameterized truth in L^*, relative to precisified possibilities i.e. objects of the form $\langle\sigma, U\rangle$ for $\sigma \in S_m$ and $U \in H_m$. The clauses then run:

1. Fn is true at $\langle\sigma, U\rangle, v, m$ iff there is a $w = \sigma(U)$ and the denotation of n at $\langle\sigma, U\rangle, v, m$ is a member of $|F|_m(\sigma(U))$.
2. $\neg\phi$ is true at $\langle\sigma, U\rangle, v, m$ iff ϕ is not true at $\langle\sigma, U\rangle, v, m$
3. $\phi \wedge \psi$ is true at $\langle\sigma, U\rangle, v, m$ iff ϕ and ψ are each true at $\langle\sigma, U\rangle, v, m$.
4. $\exists x\phi$ is true at $\langle\sigma, U\rangle, v, m$ iff ϕ is true at $\langle\sigma, U\rangle, v', m$, where v' is a simple m-variable assignment differing from v, if at all, only in what it assigns to x.
5. $\Diamond\phi$ is true at $\langle\sigma, U\rangle, v, m$ iff ϕ is true at $\langle\sigma, U'\rangle, v, m$, for some $U' \in H_m$.
6. $D\phi$ is true at $\langle\sigma, U\rangle, v, m$ iff ϕ is true at $\langle\sigma', U\rangle, v, m$ for all $\sigma' \in S_m$

We can then define truth-on-a-model, truth simpliciter and consequence (or 'following from' or '$\models$') in the obvious ways:

- For $\phi \in L$, ϕ is true-on-m iff ϕ is true at $\langle s_m, A_m\rangle, v, m$ for all m-variable assignments v.
- ϕ follows from Γ iff $\forall m$ if all of Γ are true-on-m, then ϕ is true-on-m too.
- ϕ is true *simpliciter* if ϕ is true on the intended model m.

What is the intended model? Well, mutatis mutandis, the story goes through much as before. The main difference is that the indeterminacy now surfaces not over what actual world element the intended model contains; but over what selection function it contains. In the representation above, this is all packed into the element s_m—a model m will be intended only if it contains the s_m that, for each possibility with halo U, $s(U)$ picks out that world that represents the possibility correctly. As before, we urge that one choice of s is correct—but it will be indeterminate which it is. And since the intended model needs to contain the One True s, it will accordingly be indeterminate which model is intended.[40]

[40]In the framework just given, the logic of the operator D will be S5. In effect, the 'halo' carried around with a world corresponds to accessibility relation on worlds that is an equivalence relation: dividing the worlds outside the halo from the worlds inside. But many friends of metaphysical indeterminacy will not want an S5 logic. Worried by the puzzles of higher-order indeterminacy we will shortly discuss, they would like to formulate and endorse the thought that in certain cases it is indeterminate whether or not something is indeterminate.

To generalize our framework, we need to generalize our halos. Rather than fixing S5 accessibility relations, We finish by stating some illustrative consequences of the logic. Suppose we were interested in whether the law of excluded middle held for this setting. That is to ask whether, for arbitrary ϕ, $\models \phi \vee \neg\phi$; i.e. whether for every model m, $\phi \vee \neg\phi$ is true-on-m. By the characterization of disjunction, such a sentence is true-at-$\langle s_m, A_m\rangle, v, m$ iff either ϕ is true at $\langle s_m, A_m\rangle, v, m$, or $\neg\phi$ is true at $\langle s_m, A_m\rangle, v, m$. But by the characterization of negation, that is true iff ϕ is true at the relevant string, or it is not true at the relevant string. And we know that to be the case (given that the metalanguage is classical).

What of the classical metarules that famously fail for (some) supervaluational logics? Well, consider reductio, for example. It will indeed follow from the fact that $\Gamma, \phi \models$, that $\Gamma \models \neg\phi$. For the first tells us that no model m of the language makes all of $\Gamma \cup \{\phi\}$ true-on-m. That means that any model that makes all of Γ true, cannot make ϕ true-on-m. But if ϕ is not true-on-m, then it must not be true at $\langle s_m, A_m\rangle, v, m$. But $\neg\phi$ is then true-on-$\langle s_m, A_m\rangle, v, m$, by the clause for negation above, and hence it is true-on-m. So any model that makes all of Γ true (on a model) makes $\neg\phi$ true (on that model). And that is just what it is for $\Gamma \models \neg\phi$ to hold.[41]

Turning to the modal aspects of the logic, we can show that the K, T and $S5$ axioms for $\Box$ (where $\Box\phi := \neg\Diamond\neg\phi$; and $\phi \rightarrow \psi := \neg\phi \vee \psi$) hold good in our language, and the necessitation metarule also holds:[42]

- $\models \Box\phi \rightarrow \phi$
- $\models \Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$
- $\models \Diamond\phi \rightarrow \Box\Diamond\phi$
- $\models \phi \Longrightarrow \models \Box\phi$

We also have the analogues of each of these for the D operator:

- $\models D\phi \rightarrow \phi$
- $\models D(\phi \rightarrow \psi) \rightarrow (D\phi \rightarrow D\psi)$
- $\models \neg D\neg\phi \rightarrow D\neg D\neg\phi$
- $\models \phi \Longrightarrow \models D\phi$

they need to be able to fix arbitrary accessibility relations. The natural thought is simply to replace our haloed worlds (pairs of a world w and a set of worlds containing w) with pairs of worlds and accessibility relations across non-haloed worlds. Working through the framework in the natural way, the only substantive change we need is the following:

- $D\phi$ is true at $\langle \sigma, R\rangle, v, m$ iff ϕ is true at $\langle \sigma', R\rangle, v, m$ for all $\sigma' \in S$ such that $\sigma(R)$ stands in the relation R to $\sigma'(R)$.

However, in this setting, it is perhaps more plausible that distinct possibilities will correspond to a single accessibility relation, and accordingly we may well want to move to the more general representation suggested in an earlier footnote.

[41] In supervaluational logics (at least on a standard way of presenting them), the reasoning breaks down because a sentence can fail to be true-on-m without having $\neg\phi$ true-on-m.

[42] Most of these are routine to check. The necessitation rule is perhaps the most involved (I will suppress relativization to variables since they play no role in the following proof). For this, suppose that $\models \phi$. That means that ϕ is true on all models. So, for every model m, ϕ must be true at $\langle s_m, A_m\rangle$. We show that $\Box\phi$ too is true-on-m. What this requires is that ϕ is true not just at the point $\langle s_m, A_m\rangle$ in the model, but at all $\langle s_m, U\rangle$ for $U \in H_m$. But now consider the model m' which differs from m just in which halo is picked out as actual, and in particular $A_{m'} = U$. But since ϕ is true on *all* models, it is true in particular on m'. And hence ϕ is true at $\langle s_{m'}, A_{m'}\rangle$ in the model m', which by construction of m' is just to say that it is true at $\langle s_m, U\rangle$.

Given these results, it seems plausible that classical modal logic S5 as a whole extends to our systems, and also that the logic of determinacy will map onto familiar modal systems.[43]

But there remains the question about what kind of *combined* logic we get. Now the question of *revisionism* or otherwise doesn't really have application here, since we don't have a 'classical' treatment of the combinations to be faithful too. But non-trivial claims can be read off. In particular, the following principle *fails*:[44]

$$\Box\phi \models D\phi$$

This is as it should be, if we allow for indeterminate non-contingencies (such as, perhaps, the continuum hypothesis). For contraposing the rule above, we'd have $\neg Dp \models \neg\Box p$, and $\neg D\neg p \models \neg\Box\neg p$, and hence from the indeterminacy of p we would be able to infer that p was neither necessary nor impossible, i.e. contingent.

7 Higher order indeterminacy

Recap

Let us pause to recap. In part I, we articulated our favoured account of what metaphysical indeterminacy is: the primitivist conception. That conception leaves open, we think, many other questions we would want a theory of indeterminacy to answer. In particular, it leaves open what the logic is (and so, which substantial philosophical theories are consistent and which are inconsistent). In part II, we explained a guiding conception of indeterminacy, one that is broadly classical. And now, in part III, we have worked in detail through logical frameworks. We saw that classical logic and semantics can be carried over wholesale for languages that don't involve the indeterminacy-operator. But when indeterminacy is in the language, we face some choice points.

We developed two accounts of language that can express indeterminacy as well as modality. The restricted semantics underpins a language in which quantifiers and modals cannot scope over determinacy operators (though truth functional connectives can). The unrestricted semantics, on the other hand, underpins a rich language in which determinacy interacts freely with all the other resources in the language.

The task of the present section is to illustrate how the choice between these two formal settings might emerge from substantive views about the best diagnosis to give of the phenomenon of *higher order indeterminacy*. Overall, the discussion will bring home the point that the primitivist story itself does not *force our hand* when developing these formal frameworks: we need to look to independent motivations for guidance.

Two views of higher-order indeterminacy

In the sort of classical setting in which we work, the law of excluded middle holds even in indeterminate cases. Thus, if Sparky is a borderline part of you, then we have:

(1) Either Sparky is part of you, or Sparky is not a part of you

[43] Importantly, $p \wedge \neg Dp \models \perp$ does not hold in our system. Consider a model where the world picked out by $s_m(A_m)$ makes p true, but there is a $w \in A_m$ that makes it false, and a $\sigma \in S_m$ such that $\sigma(A_m) = w$. Then $p \wedge \neg Dp$ will be true on m. If this conjunction were inconsistent, we'd get breakdowns in classical metarules just as in e.g. global supervaluational treatments of indeterminacy.

[44] Take a model at which ϕ is true at $s_m(U)$ for every $U \in H_m$; but where there's a $w \in A_m$ at which ψ fails, and a $\sigma \in S$ such that $\sigma(A_m) = w$.

Someone might worry that this conflicts with the idea that there is no fact of the matter concerning whether Sparky is part-related to you. The canonical response is to say that this confuses the claim literally expressed by the disjunction with the following:

> (2) Either Determinately, Sparky is part of you, or Determinately, Sparky is not a part of you

This latter disjunction expresses the claim that the facts of the matter concerning Sparky's relation to you are non-vague. But we've seen no reason to think that it follows from the initial disjunction (though we might admit that perhaps the former disjunction suggests or pragmatically implicates the latter).

A sceptic might note that for us, the following holds:

> (3) Either Determinately, Sparky is part of you, or Determinately, Sparky is not a part of you, or it is indeterminate whether Sparky is part of you.

She suggests: even if our classical framework makes room for indeterminate parts, the above result suggests that we replace the original dichotomy with a trichotomy. Although there may be no fact of the matter what Sparky is a part of, nevertheless (she suggests) there will be a fact of the matter about which of the following three statuses Sparky enjoys: being a determinate part of you, a determinate non-part of you, or being an indeterminate part of you.[45]

This worry seems to *repeat* the move found to be fallacious in the first case. Endorsing a disjunction with three disjuncts (that Sparky falls into one of three categories) is not to commit oneself to there being a fact of the matter about which disjunct obtains; no more so than endorsing a disjunction with two disjuncts committed one to this in the original case.

However, if we are to fully parallel the response to the original case, we must (A) admit that there is coherent thought which the objector is tracking—a 'higher order' determinacy question of whether the categorization into determinate-part/determinate-non-part/indeterminate-part is itself determinate; (B) to say that the 'first order' disjunction does not express this thought (or at least, that the objector has not yet made a case that it does express that thought).

If the friend of indeterminacy concedes (A), however, then a certain kind of regressive structure is started. She will say that the higher-order claim that the objector thinks that (3) implies is the following:

> (4) Either, (a) determinately, determinately Sparky is part of you; or (b) determinately, determinately Sparky isn't part of you; or (c) determinately, it is indeterminate whether Sparky is part of you.

But, says the friend of indeterminacy, we haven't seen reason to believe this follows from (3).

Now, of course the objector can point to a disjunction that stands to (4) as (3) stands to (2), which the friend of indeterminacy will endorse. She might (frustratingly!) worry that endorsing that disjunction will commit to some further 'fact of the matter' thesis. So the debate iterates: a new, even more refined confusion will be diagnosed by the friend of higher-order indeterminacy.

We do not think that the opponent lands a glove at any of these stages. The response by the friend of indeterminacy at each stage seems adequate—certainly enough to throw the burden of

[45] Of course, in some cases we might argue that there being facts of the matter at this level is perfectly consistent with the phenomenon being described: the friend of the open future, for example, might think that future-facts are indeterminate, without ever thinking that there must be indeterminacy in which facts are indeterminate. But we're supposing that the case at hand is one where this 'bullet-biting' response is unattractive.

proof back onto the opponent. If the opponent's strategy does not work on the first iteration, it's hard to see why one would think it would work on the more abstract iterations.

However, the dialectic does, apparently, force the friend of indeterminacy into making sense of various highly abstract iterated notions of determinacy: not just the determinate and the indeterminate, but the determinately indeterminately determinately determinate (and so on). And one might wonder whether we really have a grip on such notions. Further, even if the opponent does not succeed in getting the friend of indeterminacy to admit higher-order determinacy at any level, she has still forced her to replace the simple polar options: being a part, or not a part—by an arbitrary number of possible statuses.[46] But refining our categorization in these sophisticated ways may seem antithetical to one guiding idea about indeterminacy: that even the binary division is too rough. And of course, the opponent might choose at some stage to take up the burden of proof: to argue that one has reason to believe in some particular logical relation among determinacy operators that would bridge the gap between the analogues of (2) and (3).

This dialectic only got started because the friend of indeterminacy conceded the point labelled (A), above: that there was a coherent question about the determinacy of categorizing Sparky as a determinate part/indeterminate part/determinate non-part. This concedes part of the opponent's objection (in effect, an array of 'higher order' parthood statuses for Sparky to enjoy). The friend of indeterminacy who concedes (A) resists only the opponent's claims about what sentences express what. It is only to be expected that the debate will turn into a highly technical one about the precise logic of the notion—and the opponent can be expected to have mileage here.

But why should the friend of indeterminacy concede (A) in the first place? The original thought was that Sparky's status was unsettled between two polar options: being a part of you, not being a part of you. The friend of indeterminacy must therefore concede that first-order indeterminacy thoughts are coherent, else she loses her subject-matter. And (2) is her favoured way of articulating that thought. But why think that we must react in parallel ways to the iterated challenge? An opponent claims that (3) expressed a kind of determinacy of status for Sparky. The friend of indeterminacy might suggest that the distinction between (1) and (2) to show that this gloss on what disjunctions commit one to is tendentious, and so the burden of proof is on the opponent to argue that (3) expresses what she says it does. But, we suggest, one option for the friend of indeterminacy at this point is to suggest that there *is no coherent thought* to be expressed—the idea of 'being determinately a part of you' and 'it being indeterminately a part of you' and the rest as *statuses* which objects can have (determinately or not) is simply confused. Recall our basic thought that the indeterminacy of p is not a separate, 'third-category' status between p and not-p. One might worry that making 'indeterminately F' a separate status that things can have, determinately, indeterminately, or otherwise, loses grip on that basic thought. There's not a separate third status for p—rather, there are two ways things could be, p-wise, and it's simply unsettled which in fact obtains. It doesn't seem ad hoc, therefore, for the friend of indeterminacy to hold that iterations of determinately beyond the basic case are simply meaningless. If that is so, then no regressive structure of iterated higher-order vagueness ever gets started.

One should be reminded at this point of the 'restricted scope' semantics for the indeterminacy-invoking language we outlined earlier. By design, determinacy operators do not iterate in that setting; it is a formal articulation of the idea that iterated determinacy just doesn't make sense. Of course, there are costs of adopting such a framework: as emphasized, it is not just that the

[46] Thinking of 'determinately' as a modal operator, these statuses correspond to the notion of the 'modalities' in various modal systems.

D operator does not iterate, but it also does not embed under modal operators, or interact with quantifiers in natural ways. While declaring iterated determinacy claims meaningless may seem ok, it is harder to swallow the idea that we cannot speculate about whether it is necessarily the case that the world is determinate, or whether you, who in fact have your parts determinately, might have had indeterminate parts. But it would be nicest if an opponent could go beyond appeals to intuition, and say something about what the real theoretical costs are for one willing to bite this bullet.

If one wishes to accept that higher-order indeterminacy makes good sense, to accept (A) and fight the opponent of higher-order indeterminacy on logical grounds, then there are clear grounds for preferring the more general way of setting up the semantics where embeddings of the D operator are allowed (obviously, if one wishes to allow for higher-order indeterminacy so conceived, one will wish to develop a generalized version of haloed worlds, allowing a non-S5 logic for D (see previous footnote)). And given the options presently on the table, one might think that the costs of giving up modal and quantificational embeddings are just too much to go any other way—even if one regards the meaningfulness of iterated Ds as an unwelcome side-effect.[47]

8 Conclusion

The paper has had two central tasks: to lay out a distinctive 'primitivist' position on the nature of metaphysical indeterminacy; and to use a connection between metaphysical indeterminacy and a kind of restricted modality to develop a model theory for indeterminate languages. Our discussion of these issues is now complete, and we take it that their combination lays the foundations for a theory of metaphysical indeterminacy.

We began with the question of *what a theory of indeterminacy* should look like. The following list of tasks was offered:

1. Describing what indefiniteness is: its nature or source.
2. Describing how to *reason* with the notion of indefiniteness, or in the presence of indefiniteness: what the logic of indefiniteness is.
3. Describing the cognitive role of indefiniteness: e.g. what impact knowledge that p is indefinite should have on one's opinion as to whether p.
4. Describing paradigm instances of indefiniteness; if multiple kinds of indefiniteness (individuated by nature, logic or cognitive role) are in play, describing which account is applicable to which case.
5. …

[47] How might one who is tempted by the 'meaningless' response above think of a setting that allows iterations of determinacy? One might simply put a syntactic restriction on the language, declaring ill-formed any sentence featuring an iteration of D. But this brute-force solution seems odd given that our language appears to be one in which iterations are syntactically well-formed. Alternatively, one might think that syntactically, iterations are allowed, but that they only have vacuous truth-values. But a better idea is to sharply distinguish between higher-order vagueness and iterated D operators. One might say that D, DD, DDD etc all express the single concept of *determinacy*—so that while $DDp \vee DD\neg p \vee D(\neg Dp \wedge \neg D\neg p))$ is well-formed and true, it does not express the thought that there is a fact of the matter about whether Sparky is determinately part of you.... There just is no such thought—and what the sentence above expresses is exactly the same as the S5-equivalent sentence $Dp \vee D\neg p \vee (\neg Dp \wedge \neg D\neg p)$. Of course, there would then be a question about how to think of the relation between the formalism and ordinary language. We shall not go further into this issue here.

In this essay we have addressed the first two facets of a theory of indeterminacy: its nature and logic. But with this skeleton in place, much remains to be done.

An obvious direction for future research concerns the other items on our list. What impact would thinking that things are indeterminate (in this primitive sense) have on our cognitive lives? What kinds of first-order theories might benefit from a dose of primitive indeterminacy? Getting clear on the former point would be a step towards fleshing out some of the issues about the *concept* of indefiniteness or indeterminacy that we left hanging in section 2. On the latter point, the metaphysics of composition, of the 'openness' of the future, and of indeterminacy within fundamental physics, are salient starting points.

We also emphasized that specifying the logic of indeterminacy, in particular, was an open-ended venture. We have given a treatment of a language with modality and indeterminacy: but what of temporal and counterfactual operators? What of languages containing belief or probability-operators?

With a rich enough language to hand, we are in position to explore in detail the cogency or otherwise of prominent claims in the literature on metaphysical indeterminacy: the status and coherence of *de re* metaphysical indeterminacy; the alleged impossibility of metaphysically indeterminate existence; and the coherence or otherwise of metaphysically indeterminate identity. As an illustrative example, we discuss some puzzles surrounding *de re* indeterminacy an appendix. Otherwise, we leave these matters for another day.[48]

[48]The authors have started to address some of these questions within a similar framework in Barnes (2009), Barnes (forthcoming), Barnes (2005), Williams (2008b), Barnes & Williams (2009). See also Woodward (2010).

Appendix

A Quantifying in

The idea of vagueness in reality is often associated with *de re* vagueness, that is, vagueness expressed by certain quantified statements. For example, Merricks writes:

> One alleged variety of vagueness is metaphysical. There is vagueness of this variety if, for some object and some property, there is no determinate fact of the matter whether that object exemplifies that property. (Merricks, 2001, p.145)

And Williamson:

> . . . if reality is vague, it is vague how things are, so for some way it is vague whether things are that way; thus, for some state of affairs S, it is vague whether S obtains. (Williamson, 1994, §4)

Both authors use *de re* or 'quantified-in' formulations as a litmus test for metaphysical indeterminacy or vagueness.[49] As will be clear from our initial discussion, we do not think that *de re* indeterminacy provides a characterization of metaphysical indeterminacy, any more than *de re* modality provides a characterization of metaphysical modality. Other modalities, and other forms of indeterminacy may equally, we think, allow for *de re* formulations. However, it is particularly natural to use de re formulations in the case of metaphysical indeterminacy, and given that others may disagree with us as to its significance, we should look at how this works.

To even formulate these comparisons we require more resources than we have allowed for in L^*: they require quantification over properties or, more generally, higher-order quantification. Before we can compare the characterizations of metaphysical indeterminacy we offer with that offered by these authors, we need to explain how higher-order quantification is to work in our setting.

The most straightforward proposal is to introduce a range of variables that syntactically occupy predicate position, and to extend variable assignments so that these variables are mapped onto extensions (we shall use upper-case letters such as X for predicate-position variables). Just as with name-position expressions, we then define a general notion of denotation, so that a predicate F will denote (at $\langle\sigma(U)\rangle, v, m$) $|F|_m(\sigma(U))$—i.e. the extension that is the value at the world-component of the haloed world, of the intension assigned to 'F' by the model's interpretation—and a variable X will denote whatever extension the variable assignment v gives it. All else can go through as before.

Williamson (2003) notes (in an analogous setting) that this kind of definition will make exhaustively quantified-out statements trivial. In particular, whether Xx will be true or false at $\langle\sigma, U\rangle, v, m$ depends solely on v: whether the extension v assigns to X contains the object it assigns to x. The result is that $\neg DXx$ will always be false, at every index. For it to be true, Xx would have to be true at one $\langle\sigma, U\rangle, v, m$, and false at some $\langle\sigma, U\rangle', v, m$. But we've just seen that that can't happen, since the variable assignments fix the truth value, and they are the same. Defining ∇p as $\neg Dp \wedge \neg D\neg p$, this means that the following holds:

[49] Williamson emphasizes that he uses sentence-position quantification (expressed by quantification over 'states of affairs'). This allows the test to be applied to such states as *there being exactly three things* which falls outside the scope of Merricks' test. However, we shall concentrate on the name/predicate position forms of quantification, as the central issues we wish to discuss arise already in this case, and the technology is more familiar (in this we follow Williamson's lead).

$$\neg\exists X\exists x\nabla(Xx).$$

But this, it might be alleged, is what formally corresponds to Merricks' characterization as involving some property and some object, such that it is (m-)indeterminate whether the latter instantiates the former. Does this mean that we cannot have *de re* m-indeterminacy in our setting? That would be far too strong. Our setting allows, for example, there to be objects such that it is indeterminate of them whether they satisfy the 'perfectly natural property' *being charged*. The reason that this can happen is that such a claim has the form:

$$\exists x\nabla(Fx)$$

where F is a predicate expressing that natural property: it is not quantified out in the way explored above. It is made-true, for example, if Sparky is charged in one ontic precisification, and not charged at another, and so the extension assigned to F at these worlds differs correspondingly.

What is clear from this, however, is that the fully quantified claim above cannot be regarded as an adequate gloss on Merricks' characterization of metaphysical indeterminacy: that there be some object and some property such that it is indeterminate whether the former instantiates the latter. For we can have a paradigmatic instance of that gloss (some object indeterminately instantiating the natural property *charged*) while the fully quantified higher-order claim does not follow.

The trouble is that the higher-order variables, if handled in the way we have described, are extensional. One would get analogous results if one characterized higher-order quantification in modal contexts in the way we have done above. Some object may then be contingently charged; but there will not exist an object and some X such that it is contingent whether an object is X. By exactly the same considerations as above, we can show that the following always holds:

$$\neg\exists X\exists x \text{ Contingent } (Xx).$$

despite the fact for various predicates F we have the following:

$$\exists x \text{ Contingent}(Fx)$$

Again, one should not conclude from this that *there is no object, and no property, such that it is contingent whether the former instantiates the latter*. Rather, one should conclude that, due to the extensional treatment of the higher-order quantifiers in the fully quantified statement, the formalism does not express the thought in italics.

Can we express that thought by other means? One idea is to retreat to first-order quantification over properties and appeal to an instantiation relation. But in doing so we would leave behind those with nominalistic scruples, and it would not be obvious how to generalize the proposal.[50] There is an obvious alternative: rather than using extensional higher-order quantifiers, we can use intensional ones. Rather than a variable assignment pairing X with an extension, it will now pair it with an intension (a function from worlds to extensions). With this in place the denotation of a variable will depend on what world features in the index, and there will no longer be any a priori block on the truth of the fully quantified claim $\exists X\exists x\nabla(Xx)$.

[50] Williamson makes a persuasive case that we should generalize to allow sentence-level quantification in general if we want to see all indeterminacy in reality connected to *de re* indeterminacy.

Williamson considers this kind of proposal (again, in a slightly different context), and suggests some problems.[51] The fundamental issue that it makes it too easy to get *de re* indeterminacy on this proposal. Suppose for example that in a far off corner of the universe there is some small metaphysical indeterminacy (concerning the exact location of a subatomic particle, for example). One would not have thought that there would be any indeterminacy in which properties *you* instantiate in such a circumstance. However, consider that intension which maps one of the ontic precisifications onto the universal set, and another onto the null set. You are in one extension, and not in the other, and so this intension will be a witness for the claim:

$$\exists X \nabla X(\text{you})$$

The issue here though is a familiar one. In quantifying over all intensions (or 'abundant properties') we quantify even over highly extrinsic properties of individuals. What is witnessing the truth of this quantified statement is something like the property *being such that the far-away particle is in such-and-such a location*. Indeterminacy anywhere will lead to indeterminacy everywhere, if we count such properties.

So this point does not show anything wrong with the intensional higher-order quantification being able to express the thought behind 'there is some object and some property such that it is indeterminate whether the former instantiates the latter'. Rather, it is just highlights an ambiguity in that thought itself: if by 'properties' it means only properties in an abundant, possibly extrinsic sense, then such indeterminacy will be rather undiscriminating. If we want a more discriminating notion, the solution is easy to hand: simply restrict the range of the quantifiers to those one is interested in. So one might suggest the following as a way of expressing the thought that there is some object and some genuine or 'sparse' property which together are a locus of indeterminacy:[52]

$$(\exists x)(\exists X : \text{Natural}(X))\nabla Xx$$

The same choices are faced for quantification at each syntactic category. Extensional quantification may be useful for various purposes, but one shouldn't read into merely extensional quantifying-in deep philosophical significance. Intensional quantification-in has more of a chance of capturing the thoughts associated in the literature with *de re* indeterminacy. But one will need to pay careful attention as to what range of intensions are relevant to the particular claim one wishes to express.

References

Akiba, Ken. 2000. 'Vagueness as a modality'. *Philosophical Quarterly*, **50**, 359–370.

Akiba, Ken. 2004. 'Vagueness in the world'. *Noûs*, **38**, 407–29.

Barnes, Elizabeth. 2005. 'Vagueness in sparseness: a study in property ontology'. *Analysis*, **65**, 315–321.

Barnes, Elizabeth. 2009. 'Indeterminacy, Identity, and Counterparts,'. *Synthese*, **168**.

[51]Some of Williamson's most persuasive objections have to do with extending the proposal to name-position variables, and we will not discuss those here as there is no obvious need to take this extra step. Though see below for a proposal about how to get the same effect with intensional name-position variables.

[52]The same tactic can be used if one wishes to make all quantification intensional. To express the distinctive thought that some *object* is the locus of indeterminacy, one will need to restrict the intensions over which one quantifies to those that track the career of a single object across the space of ontic precisifications.

Barnes, Elizabeth. forthcoming. ‘Ontic vagueness: a guide for the perplexed’. *Noûs*.

Barnes, Elizabeth, & Cameron, Ross P. 2009. ‘The Open Future: Bivalence, Determinism and Ontology’. *Philosophical Quarterly*, **146**(2), 291–309.

Barnes, Elizabeth, & Williams, J. Robert G. 2009. ‘Vague parts and vague identity’. *Pacific Philosophical Quarterly*, **90**(2), 176–187.

Barnett, David. 2009. ‘Is vagueness sui generis?’. *Australasian Journal of Philosophy*, **87**(1), 5–34.

Eklund, Matti. 2008. ‘Deconstructing ontological vagueness’. *Canadian Journal of Philosophy*.

Field, Hartry H. 2000. ‘Indeterminacy, Degree of Belief, and excluded middle’. *Nous*, **34**, 1–30. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001) pp. ??–??.

Field, Hartry H. 2003. ‘No fact of the matter’. *Australasian Journal of Philosophy*, **81**, 457–480.

Fine, Kit. 1975. ‘Vagueness, truth and logic’. *Synthese*, **30**, 265–300. Reprinted with corrections in Keefe and Smith (eds) *Vagueness: A reader* (MIT Press, Cambridge MA: 1997) pp.119-150.

Fine, Kit. 2005. *Tense and Modality*. Oxford: OUP.

Graff, Delia. 2003. ‘Gap Principles, Penumbral Consequence, and Infinitely Higher-Order Vagueness’. *In:* Beall, J. C. (ed), *Liars and Heaps: New essays on paradox*. Oxford: Oxford University Press.

Greenough, Patrick. 2003. ‘Vagueness: A Minimal Theory’. *Mind*, **112**(446).

Heller, Mark. 1996. ‘Against Metaphysical Vagueness’. *Philosophical Perspectives*, **10**, 177–85.

Lewis, David K. 1983. ‘New Work for a Theory of Universals’. *Australasian Journal of Philosophy*, **61**, 343–377. Reprinted in Lewis, *Papers on Metaphysics and Epistemology* (Cambridge University Press, 1999) 8–55.

McGee, Vann, & McLaughlin, Brian. 1994. ‘Distinctions without a difference’. *Southern Journal of Philosophy*, **supp XXXII**, 203–251.

Merricks, Trenton. 2001. ‘Varieties of Vagueness’. *Philosophy and Phenomenological Research*, 145–157.

Parsons, Terence. 2000. *Indeterminate identity*. OUP.

Priest, Graham. 2001. *An Introduction to Non-classical logics*. Cambridge: Cambridge University Press.

Quine, W. V. 1953. ‘Three Grades of Modal Involvement’. *Chap. 15, pages 158–176 of: The Ways of Paradox and Other Essays*. Cambride, MA and London: Harvard University Press. First published in *Proceedings of the XIth International Congress of Philosophy*, Brussels 1953, Vol 14. (Amsterdam: North-Holland Publishing Co.).

Rayo, Agustín & Uzquiano, Gabriel. 1999. ‘Towards a theory of second-order consequence’. *The Notre Dame Journal of Formal Logic*, **40**(3), 315–325.

Rosen, Gideon, & Smith, Nick J. J. 2004. ‘Wordly indeterminacy: A rough guide’. *Australasian Journal of Philosophy*, **82**, 185–98.

Sider, Ted. 2009. ‘Ontological Realism’. *In:* Chalmers, David, Manley, David, & Wasserman, Ryan (eds), *Metametaphysics*. Oxford: OUP.

Smith, Nick J. J. 2005. ‘Vagueness as closeness’. *Australasian Journal of Philosophy*, **83**, 157–83.

Thomason, Richmond H. 1984. ‘Combinations of Tense and Modality’. *In:* Gabby, D. M., & Guenthner, F. (eds), *Handboook of Philosophical Logic*, vol. II. Dordrecht: Kluwer.

van Inwagen, Peter. 1990. *Material Beings*. Ithaca and London: Cornell University Press.

Williams, J. Robert G. 2008a. ‘Metaphysical indeterminacy and ontic vagueness’. *Philosophy Compass*, **3**.

Williams, J. Robert G. 2008b. ‘Multiple actualities and ontically vague identity’. *Philosophical Quarterly*, 134–154.

Williams, J. Robert G. 2008c. ‘Supervaluations and logical revisionism’. *The Journal of Philosophy*, **CV**(4).

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

Williamson, Timothy. 2003. ‘Vagueness in reality’. *Pages 690–716 of:* Loux, Michael, & Zimmerman, Dean (eds), *The Oxford Handbook of Metaphysics*.

Woodward, Richard. 2010. ‘Metaphysical indeterminacy and vague existence’. *Oxford Studies in Metaphysics*, **6**.

Wright, Crispin. 2001. ‘On being in a quandry’. *Mind*, **110**.