

Vagueness in a Precise World: Essays on Metaphysical Vagueness

by
Rohan Sud

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in The University of Michigan
2016

Doctoral Committee:

Professor Brian J. Weatherson, Co-Chair
Associate Professor David Manley, Co-Chair
Associate Professor Ezra R. Keshet
Professor Theodore Sider, Rutgers University
Associate Professor Eric P. Swanson

© Rohan Sud | 2016 | All Rights Reserved

To my sisters,
Anisha and Anjali Sud

Acknowledgements

First and foremost, I would like to thank my committee. I first became interested in vagueness while taking a fascinating seminar with Eric Swanson. I'm extremely grateful for the helpful advice he has given me during my time at Michigan, on the content of my dissertation and beyond. Early in graduate school, while I was still familiarizing myself with the landscape of the discipline's subfields, I was exposed to Ted Sider's work. That work initiated my decision to specialize in metaphysics: his writing is a constant source of philosophical inspiration for me and a model for my own research. I've probably spent more time talking philosophy with David Manley than any other professor and I feel very fortunate to have had that much exposure to such an insightful thinker. The enjoyment he takes in doing philosophy is contagious: talking philosophy with him is relaxed and fun – and always illuminating. Finally, I couldn't have asked for a more conscientious adviser than Brian Weatherson. He constantly encouraged me to pursue my boldest ideas and come to grips with the underlying philosophical issues buried deep in my arguments. Because of his encouragement, I'll continue that pursuit for the rest of my career.

Many of the ideas of this dissertation have been inspired by the research of a handful of philosophers working on issues at the intersection of vagueness and metaphysics – some of whom I've never had the pleasure of meeting: Elizabeth Barnes, Cian Dorr, John Hawthorne, JRG Williams, and Timothy Williamson. If this dissertation manages to add anything to the conversation on vagueness and metaphysics, it's because I had the privilege of their careful and creative work.

Thanks to the many members of the Michigan philosophy department, especially: Chloe Armstrong, Sara Aronowitz, Boris Babic, David Baker, Judith Beck, Gordon Belot, Mara Bollard, Paul Boswell, Sarah Buss, Nate Charlow, Mercy Corredor, Hobbes D. Derg, Billy Dunaway, Anna Edmonds, Alan Gibbard, Johann Hariman, Sydney Keough, Zoë Johnson-King, Ezra Keshet, Jason

Konek, Jeremy Lent, Eli Lichtenstein, Sarah Moss, Alex Plakias, Peter Railton, Cat Saint-Croix, Chip Sebens, Janum Sethi, Patrick Shirreff, Linda Shultes, Alex Silk, Daniel Singer, Nils-Hennes Stear, Jamie Tappenden, and Damian Wassell. Special thanks to Jim Joyce who gave me incisive comments on much of this (and other) work; he's been an important adviser during my time at Michigan. Thanks also to David Mark Kovacs, Alan Hájek, and Daniel Nolan.

Four graduate students have had played a particularly significant role in shaping my philosophical thought: Daniel Drucker, Dmitri Gallow, Steve Schaus, and Umer Shaikh. The many hours of conversations I've had with you four over the course of my time at Michigan have been, by far, the most valuable part of my philosophical education. You taught me how to think.

I would have never pursued a career in philosophy were it not for my undergraduate mentor Philip Kitcher. (I'm embarrassed to report to him that this is a dissertation in metaphysics (though much of it anti-metaphysical!))

Mercy: Everyone told me that, between the job market and dissertation defense, this year would be miserable. They didn't know I would have you in my life.

Many non-philosophers worked hard to keep me sane and grounded during graduate school: Michael Brener, Austin Carr, Tiara Forsyth, Matthew Harrison, Crystal Kim, Arvind Nagarajan, Nicholas Renkes, Samuel Rothschild, Steven Shlivko, Raj Vashi, Michael Vieten, and David Williams. Anisha and Anjali: I'll always look up to you two. Mom and Dad: I love you.

Contents

Dedication	ii
Acknowledgements	iii
Chapter 1: The Solution of the Many to the Problem of Vagueness	I
1 What is Supersententialism?	3
2 The Supervaluationist Machinery	7
3 What is Standard Supervaluationism?	8
4 Standard Supervaluationism and the Truth Objection	9
5 Non-Standard Supervaluationism	10
6 Non-Standard Supervaluationism and the Non-Reductive Complaint	13
7 Supersententialism and the Non-Reductive Complaint	17
8 The Lessons of the Many	20
9 The Problem of the Many Sentences	22
10 Objections Considered	25
11 Conclusion	34
Chapter 2: Vague Naturalness as Ersatz Metaphysical Vagueness	39
1 The GVAS	41
2 The Reference Magnetic Interpretation of the GVAS	45
3 Vague Naturalness as Ersatz Metaphysical Vagueness	54
4 Step One: Vague Natural Properties	55
5 Step Two: Vague Naturalness is not Metaphysical Vagueness	57
6 Step Three: Vague Natural Properties is Not Metaphysical Vagueness	64
7 Natural as Natural	72
8 Conclusion	75

Chapter 3: Ontological Deflationism, Vague Existence, and Metaphysical Vagueness		80
1	Characterization of Vague Existence	84
2	Negative Characterizations of Metaphysical Vagueness	90
3	Fuzzy Set Theory	94
4	The Modal Analogy and De Re Vagueness	97
5	Vague Properties and the Problem of the Many Properties	100
6	Vague Properties and the Return of the Modal Analogy	103
7	Vague States of Affairs	106
8	Two Deflationist Replies	109
9	Vagueness in Worldly Structure	111
10	Conclusion	114

Chapter I

THE SOLUTION OF THE MANY TO THE PROBLEM OF VAGUENESS[†]

The Problem of Vagueness (PoV) is not an easy one to solve. The most promising approach to solving the PoV is supervaluationism. On its traditional developments, however, supervaluationism is saddled with objections, the most significant of which is the Truth Objection. According to the Truth Objection, supervaluationism conflicts with the disquotational feature of the truth predicate as characterized by Tarski's T-Schema:

(T) '*p*' is true iff *p*

instances of which are generated by replacing *p* with various sentences. There is a great deal of intuitive support for the T-schema. So when a theory of truth treats instances of the T-schema as unassertable, it raises the suspicion that the theorist has failed to give an account of the concept of truth.

[†]Thanks to Daniel Drucker, Jim Joyce, Jeremy Lent, David Manley, Chip Sebens, and Eric Swanson for helpful conversations and comments. Special thanks to Brian Weatherston. Thanks also to participants of the University of Michigan Candidacy Seminar, the 10th Annual Mark L. Shapiro Graduate Philosophy Conference (and my commentators Geoffrey Grossman and Yongming Han), and audiences at the University of San Diego and the University of Michigan.

The Problem of the Many (PoM) is also not an easy problem to solve. The two most promising approaches to solving the PoM are the Solution by Vagueness and the Solution by Plentitude. According to the Solution by Vagueness, the PoM is an instance of the PoV. It therefore inherits the objections to solutions to the PoV, such as the Truth Objection. The second approach – the Solution by Plentitude – treats the PoM as largely independent of the PoV.

In this paper, I will use the tools from the Solution by Plentitude to defend a version of supervenience that I call *supersententialism*. According to supersententialism, instead of speaking a single language with multiple precisifications, we are simultaneously speaking many precise languages thereby tokening several precise sentences in a single speech act.

The core insight of my defense is that, on the Solution by Plentitude, the PoV is an instance of the PoM. In particular, I will show that, by borrowing tools originally developed for the Solution by Plentitude, supersententialism is the most plausible form of supervenience. If successful, this result would be significant for two reasons. First, because the Solution by Plentitude has been developed as an independently promising solution to the PoM, we would have an independently promising solution to the PoV. Second, if I'm right that, on the Solution by Plentitude, the PoV is an instance of the PoM, this should lend plausibility to the Solution by Plentitude over the Solution by Vagueness by massively increasing the payoff for adopting and developing the former.

Positions related to supersententialism have been considered in the past.¹ But when these nearby positions have been considered, the consideration has been brief and superficial with the result that the view is either quickly set aside as indefensible or is classified as a mere terminological variant of standard supervenience accounts (KEEFE, 2000, 1998; WILLIAMSON, 1994a).

¹Supersententialism bears important similarities (and differences) to the views discussed in SMITH (2008, §2.5), VARZI (2007), and DORR AND HAWTHORNE (2014) as well as views criticized by KEEFE (2000, 1998); WILLIAMSON (1994a). While the view of this paper draws inspiration from these, and other, sources, the view is different in important ways, which will be highlighted in the course of the paper. See footnotes 5, 6, 7, and §10.2.

This has been a mistake. Reflecting on the Problem of the Many reveals unappreciated benefits and new lines of defense from extant criticisms. In the end, we get a form of supervenience that shares many of the benefits of its rivals with none of their problems.

I What is Supersententialism?

While supersententialism retains several of the core insights – and the formal apparatus – of traditional supervenience, it is also motivated by the suggestions in Lewis’s early (1969; 1970; 1975) treatment of vagueness and recent work on speech act pluralism.

I.1 Vagueness Deported from Semantics

Supersententialism is inspired by the comments on vagueness in Lewis’s earliest works (1969; 1970; 1975). In those works, Lewis maintains that languages are formal objects that are precise in that each language assigns precise truth conditions to uninterpreted sentences. More specifically, *languages* are taken to be functions from uninterpreted sentence-formulas to sentential meanings (for example, sets of possible worlds). A community c is said to speak a language \mathcal{L} just in case there prevails, in c , a convention of truthfulness and trust in \mathcal{L} .

The conventions of a community can fail to determine a particular language as the unique language being spoken by that community. At best, the conventions of a language community delimit a class of languages being spoken. Lewis claims that this failure of convention to settle on a single language is the source of vagueness. For this reason, he suggests that vagueness arises from our *relationship* with these languages rather than from the languages themselves.

Although it is clear from these quotes that early-Lewis locates the source of vagueness in our relationship with language rather than in the language itself, his remarks don’t go much beyond this. Linda Burns (1991, §2.3, §9.5) subsequently

takes up Lewis's proposal and develops it into a form of contextualism²: highly transient facts settle which of the many languages delimited by convention is being spoken.³

And like their more modern developments (c.f. FARA (2000)), these early contextualist treatments are implausible. Without offering a detailed criticism of contextualist treatments of vagueness, I'll simply note the implausibility of the claim that contextual features are rich enough to determine a particular precise language being spoken by a community.⁴ Nevertheless, there is an important lesson to be drawn from the Lewis/Burns suggestion: vagueness is a feature of our relationship with language rather than a feature of the language we speak.⁵

1.2 Many Languages and Speech Act Pluralism

Convention only delimits a range of precise Lewisian languages as candidates for the language we speak. Yet it's incredible to think contextual non-conventional facts could help us select one among the candidate languages. So which language do we speak?

The solution to our quandary comes from reflecting on the recent proliferation of so-called speech act pluralist (SAP) views. According to these views, a single speech act – a single communicative act of producing sounds – can express multiple propositions. Some proponents of SAP (READ, 2009) use the view to

²Thanks to Brian Weatherson for helping me to see this position as a form of contextualism.

³Burns (1991) says "where there is vagueness speakers must be represented as alternating between members of a range of such languages" (182) and "speakers may adopt different languages from one another and shift from one language to another at different times" (186). It is less clear whether Lewis endorsed this contextualism, but he does make comments such as "we are free to settle these indeterminacies however we like" (1975, 188) and "the different languages of the cluster...may be differently suited for individual opinions, tastes, and conversational purposes. If everyone can pick from the cluster, incompatible preferences among languages may all be satisfied" (1969, 202).

⁴For further, and related, criticisms of the view, see KEEFE (2000, 148) and STANLEY (2003).

⁵Following Burns and Lewis, VARZI (2007) also considers the relative benefits of treating precisifications as *precise languages* rather than as ways of making a vague language precise. (He, however, does not claim that we are speaking all of the precise languages simultaneously.)

make progress on entrenched technical debates over the Liar paradox. Others (CAPPELEN AND LEPORE, 2005) are driven to SAP by reflecting on the enormous range of contents that ordinary speakers claim are being asserted by a typical speech act.

With speech act pluralism in mind, our puzzle dissolves. Instead of presupposing that we are only speaking one language among a range of languages, we should claim that we speak *all* of the languages delimited by our conventions of language use.

More carefully: languages are taken to be precise in the spirit of early Lewis. We can take languages to be set-theoretic objects that assign meanings to uninterpreted sentence types. Languages come equipped with a *grammar* where a grammar includes (1) a function from elementary lexical constituents to subsentential meanings (e.g. Carnapian intensions), and (2) a set of combination operations that build larger constituents from the elementary lexical constituents and that assign meanings to the built constituents based on the meanings of the elementary lexical constituents.

When I utter ‘Harry is bald’ I am simultaneously speaking many of these precise languages. I take the foregoing claim to come with several commitments. First: I am asserting many propositions simultaneously, each of which has precise truth conditions (i.e. for each proposition asserted, that proposition is determinately true or determinately false). Second: each time I make a particular utterance, I token many interpreted sentences – one for each precise language that I speak. That last claim has two components. There is an ontological component: I am postulating a plentitudinous ontology of interpreted sentences. There is also a linguistic component: I am claiming that a single

utterance can token each of these interpreted sentences.^{6,7} In borderline cases, one of the interpreted sentences that I utter s_1 is true, and another s_2 is false. Correspondingly, one of the propositions that I assert is true, and one of the propositions that I assert is false.

We now have a clear statement of the position. Admittedly, that position may appear incredible. We now turn to the task of motivating and defending the view. In order to do that, we must first examine the view's rivals. In the next several sections, I explore two rival forms of supervaluationism with which I will compare supersententialism. Because the formal apparatus of supersententialism is similar to these rivals, in the process of describing the alternative views, I will also give a model theory for supersententialism.

⁶In the course of discussing a problem with speech reports and semantic plasticity, DORR AND HAWTHORNE (2014, §5 (esp. pg. 333)) discuss a view of vagueness on which we are speaking many precise languages simultaneously. Because the view is of limited help with the particular problem they are investigating in their paper, it is not given a full development. In his (2014, §3.4 (esp. fn. 13)), Dorr voices his endorsement of a view on which we are speaking multiple languages simultaneously – although the view is not ultimately developed or defended.

⁷The view described as ‘plurivaluationism’ in SMITH (2008, §2.5) is most similar to the present view. Without endorsing the view, he notes the Lewis/Burns contextualist treatment of vagueness, and carves out a position in logical space on which there are multiple precisifications that are ‘correct’ or ‘intended’. The discussion in this paper advances Smith’s in two ways. *First*: it’s not clear whether Smith’s plurivaluationist is the same as supersententialism. Although his plurivaluationist claims that “there is not a unique or correct interpretation”, the plurivaluationist stops short of explicitly embracing speech act pluralism. Although that seems to be a natural consequence of the view, there is no discussion of this radical claim. And, Smith’s plurivaluationist explicitly disavows the claim that there is a notion of truth *simpliciter* that applies to sentences. Instead, there is supposed to be only a notion of truth on this-or-that acceptable interpretation as applied to uninterpreted sentences. Central to the supersententialist is the claim that there are several interpreted sentences tokened in any speech act (hence the name of the view), each of which is true or false *simpliciter*. This ontological and linguistic claim is what allows her to (1) retain the highly intuitive idea of a sentence being true *simpliciter* (2) to treat the T-schema as an instance of the Problem of the Many (see §9) and (3) to defend the view from the criticisms of Williamson and Keefe (see §10.2). *Second*: the present discussion offers a motivation and defense for the view that is not contained in Smith’s discussion.

2 The Supervaluationist Machinery

Supervaluationism is typically accompanied with a certain sort of formal machinery closely resembling that of quantified modal logic (QML) with classical ‘precisifications’ playing the analogous role of possible worlds. Consider a simple first-order language consisting of variables, names, predicates, connectives, quantifiers, and the sentential operator Δ (for ‘determinately’). The grammar of the language is defined analogously with QML with Δ in place of the necessity operator (\Box). And as with QML (with a constant domain), a model is a 3-tuple $\langle P, D, I \rangle$ consisting of a set of points (the ‘precisifications’) (P), a domain of objects (D) and an interpretation function (I) which takes as input a pair consisting of a word (names or predicates) and a precisification and has as output the referent of that word according to that precisification (an object in the case of a name or a set of n -tuples in the case of a n -place predicate).⁸ Variable assignments assign variables in the language to objects. In order to define a notion of truth *simpliciter*, we begin by defining notions of truth relative to variable assignments, precisifications, and models.

Truth-at-a-precisification is defined analogously with the notion of truth-at-a-world in QML. We first define truth-at-a-precisification-relative-to-a-variable-assignment for atomic sentences in the natural way. Next we apply the standard recursive definitions to define truth-at-a-precisification-relative-to-a-variable-assignment for complex sentences. Finally, we say that a sentence s is true-at-a-precisification p if there is some variable assignment a such that s is true-at- p -relative-to- a .

⁸For simplicity of exposition, I suppress the accessibility relation between precisifications. We’ll assume a logic of S5, such that the necessity operator quantifies over all points in the model. The model may be complicated to account for higher-order vagueness or when giving a semantics for the classical rules of inference that respects global-validity. See WILLIAMS (2008) and WILLIAMSON (1999).

3 What is Standard Supervaluationism?

Above, we defined the notion of truth relative to a particular precisification in a model. But we stopped short of defining truth in a model and truth simpliciter. How should one define these notions of truth from the stipulated definitions of truth-at-a-precisification? Supervaluationists define the technical notions of *supertruth* (and *superfalsity*):

(*Supertruth*) s is supertrue (superfalse) on a model $\mathcal{M} = \langle P, D, I \rangle$ iff for all precisifications $p \in P$, s is (not) true-at- p

Standard supervaluationism identifies truth (falsity) on a model with the defined notion of supertruth (superfalsity) on that model:

($T=ST$) s is true (false) on a model \mathcal{M} iff s is supertrue (superfalse) on \mathcal{M} .

And, as usual, truth (falsity) simpliciter is taken to be truth (falsity) on the intended or correct model. What determines which model is the intended one is a matter of substantial debate among meta-semanticists, a debate over which the present account remains neutral.

($T=ST$) allows the standard supervaluationist to give an account of many of the features that distinguish borderline sentences. Most importantly, the thesis allows the standard supervaluationist to explain the unassertibility of borderline sentences. Suppose s is a sentence that is true on some but not all precisifications of the intended model. For example, take s to be

(1) Bob is bald

where Bob is a borderline case of baldness. Given ($T=ST$), (1) is not true. And, given ($T=ST$), the negation of (1) is not true. This result, combined with a maxim to assert only truths

(*Truth-Assertion*) Assert s only if s is true

allows the supervaluationist to explain our hesitation in asserting either (1) or its negation.

4 Standard Supervaluationism and the Truth Objection

Standard supervaluationists' adherence to $(T=ST)$, however, leads to the Truth Objection – some instances of the T-schema are not assertible. From a borderline sentence like (1), there are two paths to arriving at the unassertibility of an instance of the T-schema.⁹

4.1 Path One

Note that the law of excluded middle holds for the supervaluationist. In particular, the relevant instance

(2) Bob is bald or it's not the case that Bob is bald

is supertrue. Now suppose for *reductio* that the supervaluationist accepts the relevant instances of the disquotational schema:

(3) 'Bob is bald' is true iff Bob is bald.

(4) 'It's not the case that Bob is bald' is true iff it's not the case that Bob is bald.

Then, substituting the left biconjuncts from (3) and (4) into (2) we get:

(5) 'Bob is bald' is true or 'It's not the case that Bob is bald' is true

By $(T=ST)$, (5) is claiming that either a sentence or its negation is supertrue. But we concluded above that neither (1) nor its negation is supertrue when the sentence is true on some but not all precisifications of the intended model. Short of giving up some classical rule of logic, supervaluationists are forced to admit that some instances of the T-schema – like (3) and (4) – are not true and, by (*Truth-Assertion*), are therefore not assertible.

⁹WILLIAMSON (1994b, 162) points out the first path and KEEFE (2000, 214) points out the second.

4.2 Path Two

Consider again the relevant instance of the T-schema:

(3) ‘Bob is bald’ is true if and only if Bob is bald.

From $(T=ST)$ we learned that (1) is not true (because it’s true on some precisifications and false on others):

(6) It’s not the case that ‘Bob is bald’ is true

Importantly, the claim made by (6) is not borderline – if it were borderline, then it would not be true.¹⁰ But of course, (6) is true. By $(T=ST)$ it’s supertrue, which is to say it’s true on every precisification. Because (1) is not superfalse, there is some precisification p^* on which it is true. But then note that (3) is false on p^* : its left conjunct is false (because it’s false on every precisification) and its right conjunct is true on p^* . Supervaluationists are therefore forced to admit that some instances of the T-schema – like (3) – are not true and, by *(Truth-Assertion)*, are therefore not assertible.

But plainly, the objection goes, all instances of the T-schema are true and are assertible. Indeed, the T-schema features so centrally in our concept of truth that its instances appear to be trivialities. To give up on the T-schema would be to give up on a central feature of truth and raises the suspicion that standard supervaluationists have changed the subject in their attempt to elucidate ordinary truth.

5 Non-Standard Supervaluationism

Non-standard supervaluationists¹¹ attempt to retain the original T-schema while maintaining some of the core ideas of the standard supervaluationist framework.

¹⁰Again: we’re setting aside higher-order vagueness here. With higher-order vagueness, the point is simply that (6) does not adopt the first-order vagueness of (1).

¹¹I have in mind here the view outlined in FIELD (1994).

In particular, they agree with the standard supervaluationist on at least two points related to the formal models. They accept that

(*Core Idea 1*) Typical borderline sentences are not supertrue; there is some precisification p such that the sentence is not true-at- p .

(*Core Idea 2*) A sentence is assertible only if it's supertrue.

Yet, non-standard supervaluationists wish to hold on to the view that all instances of the T-schema are assertible. Therefore, non-standard supervaluationists must accept the following: for any instance of the T-schema, that instance is true on any precisification in the intended model. Recall our borderline sentence (1) which is true on some but not all precisifications of the intended model and the relevant instance of the T-schema (3). By stipulation, the right bijkunct of (3) is true on some but not all precisifications. For the bijkunct to be true on all precisifications, the left bijkunct must be true on exactly those precisifications the right bijkunct is true on – the sentence “Bob is bald’ is true’ is borderline in the same way (1) is. As non-standard supervaluationists would put it: attributions of truth must inherit the vagueness of the quoted sentence.

Non-standard supervaluationists agree with the standard supervaluationist on *Core Ideas 1* and *2*. So over what do the non-standard supervaluationists disagree with the standard supervaluationists? *Core Ideas 1 and 2* give an account of the unassertibility of borderline sentences in terms of the *formal apparatus* of supervaluationism, but don't connect that apparatus with notions of truth. It is here that the standard supervaluationist and the non-standard supervaluationist disagree: they give different accounts of the notions of truth in a model and, more importantly, truth simpliciter.

There appear to be only two feasible options for an account of the notion of truth-in-a-model.¹² Either a supervaluationists can take truth in a model to be truth at all the precisifications, thereby accepting ($T=ST$). Or they can take

¹²There is third option: one could take truth to be truth at *some privileged subset* of precisifications (c.f. WILLIAMS (2008)); this version also succumbs to the truth objection.

truth in a model to be truth at a particular precisification. As explained above, the first option leads to the unassertibility of the T-schema. The T-schema and classical logic implies that a sentence or its negation is true. Together with $(T=ST)$ we would get the result that every sentence is true on all precisifications or false on all precisifications. But (with *(Core Idea 1)*) that would rule out the existence of borderline sentences.

Thus, the non-standard supervaluationist must pursue the second option by modifying the definition of truth in the following way. Add to the model another element so that models consist of a 4-tuple $\langle P, c, D, I \rangle$ where $c \in P$. Intuitively, c is taken to be the ‘correct’ precisification. Then define truth (falsity) in a model as truth (falsity) at c and truth (falsity) simpliciter as truth (falsity) in the intended model. As a result, each sentence is either true or false simpliciter – bivalence is preserved. The model theory, then, is exactly analogous to the model theory of epistemicists or that of QML.

On this definition of truth, both paths to the Truth Objection are blocked. With respect to the first path: because bivalence is preserved and each precisification is classical, we get the result that either a sentence or its negation is true. Thus (5) is no longer a problematic prediction of the view. With respect to the second path: we can admit that (6) is borderline and not assertible.

Does the acceptance of bivalence amount to a denial of vagueness? No, for it can be indeterminate which model is the intended one. So, while the notion of truth-in-a-model is a precise notion, attributions of truth simpliciter can be indeterminate.¹³

¹³Many non-standard supervaluationists – like the one represented in FIELD (1994) – are skeptical of the explanatory significance of the semantic notions formalized by the model. While I think the model I present is the best way to supplement the non-standard supervaluationist view with a formal model, I do not mean to commit myself to any claim as to the formal apparatus’ explanatory import.

6 Non-Standard Supervaluationism and the Non-Reductive Complaint

When non-standard supervaluationism is discussed in the literature, the view that one typically has in mind is the one offered in Hartry Field's early (1994) work.¹⁴ According to the non-standard supervaluationist, vagueness of the truth attribution is explained by taking the truth predicate to be a vague one and positing a penumbral connection between the vagueness in the truth predicate and the vagueness in the quoted sentence. Allowing the predicate 'is true' to be vague opens space for the necessary flexibility to allow for the posited penumbral connection, so that the truth of the left biconjunct of the T-schema can vary with the truth of the right biconjunct, across all precisifications.

Considering non-standard supervaluationism, Williamson (1994b, 164) writes:

What then remains of supervaluationism? There remains the 'definitely' operator with its semantics of admissible interpretation. However, this apparatus has lost its privileged connection with the concept of truth. Of any admissible valuation, we can ask whether it assigns truth to all and only the true sentences of the language and falsity to all and only the false ones. At most one valuation has that property. But then any other valuation will assign truth-values incorrectly, so how can it be admissible? It might be replied that no interpretation is definitely the one with the desirable property.

According to the non-standard supervaluationist there is one, and only one, precisification in the intended model that assigns truth to all and only the true sentences. Call that precisification the *correct* precisification. Truth simpliciter is defined in terms of truth at this correct precisification. All other precisifications are incorrect, so their role is restricted to their role in giving the semantics for the determinately operator (Δ). And that role is consistent with the non-

¹⁴How to interpret Field's stance in his (1994) is not straightforward. While he outlines the view I am discussing as a solution to the Truth Objection available to a deflationist about truth, he stops short of endorsing it. As I read him, the view he ends up endorsing in that paper is closer to the view of MCGEE AND McLAUGHLIN (1994) according to which the truth predicate has two competing meanings. (However, he later gives up on supervaluationism altogether and adopts three-valued logic with a sophisticated account of the conditional. See his FIELD (2003).)

standard supervenientist's definitions of truth. Recall that the correct precisification is not *determinately* correct – there are unintended models that are not determinately unintended, according to which the alternative precisifications in the intended model have the property of being the correct precisification. So, it's no wonder that the alternative not-determinately-incorrect precisifications figure into the semantics for the determinately operator (Δ).

It's helpful to compare the point with the semantics given for modal languages. In the intended model for modal languages, there is one, and only one, actual world and truth simpliciter is defined in terms of truth at the actual world. All the other worlds in the intended model are non-actual so their role in the semantics is restricted to their role in giving the semantics for the necessity operator (\Box). And that role is consistent with the modal logician's definition of truth simpliciter as truth at the actual world. There are unintended models that *could have been* the intended models, according to which the other possible worlds in the intended model have the property of being actualized. So, it's no wonder that other possible worlds figure into the semantics for the necessity operator (\Box).

Nevertheless, as Williamson (1994b; 1994a) has taken pains to point out, the view suffers from a less obvious (yet no-less-important) issue.¹⁵ Continuing from the above quote, Williamson (1994b, 164) writes:

It might be replied that no interpretation is definitely the one with the desirable property. Once definiteness has been separated from truth, the reply is without force. If an interpretation does have the desirable property, why should it matter if it does not definitely have it? Indeed, the reply is in danger of losing its sense as well as its force. If we cannot grasp the concept of definiteness by means of the concept of truth, can we grasp it at all? No illuminating analysis of 'definitely' is in prospect. Even if we grasp the concept as primitive, why suppose it to be philosophically significant?

Earlier, we suggested that the non-standard supervenientist take the determinately operator to quantify over the not-determinately-incorrect precisifications. This however does not constitute an *account* of the determinately

¹⁵A similar complaint can be found in KEEFE (2000, 205-206)

operator because it appeals to the notion of determinateness in explaining the relevance of the incorrect precisifications in giving the semantics for the determinately operator. We have a formal apparatus for understanding talk of determinateness, and while that apparatus is illuminating in some ways, it does little to illuminate the source of indeterminacy. The same point can be made in the modal case. A possible-world semantics for modal talk, as enlightening as it is, appeals to primitive modal notion and leaves the source of modality open to interpretation.

One option is to take the notion as a primitive. This is the option pursued in FIELD (1994) and BARNETT (2009). Williamson's objection to this strategy seems to be that, if taken as primitive, an explanation must be given of why the concept of determinateness plays its distinctive philosophical role. I take that role to be primarily the role determinateness plays in our patterns of assertion: if a sentence *s* is indeterminate, that sentence is unassertible.

An analogous criticism is inapplicable to the standard supervenientist. Because they identify determinateness with the everyday notion of truth, their story as to why determinateness plays the philosophical role it does goes via the philosophical role that truth plays. For instance, we have a simple explanation of the unassertibility of indeterminate sentences in terms of truth – namely that we shouldn't assert sentences that aren't true (codified in *Truth-Assertion* above).

The challenge Williamson lodges on non-standard supervenientists who take the notion as primitive, however, is not obviously a fair one. In particular, for those of us that take concepts to be identified by their conceptual role, Williamson's complaint is impotent. Non-standard supervenientists take a concept as primitive and outline a conceptual role that this concept is meant to play. That role is 'philosophically significant' in that it relates to concepts like assertibility which philosophers tend to study. There is nothing more that needs to be explained. Once both disputants acknowledge that a concept that plays the specified role exists, and that the concept is primitive, there is no further explanatory demand as to why *that* particular primitive concept plays *that*

particular role. That's not to say that positing primitive concepts is cost-free. We have reason to prefer theories of our conceptual structure that appeal to fewer primitive concepts.

But I think Williamson's complaint obscures the more serious reason for being unsatisfied with the account of determinateness given by the non-standard supervenient. Talk of concepts is a red herring. Consider a dispute between the epistemicist and a standard supervenient. That dispute is not – or at least not *only* – about our *concept* of determinateness. Either party may take the concept of determinateness to be *conceptually* primitive (however conceptual primitiveness is best understood). The important question is the *metaphysical* one: how do facts about indeterminacy reduce to more metaphysically fundamental facts. The epistemicist has (the start of) a story: facts about indeterminacy are reduced to facts about knowability and facts about knowability are reduced to facts including those about the plasticity of reference. And the standard supervenient has (the start of) a story: the fundamental facts that fix the facts about which interpretations are admissible fails to privilege a particular interpretation – a range of interpretations conform equally well with the reduction base and have equal claim to being the intended interpretation.

Presumably the non-standard supervenient will not want to claim that the determinately operator or facts involving the operator are metaphysically fundamental, even if she admits its conceptual fundamentality. Taking the determinately operator as metaphysically fundamental risks collapsing the view to a highly implausible form of metaphysical indeterminacy.¹⁶ But if the operator is not fundamental, then the account is incomplete: admitting conceptual primitiveness does *nothing* to explain the operator's reductive basis. That's the real reason to be unsatisfied with non-standard supervenientism. They simply haven't given us a complete account. Compare the analogous point in the

¹⁶This is, for example, how BARNES AND WILLIAMS (2011) understand their view: they take the determinately operator or facts involving determinacy to be metaphysically fundamental. They claim that this makes them invulnerable to Williamson's complaint (which indicates that they too view Williamson's complaint as metaphysical, rather than conceptual, in nature).

modal case. The dispute over various accounts of modality is over the metaphysical reduction of modal facts (the linguistic ersatz was never making a claim of conceptual reduction!)

Let's call this the *Non-Reductive Complaint*. The non-reductive complaint is especially troubling once we recognize that epistemicism *does* give an account of the source of indeterminacy in a way that is entirely consistent with the non-standard supervaluationist position as elucidated so far. The epistemicist accepts the same claims as the non-standard supervaluationist – perhaps even including the conceptual primitiveness claim. For instance, they claim that there is a correct precisification in the intended model and that the other precisifications are not determinately incorrect. But, the epistemicist can give a more complete account by filling in the details as to the metaphysical grounds of indeterminacy.¹⁷ The appeal of epistemicism, then, is its ability to avoid both the Truth Objection and the Non-reductive Complaint. But, like the early contextualist solution suggested by Lewis and Burns, the position rests on the incredible claim that the non-semantic facts are rich enough to select a unique precisification as correct.

7 Supersententialism and the Non-Reductive Complaint

The supersententialist, I'll now argue, can resist the difficulty that Williamson pointed out for attempts to develop non-standard supervaluationism. Let's take a moment to recall that difficulty. The non-standard supervaluationist's woes began at the final step of the specification of the model-theory for vague languages. She defined 'truth-in-a-model' from truth at the correct precisification, and then took truth simpliciter to be truth in the intended model. Standard supervaluationists defined truth-in-a-model in such a way that ruled out instances of the T-schema as applied to borderline sentences. The non-standard supervaluationist was able to make some progress on this front by adding a parameter to

¹⁷This is how WILLIAMSON (1994b) uses the complaint to motivate epistemicism.

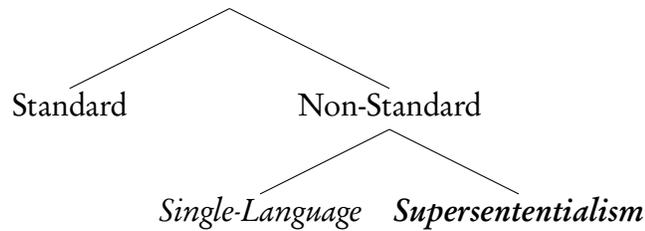


Figure 1.1: Varieties of Supervenience

the models (c , the ‘correct’ precisification) and re-defining truth-in-a-model using that parameter. However her project ran into problems when she attempted to define truth simpliciter as truth in the intended model. Because there is only one such model, we lost our grip on the relevance of the (not determinately) unintended models corresponding with the other precisifications in the intended model.

The appropriate point of resistance is located at the final step of Williamson’s argument: there is no need to assume that there is only *one* intended model. If the community speaks many languages simultaneously, there are several models that are intended by a community: one for *each* of the languages that the community speaks. Each intended model will share the same set of precisifications and for each precisification, there will be an intended model that selects that precisification as correct.

The supersententialist is a non-standard supervenience theorist – she denies that truth in a model is supertruth. Unlike other non-standard supervenience theorists, however, she thinks that there are many intended models. She should therefore adopt wholesale the formal apparatus of the non-standard supervenience theorist, in which a single model represents the semantic properties of a single one of the many interpreted sentences uttered. We can take models to be the same as those described above: models consist of a 4-tuple $\langle P, c, D, I \rangle$ where $c \in P$. Intuitively, c is taken to be the ‘correct’ precisification and we define truth (falsity) in a model as truth (falsity) at c .

If there are several intended models, related in the way described above, then

it's clear why, for any particular intended model, the other precisifications are relevant. The other precisifications are relevant because there are *other* intended models according to which these other precisifications are correct. More simply: the other precisifications are relevant because we are also speaking languages on which the truth is given by the other precisifications!

Most importantly, this account of determinacy explains the 'philosophical significance' of the determinately operator. Consider the role that determinacy plays in assertion. We don't assert $\ulcorner p \urcorner$ when p is indeterminate. But that follows from (*Truth-Assertion*). If p is indeterminate, then to utter $\ulcorner p \urcorner$ would be to assert *something* false – at least one of the interpreted sentences we would be asserting would be false, even if some other interpreted sentences are true. Our goal in assertion is not only to speak the truth, but also to avoid speaking falsehoods which explains why the determinately operator plays the role it does in our practice of assertion.

The supersententialist, I've argued, does not succumb to Williamson's criticisms of non-standard supervaluationism. But have we lost sight of our original goal? If supersententialism is to be preferred to standard supervaluationism, it must avoid the Truth Objection. In other words, supersententialists must explain the assertibility of all instances of the T-schema. In particular, they need to explain the assertibility of claims like (3). If each utterance of the T-schema is to be assertible, the uninterpreted sentence must be true on all of the intended semantic models. Equivalently, it must be true on all of the precisifications of some intended model. That is to say, it must be true in all the languages being simultaneously spoken.

This is no easy task for the supersententialist. Unlike the non-standard supervaluationist, the supersententialist cannot appeal to the vagueness of the truth predicate: any interpreted sentence is precisely true or precisely false. At this point it may appear as though we've taken one step forward and several steps backwards. The supersententialist doesn't succumb to Williamson's arguments against non-standard supervaluationism. But it's not clear that the view

rebutts the Truth Objection. And what's more, the view required us to adopt the strange claim that we are simultaneously speaking many languages!

The key to developing the position is to examine a different debate – the debate over the Problem of the Many – and apply the lessons from that debate to our present dialectic.

8 The Lessons of the Many

Consider an ordinary cat, Tibbles, on a mat. Tibbles is shedding and some of her hairs are in the process of falling off of her body. Consider one such hair, *h*. For that hair *h*, it's indeterminate whether it is part of Tibbles – it's a borderline part of Tibbles. Now consider the merological sum of Tibbles and *h*. Call that object Tibbles+. Consider the merological difference of Tibbles and *h*. Call that object Tibbles-.

Now consider the following claims:

- (7) a. There is exactly one cat on the mat.
- b. Tibbles is on the mat.
- c. Tibbles is Tibbles+.
- d. Tibbles is Tibbles-.

Both Tibbles+ and Tibbles- have equal claim to being a cat: there is no feature of catiness that could distinguish between Tibbles+ and Tibbles-. Yet, the first claim seems plainly true: there is only one cat on the mat. The second claim also appears to be true. Together with the first claim, we get the result that there is one and only one cat on the mat, and that cat is Tibbles. Yet, the last two claims are unassertible: we hesitate to identify Tibbles with Tibbles+ or to identify her with Tibbles-. This is one case of the infamous Problem of the Many: we must explain the assertibility of the first two claims and the unassertibility of the second two.

One attempt to solve the problem posits a sort of dualism between Tibbles and the objects Tibbles+ and Tibbles-. On one development, there is a metaphysically vague object that ‘Tibbles’ determinately refers to (VAN INWAGEN, 1990; TYE, 1996). On another development, Tibbles is not identical to any material object, but is rather indeterminately constituted by one of the material composites (LOWE, 1995) or constituted by multiple material composites (JONES, 2015).

LEWIS (1993) claims, and I agree, that dualism is a solution of last resort. This is the first lesson to draw from the Problem of the Many.

No Dualism Don’t posit a dualism between a ‘vague object’ and a ‘precise object’; there are no vague objects.

Instead, Lewis advocates that a less costly solution locates the source of the Problem in our language rather than in fundamental metaphysics, and proposes two such solutions. On the first solution, Lewis takes the Problem to be a case of vagueness, and applies the standard supervaluationist machinery to the Problem. On this solution, the predicate ‘cat’ is vague; on one admissible precisification, Tibbles+ is in the extension of ‘cat’ and on another precisification Tibbles- is in the extension of ‘cat’ (but in no precisification are both Tibbles+ and Tibbles- within the extension). In order for (7b) to come out as assertible, the standard supervaluationist must also claim that the name ‘Tibbles’ is vague while positing a penumbral connection between the name and the predicate ‘cat’ such that, for any precisification, the referent of ‘Tibbles’ according to that precisification is in the extension of ‘cat’ according to that precisification. This suggestion immediately predicts the desired pattern of assertibility for the four claims above. Call this the *Solution by Vagueness*.

On Lewis’s second solution, Tibbles+ and Tibbles- are *both* cats. Yet, Lewis claims that ordinary notions of counting are by ‘almost identity’ rather than identity, where two objects are almost identical just in case they massively overlap. Call this the *Solution by Plenitude*. The solution by plenitude must also help

itself to some of the resources of the Solution by Vagueness.¹⁸ The solution by plenitude admits that, strictly speaking, there are many cats on the mat. The plenitudinous-philosopher, therefore, rejects the vagueness-philosopher's insistence that the predicate 'cat' is vague.¹⁹ However, the plenitudinous philosopher ought to continue to admit that the singular terms 'Tibbles' and 'the cat on the mat' are vague. That explains our hesitation in asserting (7c) or (7d).

There has been an ongoing debate over which is the right solution to the Problem of the Many. Lewis, for instance, claimed that both solutions are adequate, and that each required elements of the other. I happen to think that both solutions to the Problem of the Many are serious contenders.²⁰ I therefore remain agnostic between the two solutions as solutions to the Problem of the Many. However, the research underlying the Solution by Plenitude can be fruitfully applied in a defense of supersententialism.

There are at least two lessons from the Solution by Plenitude worth distilling.

Vague Singular Terms Our singular terms vaguely refer to one of the many cats on the table.

Almost-Identity We count by almost-identity rather than strict identity.

9 The Problem of the Many Sentences

The supersententialist is facing an instance of the Problem of the Many, with tokened sentences instead of cats. Drawing out this parallel, we can see that the

¹⁸LEWIS (1993) makes this point. He also argues that the Solution by Vagueness should help itself to elements of the Solution by Plenitude, but that is not relevant for my argument here.

¹⁹At least vague in the same way that the vagueness philosopher claims – the plenitudinous philosopher might admit vagueness as to which plurality of cat-candidates count as the many cats on the mat.

²⁰See WILLIAMS (2006) for one powerful argument in favor of the Solution by Plenitude. See WEATHERSON (2003) for an excellent defense of the Solution by Vagueness.

supersententialist has the tools to rebut the Truth Objection. The supersententialist claims that we speak several languages, one for each of the precisifications in the supervaluationist model. When I make a particular utterance, I am in fact uttering many sentences. When I make a borderline utterance of “Bob is bald”, I am producing several interpreted sentences, each with precise truth conditions. Some of these sentences are true; some of these sentences are false. Because we refrain from asserting falsehoods, and we cannot assert the true sentences without also asserting the false sentences, we hesitate to utter any of the sentences. Thus truth under all of the precisifications – or truth under all of the languages – is a necessary condition for assertibility.

Consider once more (3):

(3) ‘Bob is bald’ is true iff Bob is bald.

If our theory is to avoid the Truth Objection, we must show how (3) is true on all of the precisifications (equivalently: true on all of the languages being spoken). The right-hand side of the biconditional is true on some precisification and false at others; our challenge is to explain how the left-hand side of the biconditional is true at exactly those precisifications. The non-standard supervaluationist solved her analogous problem by taking the truth predicate to be vague in the same way as the quoted sentences. But, we noted, the supersententialist cannot make the parallel move. Because each of the sentences have precise truth conditions, the truth predicate (as applied to these interpreted sentences) cannot be vague.

By positing a plenitude of sentences being uttered in a single utterance, however, the supersententialist can apply the lesson **Vague Singular Terms** from the Problem of the Many. Although there is no room to posit vagueness in the predicate, we can posit vagueness in the singular term – the sentence being named by the quotation marks.²¹ The supersententialist can also posit a

²¹Compare the move made at the end of WEATHERSON (2003) in which Weatherston takes the name of a predicate to be vague in order to rebut an argument by MERRICKS (2001) for metaphysical indeterminacy.

penumbral connection between the quotation name “Bob is bald” and the right bijnunct ‘Bob is bald’ such that the sentence named under a precisification is true just in case ‘Bob is bald’ is true on that precisification. In less standard but more perspicuous terms: for any language L being spoken, the interpreted sentence being named in L by the quotation in an utterance of (3) has the same truth conditions as the sentence ‘Bob is bald’ as interpreted by L .²²

This move is available to the supersentialist because she posits several languages being spoken simultaneously. Without a plentitudinous ontology of interpreted sentences, there is no plentitude of candidates for which the quotation-name can vaguely refer to. Such a plentitude is vital to opening up the possibility of positing vagueness in the singular term. The failure to appreciate this significant advantage is perhaps what has led many supervalutionists to be dismissive of proposals similar to supersentialism. For example, KEEFE (2000, 1998) accuses a proposal akin to supersentialism as collapsing into standard supervalutionism or mere notational variants of the view.²³ I agree that supersentialism has certain structural and formal parallels with supervalutionism – namely it endorses (*Core Idea 1*) and (*Core Idea 2*) – which is why I see the view as a form of supervalutionism. But by claiming that we speak many precise languages, the view posits a plentitudinous ontology of sentences that are uttered together. Once this advantage is appreciated, however, the suggested way around the truth objection becomes obvious. And, the supersential development of non-standard supervalutionism, unlike other developments, is not susceptible to the non-reductive complaint as we established in §7.

²²Note that the same sort of solution can be applied to a propositional disquotational scheme – it’s true that p iff p – according to which the that- clause is a vague singular term between various propositions with precise truth-conditions.

²³Keefe directs her attack against BURNS (1991) as interpreted as a supersentialist. However, Burns’ actual theory differs from supersentialism, as explained above (§1.1).

10 Objections Considered

In this section, I reply to various objections to the view outlined above. One theme running throughout these replies is an application of research done on the Problem of the Many in service of rebutting the present objections.

10.1 We Speak One Language

Objection. According to you, I am speaking many languages at once. As much as I might wish that I were multilingual, I am not. I speak one and only one language: English! Yet according to the supersentialist, I am in fact speaking many languages.

Reply. Recall our lesson **Almost-Identity** from Lewis's Solution by Plenitude to the Problem of the Many. Lewis accepts that there are in fact many cats on the mat but claims that we count only one cat on the mat because we count by almost-identity rather than strict identity, where two objects are almost-identical if they massively overlap. Together with the claim that the name 'Tibbles' indeterminately refers to one of the many cats, we predict the relevant data.

The supersentialist can make use of the same strategy. According to the supersentialist, there are in fact many languages that we speak, but we count only one language because we count languages by almost*-identity, where two languages are almost*-identical when the intensions for the terms massively overlap.²⁴ Together with the claim that the name 'English' indeterminately refers to one of the many languages, we predict the relevant data. (Speaking loosely) we do speak one and only one language: English! The same solution can be extended from number of languages spoken to number of sentences ut-

²⁴Of course the relevant notion of overlap used to define almost*-identity is not the mereological notion Lewis used to define almost-identity, but I take it that the notion is obvious enough. For example, take intensions to be functions (i) from worlds (w) to sets of objects and consider a suitable measure m . Two intensions i_1 and i_2 massively overlap if for all w in the domain of both i_1 and i_2 , $m(i_1(w) \cap i_2(w))$ is very close to $m(i_1(w))$ and very close to $m(i_2(w))$.

tered, in the obvious way, to generate the result that I assert exactly one sentence.

Indeed, I think there is a good case to be made, independently of our discussion of vagueness, that we already have a grip on the notion of almost*-identity and count languages by something like almost*-identity. Compare the language spoken in the US according to which ‘torch’ means something that produces fire with the language spoken in England according to which ‘torch’ means flashlight.²⁵ By Leibniz’s law, these languages are not strictly identical. Yet no American-English speaker claims to be bilingual as a result of speaking British-English.

10.2 What About the Unsharp Sentence?

Objection. You’ve shown that the T-schema holds for each of the precise languages that we speak. But that doesn’t solve the original problem. We were supposed to show that the T-schema holds for the *imprecise sentence*, not merely the precise sentences! Williamson (1994a) puts the objection best when he considers a proposal similar to supersentialism:

The utterance ‘*a* is a heap’ is by ordinary standards the utterance of a single sentence. If it is held that in reality many sentences were uttered, a special way of individuating sentences must be in play...Call the many nonstandardly individuated co-uttered sentences *sharpened*, and the original standardly individuated sentence *unsharpened*. One can still ask for the truth-value of the unsharpened sentence...the disquotational principle must fail for unsharpened sentences...According to the original intuition, if *a* is a heap then the unsharpened ‘*a* is a heap’ is true...(180)

Reply. My opponent, and Williamson, appear to have a different view in mind than the view I am endorsing. According to my view, *there is no such thing* as the ‘unsharpened sentence’. There are only a plurality of sharp sentences. That follows from the lesson **No Dualism** from the Problem of the Many: don’t adopt dualism unnecessarily. To posit an additional object – the ‘unsharpened

²⁵Thanks to Chip Sebens for this helpful example.

sentence’ – over and above the many sharpened sentences is to adopt a sort of ‘dualism’ about sentences that is no part of my theory. Just as I don’t posit an unsharp cat over and above the many sharp cats when solving the Problem of the Many, I do not posit an unsharp sentence over and above the many sharp sentences when solving the Problem of Vagueness.

Williamson appears to have assumed that only a dualism between unsharp and sharp sentences can accommodate our original intuition behind the T-schema, which appears not to be directed towards these precise sentences. But again, the analogy with the Problem of the Many is useful here. If there is a sense in which the original intuition behind predications of the truth predicate didn’t apply to sharp sentences but rather to an unsharp sentence, then it’s the same sense in which our original intuition behind predications of the cat predicate didn’t apply to sharp cat-candidates but rather an unsharp cat. I’ve successfully accounted for both the unsharp-cat intuition and the unsharp-sentence intuition by positing vagueness in the relevant singular terms.

Objection (cont’d). But consider the sentence that has the truth-conditions as posited by the standard-supervaluationist, complete with truth-gaps. Denying the existence of such an unsharpened sentence is unmotivated. Here’s an argument from Keefe (2000) given when considering a proposal similar to supersententialism according to which there are no unsharp sentences:²⁶

[The supersententialist might insist] that talk of ‘the vague language’ is not acceptable...[but she] has no grounds on which to maintain this insistence. For she adopts Lewis’s account of the role of beliefs and intentions in fixing the language spoken and hence the meanings of sentences, and if we are seeking to be truthful and trusting in the whole cluster of languages then our beliefs and intentions are related in exactly the appropriate way to the corresponding language whose truth-conditions are given by the [standard] supervaluationist account. (144)

There is some object that is very much *like* a Lewisian language in that it is a function from words to intension-like things. For instance, this object is

²⁶Keefe misattributes the view to Burns (1991). As I pointed out above, Burns actual view is that we speak one of the many languages delimited by our conventions. See §1.1.

a function from words to functions from worlds to pairs of extensions and anti-extensions. This object plays all of the *theoretical roles* of a language. In particular, we obey a convention of truthfulness and trust with respect to this language-like object. So this object must qualify as a language! But then we are back to Williamson's earlier complaint: the T-schema does not apply to the sentence uttered in this unsharp language.²⁷

Reply. Suppose we decide to call such an object a language. Why think that we speak such a language? If we did speak such a language, we wouldn't assert the T-schema. We don't speak a language on which 'cat' means table, because we don't apply the predicate 'cat' to tables. Similarly, we don't speak a language on which truth is super-truth, because we make utterances like (3).

Objection (cont'd). But plainly there is *something* that is an 'unsharpened sentence'. Consider the individual's *utterance* as distinct from the many sentences uttered. The T-schema fails for utterances. Or, consider the uninterpreted sentence uttered. The T-schema fails for uninterpreted sentences (u-sentences).

Reply I. It's not clear that utterances are bearers of truth. Note the awkwardness of predicating truth to utterances when we are explicit about referring to the utterance rather than the sentence²⁸:

(8) # That speech act was true.

(9) # What he did in asserting that sentence was true.

There is even less reason to treat uninterpreted sentences as bearers of truth – at least as bearers of truth *simpliciter*, which I take to be what is predicated in instances of the T-schema. The sentence is supposed to be uninterpreted, so it doesn't have a meaning – it lacks intentionality. But such meaningless objects plainly cannot be true *simpliciter*. This is especially apparent when we notice that the same uninterpreted sentence can correspond to a true interpreted sentence and a false interpreted sentence. Again, ordinary folk talk of truth can

²⁷Thanks to Daniel Drucker for discussion here.

²⁸This example is from MACFARLANE (2014, §3.I.2)

be interpreted as being predicated of interpreted sentences rather than uninterpreted sentences.

Reply 2. Suppose we *do* decide to posit utterances and uninterpreted sentences as linguistic entities that are truth-apt. On this supposition, our practice of quotation-naming is ambiguous between naming utterances, uninterpreted sentences or interpreted sentences. I've shown that the T-schema holds for the interpreted-sentences and thus for one disambiguation of the truth predicate and quotation-name. That is sufficient to explain the assertibility of instances of the T-schema.

10.3 Did She Really Say That?

Objection. Mary observes Hsiao making the following utterance:

(10) Harry is bald.

According to your view, when Hsiao utters 'Harry is bald' she tokens many non-identical (counting by strict identity) sentences, each with precise truth conditions. And thus Hsiao asserts many propositions when she utters 'Harry is bald'. For example, she asserts that Harry has less than 10,000 hairs on his head (assuming this is one precisification of 'bald'). But Mary cannot say:

(11) # Hsiao said that Harry has less than 10,000 hairs on his head.

There are two replies to this objection. The first is motivated by pragmatic considerations. The second follows the theme of the rest of the paper by reflecting on lessons from Lewis's solution to the Problem of Many. A warning before proceeding: theorizing about speech reports is a messy business. Indeed, the messiness has led some (CAPPELEN AND LEPORE, 2005) to give up on simple, systematic theories of speech reports. My response, therefore, is sometimes forced to mirror some of this messiness.

Reply 1. We can accept the truth of this claim, while explaining the awkwardness of the speech report pragmatically. Consider scalar implicatures. Imagine a scenario in which I've seen exactly how many chairs are in the room. If I utter:

(12) There are 3 chairs in the room.

I implicate that there are *exactly* 3 chairs in the room. For otherwise, if there were exactly 5 chairs in the room, I would have made the strictly stronger assertion:

(13) There are 5 chairs in the room.

That follows from the Gricean maxim of maximal informativeness. Note just how strong the implicature is. Indeed, the implicature is so strong, meta-linguistic negation is warranted. If my friend walks into the room and observes exactly 4 chairs, it seems appropriate for him to utter:

(14) There aren't 3 chairs in the room – there are four.

In the case of Hsiao's assertion, the maxim of maximal informativeness is also violated. If Mary had uttered:

(15) Hsiao said that Harry is bald.

she would have asserted strictly more propositions. The fact that she chose to perform the less-informative speech act implicates that Hsiao didn't say that Harry is bald. This explains the awkwardness of asserting (11). And as with the chair example, the implicature may be strong enough to warrant meta-linguistic negation.

Indeed, if this pragmatic explanation of the unassertibility of (11) were true, we would expect the unassertibility of (11) to wane as the conversational purpose of the assertion (or the 'question under discussion') is made specific to a particular precise proposition, so that there is less reason to be maximally informative. And this prediction seems to be satisfied. For instance, suppose that we're in a context in which the term 'tall' definitely applies to a child if she is taller than 5'0" and definitely doesn't apply if the child is shorter than 4'7". (So, if a child is between 4'7" and 5'0", it is indeterminate whether she is tall, and

(by *Core Idea 2*) inappropriate to assert that she is tall.) A parent is planning a trip to an amusement park for which, let's say, one cannot enter unless they are at least 4'9". Arvind, speaking to the parent says:

(16) Regina is tall.

When asked whether Regina will be tall enough to ride the rides, the parent reports:

(17) Don't worry: Arvind said that Regina is taller than 4'9".

(17) strikes me as perfectly assertable. It is certainly much *less* awkward than if the conversational purpose of the assertion weren't so specific, which is exactly what the pragmatic story would predict. (Because vague words tend to be highly context sensitive, when imagining the scenario, it's important to keep at the forefront that the context picks out the range of meanings for 'tall' as specified above.) Indeed, the parent could follow up the report by canceling the implicature:

(18) I don't mean to say that's all Arvind said. He said more than just that: he said that she is tall.

Reply 2. My second reply continues the theme of drawing lessons from the Solution by Plenitude from the Problem of the Many, and explains our reluctance to accept claims like (11) using a psychological theory of counting. Recall the lesson **Almost-Identity**: when we say 'there is exactly one cat on the mat', we are counting by almost-identity rather than strict identity. Counting by almost-identity comes with a linguistic dimension and a psychological dimension.

According to the linguistic dimension, my utterance

(19) There is exactly one cat on the mat.

can be adequately paraphrased as (something like)

(20) Every cat on the mat massively overlaps each other.

either because (19) is true in virtue of expressing a proposition with the truth-conditions of (20), or (19) is strictly speaking false but the availability of the nearby proposition expressed by (20) blunts the incredulity of the falsehood.

There is also a psychological dimension to counting by almost-identity. Plainly, when we look at Tibbles on a mat, we are not simultaneously representing the many cats on the mat. Rather, we are only representing a single cat on the mat and are ignoring the distinctions between the many cats. That's why, when the differences between the many cats become relevant, we are initially struck with puzzlement. When the differences between the cats are made salient, we form a more fine-grained representation of the situation. Representing the situation fine-grainly, we realize that we've been ignoring the many cats, and have been counting by almost-identity all along, despite thinking that we were counting by strict identity. And, that fine-grained representation is short-lived – outside of philosophical contexts, we switch back to the coarse-grained representation of a single cat on the mat.

Now consider how this psychological story applies to the supersententialist. When Mary says "Hsiao said that Harry has less than 10,000 hairs on his head" it naturally strikes us as very odd. That's because we've been operating under a coarse-grained representation of the many propositions expressed, treating them as one. Because we've been representing all of these propositions as one, when the difference between the propositions are made salient by Mary's utterance, we are struck with puzzlement. It's only after philosophical reflection that we realize that we've been counting them by almost-identity all along.

10.4 Higher Order Vagueness

Objection. You've presented your view while ignoring higher-order vagueness. But the phenomenon of higher-order vagueness is real, and your view is incompatible with it. According to your view, when I make an utterance, I am in fact

uttering many sentences in many languages. And, an assertion is permissible just in case it expresses only true sentences. A borderline sentence is one that is simultaneously uttered with other sentences of differing truth values. But then consider the set of sentences S which are expressed by a given assertion. Either (1) each of the sentences are true, (2) each of the sentences are false, or (3) some are true and some are false. In the first and second cases, the sentence is not borderline. In the third case, the sentence is borderline. But then there is no room for higher-order vagueness: no scenario corresponds with *borderline* borderline sentences.

Reply. I will not attempt to directly show how my account can give an adequate account of higher-order vagueness. I've been arguing that supersententialism is the best form of supervaluationism. So, I will be content to show that higher-order vagueness poses no *unique* problem for my account, as compared to the standard supervaluationist. There is parallel complaint that can be leveled against the standard supervaluationist, and the same replies she gives to her analogous objection is available to me.

Compare the complaint above with the following analogous complaint directed against the standard supervaluationist:

Consider a sentence s and the set of admissible precisifications. Either (1) s is true on each of the admissible precisifications, (2) s is false on each of the admissible precisifications, or (3) s is true on some precisifications but false on other precisifications. In the first and second cases, the sentence is not borderline. In the third case, the sentence is borderline. But there is no room for higher-order vagueness: no scenario corresponds with *borderline* borderline sentences.

The standard supervaluationist has two responses. The first is to deny the phenomenon of higher-order vagueness; the supersententialist can, of course, also deny higher-order vagueness. The second response is to argue that the meta-language term 'admissible precisifications' is itself indeterminate.²⁹ Although there is one and exactly one set of admissible precisifications, it's indeterminate which set that is. In other words, which standard supervaluationist model is

²⁹See KEEFE (2000, §8) for this response.

the intended one is indeterminate. Thus, the standard supervaluationist can explain cases of borderline borderline sentences as cases in which the sentence is borderline on one model \mathcal{M}_1 and not borderline in \mathcal{M}_2 , where \mathcal{M}_2 and \mathcal{M}_1 are not determinately not the intended model. A parallel reply is open to the supersententialist. The supersententialist claims that the theorist’s term “languages being spoken” is itself indeterminate. Although there is one and exactly one set of languages being spoken, it’s indeterminate which set that is. In other words, which set of supervaluationist models are the intended set is indeterminate.

II Conclusion

The goal of this paper was to defend supersententialism as the most plausible form of supervaluationism. My strategy was to treat the Problem of Vagueness as an instance of the Problem of the Many and draw on the Solution by Plentitude for the Problem of the Many.

Let’s take a moment to compare the advantages of supersententialism with rival forms of supervaluationism. The supersententialists’ main advantage over the standard supervaluationist is that they retain the assertibility of all instance of the T-schema. According to the standard supervaluationist, some instances of the T-schema are untrue and therefore unassertible.³⁰ The supersententialist’s main advantage over rival forms of non-standard supervaluationist is that

³⁰There are several other advantages that the supersententialist, as a non-standard supervaluationist, has over the standard supervaluationist. First: it solves a puzzle for speech reports described by SCHIFFER (1998). (Although, see WEATHERSON (2003) for a solution on behalf of the standard supervaluationist.) Second: because truth isn’t supertruth, local validity is the correct understanding of truth-preservation and the classical rules of inference are preserved. Because the standard supervaluationist must treat truth-preservation as global validity, they lose several classical rules of inference. (C.f. VARZI (2007); although see WILLIAMS (2008) for a solution on behalf of the standard supervaluationist.) Third: non-standard supervaluationists have a natural response to the objection to supervaluationism offered by FODOR AND LEPORE (1996) according to which the precise precisifications of a vague language are simply irrelevant to discovering the meaning of the original vague language. The objection doesn’t touch the non-standard supervaluationists like the supersententialist, because they deny that there is a vague language distinct from the precise languages (c.f. VARZI (2007)).

it avoids the non-reductive objection. The supersententialist gives a fully reductive explanation of determinateness and vagueness which neatly explains the philosophical significance the notions have for our patterns of assertion.

Supersententialism, no doubt, makes some seemingly incredible claims. But all theories of vagueness make some seemingly incredible claims. Drawing on the Problem of the Many, I hope to have shown that these claims are not as incredible as they may have seemed. And, with these claims, I hope to have offered the most promising form of supervenience.

References

- Elizabeth Barnes and J.R.G. Williams. A theory of metaphysical indeterminacy. *Oxford Studies in Metaphysics*, 6, 2011.
- David Barnett. Is vagueness sui generis? *Australasian Journal of Philosophy*, 87 (1):5–34, 2009.
- Linda Claire Burns. *Vagueness: An Investigation into Natural Language and the Sorites Paradox*. Springer, 1991.
- Herman Cappelen and Ernie Lepore. *Insensitive Semantics*. Blackwell Publishing, 2005.
- Cian Dorr. Quantifier variance and the collapse theorems. *The Monist*, 97(4): 503–570, 2014.
- Cian Dorr and John Hawthorne. Semantic plasticity and attitude reports. *The Philosophical Review*, 123(3), 2014.
- Delia Graff Fara. Shifting sands: An interest relative theory of vagueness. *Philosophical Topics*, 28(1):45–81, 2000.
- Hartry Field. Disquotational truth and factually defective discourse. *The Philosophical Review*, 103:405–52, 1994. Reprinted in FIELD 2001, 222–258.

- Hartry Field. *Truth and the Absense of Truth*. Oxford University Press, 2001.
- Hartry Field. No fact of the matter. *Australasian Journal of Philosophy*, 81(4): 457–480, December 2003.
- Jerry A. Fodor and Ernest Lepore. What cannot be evaluated cannot be evaluated and it cannot be supervalued either. *The Journal of Philosophy*, 93(10): 516–535, October 1996.
- Nicholas Jones. Multiple constitution. *Oxford Studies in Metaphysics*, 9, 2015.
- Rosanna Keefe. Vagueness and language clusters. *Australasian Journal of Philosophy*, 76(4):611–620, 1998.
- Rosanna Keefe. *Theories of Vagueness*. Cambridge University Press, 2000.
- David Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
- David Lewis. General semantics. *Synthese*, 22:18–67, 1970. Reprinted with Postscript in LEWIS 1983, 189–232.
- David Lewis. Languages and language. *Minnesota Studies in the Philosophy of Science*, 7:3–35, 1975. Reprinted in LEWIS 1983, 163–188.
- David Lewis. *Philosophical Papers, Volume I*. Oxford University Press, Oxford, 1983.
- David Lewis. Many but almost one. In John Bacon Keith Campbell and Lloyd Reinhardt, editors, *Ontology, Causality, and Mind: Essays on the Philosophy of D.M. Armstrong*. Cambridge University Press, 1993. Reprinted in LEWIS 1999, 164–182.
- David Lewis. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 1999.

- E.J. Lowe. The problem of the many and the vagueness of constitution. *Analysis*, 55, 1995.
- John MacFarlane. *Assessment Sensitivity: Relative Truth and its Applications*. Clarendon Press, 2014.
- Vann McGee and Brian McLaughlin. Distinctions without a difference. *The Southern Journal of Philosophy*, 33, 1994.
- Trenton Merricks. Varieties of vagueness. *Philosophy and Phenomenological Research*, 62(1):145–157, January 2001.
- Stephen Read. Plural signification and the liar paradox. *Philosophical Studies*, 145:363–375, 2009.
- Stephen Schiffer. Two issues of vagueness. *The Monist*, 81, 1998.
- Nicholas J.J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, 2008.
- Jason Stanley. Context, interest relativity, and the sorites. *Analysis*, 63(4), October 2003.
- Michael Tye. Fuzzy realism and the problem of the many. *Philosophical Studies*, 81, 1996.
- Peter van Inwagen. *Material Beings*. Cornell University Press, 1990.
- Achille C. Varzi. Supervaluationism and its logics. *Mind*, 116(463):633–675, July 2007.
- Brian Weatherson. Many many problems. *Philosophical Quarterly*, 53(213), 2003.
- J.R.G. Williams. An argument for the many. *Proceedings of the Aristotelian Society*, 106, 2006.

- J.R.G. Williams. Supervaluationism and logical revisionism. *The Journal of Philosophy*, pages 192–212, 2008.
- Timothy Williamson. Definiteness and knowability. *The Southern Journal of Philosophy*, 33:171–191, 1994a.
- Timothy Williamson. *Vagueness*. Routledge, 1994b.
- Timothy Williamson. On the structure of higher-order vagueness. *Mind*, 108 (429):127–144, January 1999.

Chapter 2

VAGUE NATURALNESS AS ERSATZ METAPHYSICAL VAGUENESS[†]

There is a certain argument schema that is popular in contemporary metaphysics. It focuses on some term t and argues that the following claims are inconsistent:

(V1) t is vague

(V2) The property picked out by t is perfectly natural

(V3) There is no metaphysical vagueness

Call the above schema the Generalized Vagueness Argument Schema (GVAS). The GVAS is applied to different types of properties and to different dialectical ends. Authors including Dougherty (2013), Dorr and Hawthorne (2014), Sider (2009b; 2011), Williams (2008a), and, arguably, Barnes (2014) and Parfit (1984) apply the schema in debates over ethical properties, personal identity,

[†]Thanks to Daniel Drucker, David Manley, Eric Swanson, JRG Williams, and two anonymous referees. Thanks also to audiences at Michigan for helpful questions. Special thanks to Ted Sider for several early meetings and comments on this essay, and Brian Weatherson for extensive feedback throughout the development of this essay.

consciousness, special science properties, unrestricted mereological composition, four dimensionalism, ontological realism, and beyond.¹ Sometimes the schema is used to argue that there is metaphysical vagueness (from (V₁) and (V₂)). Other times, it's deployed in order to establish that a property can't be perfectly natural (from (V₁) and (V₃)) or that a term isn't vague (from (V₂) and (V₃)). And sometimes it's used in a reductio against a position that is supposedly committed to all three theses.

The GVAS, if it works, ends up having significant upshots for our theorizing. That's because often each of the three theses (V₁)-(V₃) looks fairly compelling. In the case of ethical properties, for example, we might have good reason to acknowledge that our moral terms are vague (V₁) while adopting robust moral realism, understood in terms of perfect naturalness (V₂) – all without having to accept spooky metaphysical vagueness (V₃).

The prevalence of the GVAS has not been fully appreciated. That's partly because authors present the GVAS with varying degrees of explication. On its more explicit presentations, proponents of the GVAS gesture at the role perfect naturalness plays in resolving semantic indeterminacy – a meta-semantic program that has come to be known as reference magnetism. Often, however, the details of the argument from reference magnetism to the GVAS are kept implicit.

The first task of this paper is to document the prevalence of the GVAS and offer a detailed presentation of the reference magnetic argument for the incompatibility of (V₁)-(V₃). This presentation alone is significant. It unifies a set of seemingly disparate arguments under one hidden schema and develops the strongest version of that schema, which can then be extended to instances where exposition is lacking. With a clear presentation in hand, we are also well positioned to chart out the various possible strategies for resisting the GVAS.

The second task of this paper is to develop one such strategy that I find par-

¹Neither Parfit nor Barnes make explicit appeal to Lewisian naturalness but rather to fundamentality and irreducibility. As I discuss, their arguments are incomplete; understanding them as instances of the GVAS allows us to complete them.

ticularly fruitful. I propose that those attracted to (V₁)-(V₃) adopt a position that I call *ersatz metaphysical vagueness*. According to the ersatzer, the term ‘perfectly natural’ is vague and the vagueness is not metaphysical. I’ll argue that adopting ersatz metaphysical vagueness allows us to mimic genuine metaphysical indeterminacy, giving us a highly general strategy for resisting GVAS arguments. By taking the term ‘perfectly natural’ in (V₂) to be non-metaphysically vague, (V₁)-(V₃) can be made consistent. Importantly, we can continue to treat the vague term as picking out a perfectly natural, reference magnetic property.

This will be a welcome result for those who are attracted to (V₁)-(V₃) in a given domain. Moreover, because GVAS arguments that go from (V₁) and (V₂) to the negation of (V₃) are a popular way to motivate metaphysical vagueness, ersatz metaphysical vagueness gives us a non-objectionable alternative to genuine metaphysical vagueness.²

I The GVAS

The GVAS wouldn’t be of much dialectical interest if there were little initial plausibility to one or other of (V₁)-(V₃). But, for a range of properties, each of the theses do have *some* independent plausibility – none are *obviously* false. Thus, the GVAS, if successful, would have substantial theoretical upshots.

Consider (V₃), the claim that there is no metaphysical vagueness. Metaphysical indeterminacy is often stated as the view that the “world itself, independently of what we know about it or how we represent it, is metaphysically indeterminate” (WILLIAMS, 2008b, 1). Giving a theory of metaphysical vagueness – a theory of how the ‘world itself’ can be vague – is notoriously challenging. Fortunately, for our purposes, we need not dive into the nature of

²Throughout the paper, I will be a bit sloppy in my use of the terms ‘indeterminate’ and ‘vague’, treating them as synonymous. Strictly speaking, vagueness might require more than indeterminacy. However, the GVAS seeks to show that terms that pick out perfectly natural properties cannot be non-metaphysically indeterminate (and therefore cannot be non-metaphysically vague). So the focus of the paper, in the first instance, is indeterminacy not vagueness.

metaphysical vagueness. Following BARNES (2010), it suffices to offer a *negative characterization* of metaphysical vagueness according to which some sentence is indeterminate and the source of the indeterminacy is neither semantic nor epistemic. More carefully³:

MV A sentence *S* is metaphysically vague iff, were all its representational content precisified, there is an admissible precisification of *S* such that according to that precisification the sentence would still be non-epistemically indeterminate.⁴

How should we assess the plausibility of a metaphysically vague sentence (as defined by (MV))? Some have complained that metaphysical vagueness is incoherent or unintelligible.⁵ Coherence and intelligibility aside, I can think of at least two reasons for doubting the claim that there is metaphysical indeterminacy. The first reason: the view strikes many (myself included) as incredible. Although this may be good reason to reject metaphysical indeterminacy, incredulity arguments are, of course, of limited persuasiveness. The second reason: metaphysical indeterminacy requires positing irreducible ideology that is objectionable. The specifics of this worry, of course, depend on the particular theory of metaphysical vagueness under examination. But the most plausible developments do seem to succumb to this objection. For instance, according to the development in BARNES AND WILLIAMS (2011), indeterminacy is taken as an irreducible piece of ideology: there are irreducible facts that involve the indeterminacy operator. And the indeterminacy operator just seems like the wrong sort of irreducible ideology. Just as many balk at the idea that the world is irreducibly *modal*, there is something deeply off-putting about the idea that the world is irreducibly indeterminate.

³In Chapter 3 of this dissertation, I argue that (MV) is a flawed extension of the *exclusion test* for metaphysical vagueness. Because offering an analysis of metaphysical vagueness would take me far afield from the goals of this paper, I set aside those complaints here and treat (MV) as a proxy for metaphysical vagueness. (Note that the exclusion test suffices to establish premise (v) of the argument below.)

⁴C.f. BARNES (2010, 604)

⁵See BARNES AND WILLIAMS (2011) and WILSON (2013) for a reply.

Of course, the literature on metaphysical vagueness is large and camps are entrenched – no doubt some find metaphysical vagueness neither incredible nor objectionable. But even fans of metaphysical vagueness should admit that the position isn't *obviously* true and thus have reason to be interested in the success of the GVAS. Moreover, (V₃) can be refined to claim that the *particular* terms targeted by the GVAS (e.g. 'wrong', 'person', 'conscious', etc.) are not metaphysically vague. Anyone who doesn't take this more specific claim to be obvious should take an interest in the GVAS.

The reasons to prefer (V₂) vary from dialectic to dialectic. But the sorts of properties that the GVAS is typically directed against include normative properties, consciousness, special science properties, personal identity, and existence. And while there is no general argument that purports to establish the perfect naturalness of all of these properties, many theorists have found domain-specific reason to accept (V₂). For instance, Dougherty (2013) and Schoenfield (2016) are sympathetic to moral realism, which they understand as the claim that moral properties are "part of the deep metaphysical structure of the world," where this is glossed in terms of Lewisian naturalness. ZIMMERMAN (2010) understands dualist positions in the philosophy of mind in terms of perfectly natural qualia properties. SIDER (2001b, 2009a,b, 2011) finds good reason to accept ontological realism, which he understands as the claim that the quantifier is perfectly natural (or in his terminology 'structural'). EKLUND (2004, 492) considers the thesis that *being a person* is a natural property. And WILLIAMS (2008a) understands 'metaphysically primitive' special science properties in terms of perfect naturalness.

Turn now to (V₁): the claim that the terms for these properties are vague. There are two main avenues for generating the plausibility of (V₁) for predicates like 'is permissible', 'is a person', 'is conscious', etc.

The first avenue is loosely inspired by the vagueness argument for composition given in LEWIS (1986, 212). Call the property that is supposed to be the referent of the term *t* in the schema the *target property*. So, if the term targeted

by the GVAS is ‘is a person’, then the target property is *personhood*. We’ll say that the property named by t is vague just in case there is some object such that it is indeterminate whether that object instantiates the property named by t . So, we’ll say that *personhood* is vague just in case there is some object such that it’s vague whether that object instantiates *personhood*. Crucially, *personhood* may be vague without the vagueness being metaphysical. For all I’ve said, the indeterminacy might be entirely semantic; it may be that we haven’t settled which precise property is being picked out by the term ‘*personhood*’.

For each target property, there is supposed to be some intuitive desideratum for the property’s instantiation, such that the target property is instantiated iff some *base property* is instantiated. It’s then noted that any reasonable way of spelling out the base property will be vague: there will be cases in which it is indeterminate whether the base property is satisfied. A more precise specification of the base property is supposed to be absurd. There is thus pressure to conclude that the target property is indeterminate (VI).⁶ Examples of target properties and their corresponding base properties include:⁷

- Mereological composition and life, consciousness, sufficient stuck-togetherness, etc.

⁶We can firm up the argument a bit by considering its form: $\{\phi \leftrightarrow \psi, \nabla(\phi)\} \vdash \nabla(\psi)$ where ‘ $\nabla(\phi)$ ’ stands for ‘indeterminately ϕ ’ or equivalently $(\neg \Delta(\phi)) \wedge (\neg \Delta(\neg\phi))$ where $\Delta(\phi)$ stands for ‘determinately ϕ ’. The formulas ϕ and ψ will assert the existence of an instance of the base property and target property respectively. Thus, they will be of the form: $\exists xFx$. The conclusion of the argument will be $\nabla(\exists xFx)$ where F is the target property. Assuming that the quantifier is not vague, the vagueness looks to be in the target term F . In other words: $\nabla(\exists xFx) \vdash \exists x\nabla(Fx)$. Focusing on supervaluationism and epistemicism, the question as to whether or not the inference $\{\phi \leftrightarrow \psi, \nabla(\phi)\} \vdash \nabla(\psi)$ is valid will depend on whether or not the correct logic of vagueness validates $p \vdash \Delta(p)$. On simple supervaluationist models, of the kinds discussed by WILLIAMSON (1994), this inference holds. However, this inference also forces the supervaluationist to give up some classical rules of inference (contraposition, conditional proof, argument by cases, and *reductio ad absurdum*). Theories for which $p \not\vdash \Delta(p)$, including epistemicism and more sophisticated forms of supervaluationism (like those discussed in FIELD (1994); MCGEE AND MCLAUGHLIN (1994); WILLIAMS (2008c)) will require the stronger premise that the intuitive desideratum *determinately* holds: $\Delta(\phi \leftrightarrow \psi)$. It’s not clear to me how severely this stronger requirement weakens the argument (although I think this is a question worth pursuing). See also NOLAN (2006).

⁷C.f. PARFIT (1984, 231-243), DOUGHERTY (2013), and GOFF (2014, 184-187).

- Normative properties and normatively relevant descriptive properties
- Personhood and psychological or bodily continuity
- Consciousness and patterns of neural firings
- Special science properties and microphysical properties

The second avenue is closely related to the first. Instead of focusing on a biconditional that relates the instantiation of the base property and the target property, we use a range of properties of the same type as the base property to construct a sorites sequence for the target property. Consider, for instance, the GVAS as applied to normative properties. We tend to think that the instantiation of normative properties comes with some descriptive base property. As Miriam Schoenfield points out, this allows us to construct a sorites sequence:

Darryl is watching his two year old daughter play in a city park. It is permissible to divert his attention from her for 1 second. it is not permissible to divert his attention from her for 5 minutes. Is it permissible to divert his attention for 30 seconds? 31? 32? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of diversions whose lengths differ by a second. (262)

The Darryl case draws our attention to a sort of tolerance principle that seems to hold for the normative term ‘permissible’ which undergirds a sorites premise: if looking away for n seconds is permissible, then looking for $n + 1$ seconds is permissible. And similar cases have been constructed for personal identity (PARFIT, 1984, 239), composition (SIDER, 2001b, 124), and consciousness (GOFF, 2014, 184-187).

2 The Reference Magnetic Interpretation of the GVAS

Taking the term picking out the target property to be vague, however, creates some tension with the other two theses. Different proponents of the GVAS spell out this tension with varying degrees of explication. Consider, for instance, an

appeal to the GVAS by WILLIAMS (2008a), in which he argues from (V1) and (V2) to the negation of (V3).⁸

Intuitively, it can be a vague matter at any given moment whether some cluster of particles is uranium...so for an *a* such that it is vague whether *a* is uranium, is this a matter of semantic vagueness?...if *being uranium* is primitive, this line of thought is blocked....There seems no room to wriggle out of this by appeal to semantic phenomena. What goes for uranium, of course, goes also for other kinds in which we might take a metaphysical interest: *being a person*, *being water* or *being conscious*....If *uranium*, or *water*, or *consciousness*, or *being a person*, are among the perfectly natural properties, in Lewis' sense, then to avoid ontic vagueness one would have to insist, counter-intuitively, that there are sharp boundaries to the application of such properties....Such 'metaphysical brutalism' might still be consistent with an epistemic account of the vagueness...But unless brutalism is an option, there seems no way in appropriate cases to avoid regarding the vagueness of whether or not something is uranium as an instance of ontic vagueness. (2008a, 141-142)

Williams is plainly appealing to the GVAS. And although several crucial steps in his argument are left implicit, we can use his discussion as a starting point for explicating the tension between (V1)-(V3). To show that (V1)-(V3) are inconsistent, it suffices to construct an argument from two of the claims to the negation of the third. Williams takes (V1) and (V2) as premises and shows that (V3) is false. That is, he takes as premises that the predicate 'uranium' is vague and that *being uranium* is perfectly natural, and argues from these claims that there must be metaphysical vagueness. More generally he takes as premises:

- (i) The predicate '*F*' is vague
- (ii) The property picked out by '*F*' is perfectly natural

He then argues that the vagueness is neither semantic nor epistemic, concluding that the vagueness must be metaphysical:

- (iii) If '*F*' picks out a perfectly natural property, '*F*' is not semantically vague

⁸While Williams is pointing to the tension between (V1)-(V3) as a way to motivate metaphysical indeterminacy, he officially remains neutral as to which of (V1)-(V3) to reject.

- (iv) If ‘*F*’ picks out a perfectly natural property, ‘*F*’ is not epistemically vague
- (v) If ‘*F*’ is neither semantically vague nor epistemically vague, then there is a sentence that is metaphysically vague
- (vi) Therefore: there is a sentence that is metaphysically vague.

The argument is clearly valid and we’re assuming premises (i) and (ii). Premise (v) follows from (MV). If ‘*F*’ is vague, there plausibly is some sentence *S* that includes ‘*F*’ and inherits the vagueness of ‘*F*’. For example, there will be some vague sentence of the form ‘*Fa*’ where ‘*a*’ is not vague. If ‘*F*’ is neither semantically vague nor epistemically vague, then removing the semantic and epistemic vagueness won’t remove the vagueness of *S* and, by (MV), we have metaphysical vagueness. The controversial premises, then, are (iii) and (iv).

2.1 Naturalness and Semantic Vagueness

Why does Williams think that regarding *being uranium* as a vague perfectly natural property leaves ‘no room to wriggle out of this by appeal to semantic phenomena’. Why can’t the vagueness in ‘*F*’ be semantic?

While Williams kept the details of this step in his argument implicit, other philosophers who have appealed to the GVAS have been more explicit about the difficulty of treating ‘*F*’ as semantically vague. When this step is made explicit, proponents of GVAS appeal to a meta-semantic theory known as *reference magnetism*. Consider, for example, Dougherty’s (2013) discussion of the GVAS:

...a natural kind can act as a ‘reference magnet’ for a term. Let us suppose for now that there is a unique set of things that constitutes the natural kind, orangutan...Now, assuming we reject a metaphysical view of vagueness, this natural kind has a precise boundary: everything is in this set or it is not. Therefore, our use of the word ‘orangutan’ would pick out a precise set of things. (6-7) [...] Robust realism can be understood as the view that ethical kinds have the same metaphysical status as natural kinds: they form part of the metaphysical structure of the world, which obtains entirely independently of humans’ thoughts and

practice. (9)⁹

Sider is also explicit about the role reference magnetism plays in his use of the GVAS. He argues that ‘existence’ cannot be vague – a claim that is used to great effect in the vagueness arguments for unrestricted composition and four dimensionalism, and in order to respond to ontological deflationists that take ontological questions to be indeterminate. Sider (2009b, 410) writes: “[The thesis of] Indeterminacy says that ‘there exists’ in English is semantically indeterminate over various candidates; but if [perfectly natural] existence is a reference magnet, then ‘there exists’ determinately means [perfectly natural] existence.”^{10,11}

Each of these authors appeal to a meta-semantic view dubbed ‘reference magnetism’. As Weatherson (2003) describes the view:

Sometimes our words refer to a particular property or object rather than another, not because our dispositions make this so, but because of some particular feature of that property or object. David Lewis calls this extra feature ‘naturalness’: some properties and objects are more natural than others, and when our verbal dispositions do not discriminate between different possible contents, naturalness steps in to finish the job, and the more natural property or object becomes the content. (484)

Our linguistic dispositions only go so far in establishing the reference of our words; the naturalness of the possible referents of our terms also plays a role in determining reference. In particular, if there are a range of properties p_1, \dots, p_n that are consistent with our use of predicate ‘ F ’, then if one of those properties

⁹Two notes are important. First: Dougherty doesn’t distinguish between perfectly natural kinds and relatively natural kinds. Second: Dougherty suggests several ways in which to avoid the GVAS – for instance by giving up the claim that there is a unique natural kind or adopting epistemicism; I discuss these avenues of escape below. What’s important here is that he seems to think that if there is a single natural kind, then terms picking out that natural kind won’t be semantically vague.

¹⁰See also SIDER (2011, §5).

¹¹While Sider formulates his argument using a proprietary structure operator (\mathcal{S}) that attaches directly to the quantifier, we can follow HEIM AND KRATZER (1998) and DORR (2005) in treating existential quantifiers as second-order properties, allowing us to treat Sider’s appeal as another instance of the GVAS.

is more natural than the rest, ‘*F*’ refers to that property. In this way, naturalness can resolve semantic indeterminacy left over from our linguistic dispositions. But, as Weatherson points out, sometimes naturalness fails to resolve the semantic indeterminacy:

Well that is what happens when things go well. Vagueness happens when things do not go well. Sometimes our verbal dispositions do not discriminate between several different contents, and no one of these is more natural than all the rest. In these cases there will be many unnatural contents not eliminated by our dispositions which naturalness does manage to eliminate, but there will still be many contents left uneliminated....Hence there is no precise fact about what the phrase denotes. Hence it is vague. (484)

Semantic indeterminacy results when neither our linguistic dispositions nor naturalness privilege one of a range of candidate referents for a term. Suppose *being uranium* is perfectly natural and ‘uranium’ is vague. Assume for *reductio* that the vagueness is semantic. That would mean that there is range of properties p_1, \dots, p_n that are consistent with our usage of the term ‘uranium’, none more natural than the rest. But at least one of the properties would have to be *perfectly* natural. (After all, ‘uranium’ refers to one of those properties, so one of those properties is *being uranium*, which by hypothesis is perfectly natural.) If *only* one of the properties were perfectly natural, then one of the properties would be more natural than the rest and there wouldn’t be any semantic vagueness.

There appears to be only one way to posit semantic vagueness in ‘uranium’: treat all of the properties p_1, \dots, p_n in the range as perfectly natural.¹² But I doubt we should ‘wriggle out’ of the GVAS in this way. Such an abundance of perfectly natural properties is an over-abundance: it violates considerations of ideological parsimony.

¹²DOUGHERTY (2013) discusses this possibility with vague moral properties. He does not distinguish between perfectly natural kinds and relatively natural kinds and my criticism here only applies to those that take the term in the GVAS to pick out *perfectly* natural properties – positing a range of relatively natural kinds does not violate considerations of parsimony. Plausibly, though, a thorough-going moral realist (of a Moorean stripe) does take moral properties to be perfectly natural. As Dougherty himself describes the view he calls ‘robust moral realism’: “ethical properties are part of the deep underlying metaphysical structure of the world” (7).

Theorists disagree over how parsimony considerations are best cashed out in terms of natural properties. As Lewis originally thought of perfectly natural properties, the perfectly natural properties are supposed to form a minimal supervenience basis. On this more traditional understanding, parsimony considerations tell against hypotheses that include modally redundant perfectly natural properties. According to more modern developments of the notion (c.f. SIDER (2011)), the perfectly natural properties are supposed to form a minimal *reductive* basis; i.e. it is thought that all facts reduce to facts involving only perfectly natural properties, where different theorists disagree over the relevant notion of reduction. On this more contemporary understanding, parsimony considerations tell against hypotheses that include perfectly natural properties that are explanatorily idle – where the form of explanation is supposed to be distinctively metaphysical.

On either understanding of parsimony, it is highly implausible that there exist a range of perfectly natural properties that serve as referents for terms like ‘person’, ‘wrong’, ‘conscious’, or ‘uranium’. The first understanding of parsimony simply appears to be a non-starter for claiming that there is even a single perfectly natural property that serves as a candidate referent for our term, at least for properties like *being uranium*, *being a person*, or *being wrong*, which seem plainly to supervene on microphysical properties. This supervenience would create inexplicable, brute connections between the perfectly natural property and the microphysical.

Even if we were comfortable with this brute connection, or we allowed duplicate microphysical worlds in which the putative natural property is not instantiated (which some may want to do, for instance, in the case of consciousness), it is even *more* difficult to maintain that there are a *range* of such properties. Consider the case of personal identity. Suppose we have a situation in which there are a range of properties – *personhood*₁, *personhood*₂, *personhood*₃... – each of which is perfectly natural. On the more traditional supervenience understanding of natural properties, these properties would have to be modally

independent. That would require there to be pairs of worlds with the same microphysical structure and with the same instantiation pattern of *personhood*₂, *personhood*₃, *personhood*₄, etc. that differed with respect to whether *personhood*₁ were instantiated. The result would do significant violence to our ordinary intuitions of what is possible. But the cost of salvaging our modal intuitions is to posit an *additional* set of inexplicable, brute connections between the natural properties.

The situation is no better on a reductive understanding of parsimony. Perhaps facts about wrongness, consciousness, etc. cannot be reduced to a microphysical basis, thus forcing us to accept such properties as perfectly natural. But there are no facts that require a range of natural properties in order to be explained – at least no facts that don't involve the indeterminacy operator.

If positing a range of perfectly natural properties is out of the question, and there is only one perfectly natural property – *being uranium* – in the range of properties that are consistent with our use of the term 'uranium', then that property will attract the reference of our terms, closing off room for semantic vagueness.

2.2 Naturalness and Epistemicism

The last option for avoiding the GVAS is to claim that the vague predicate '*F*' is *epistemically* indeterminate. Williams hints at a strategy for rejecting this option in order to secure premise (iv): epistemicism requires a sort of objectionable 'metaphysical brutalism'. The epistemicist must admit that a single atom can make the difference between consciousness and lack of consciousness or the difference between being a person and not being a person. She will accept that there are two scenarios that are incredibly alike, save the location of a single atom, in which the atoms form a person in the first case but not in the second. But that seems incredible. Sider puts the worry best in his discussion of composition:¹³

¹³Similar complaints of brutalism also appear in DOUGHERTY (2013).

To postulate such a sharp cut-off would be to admit that the realm of the macroscopic is in some sense ‘autonomous’ of the microscopic....there would seem to be something ‘metaphysically arbitrary’ about a sharp cut-off in a continuous series of cases of composition. Why is the cut-off here, rather than there? SIDER (2001b, 124)

The very same considerations that are in play when deciding whether or not to accept brutalism for composition are in play when deciding whether or not to accept brutalism for properties like personhood, etc., *assuming that those properties are perfectly natural*. That is, it would be metaphysically arbitrary for a fundamental macro-property like personhood or permissibility to have such a fickle relationship with the microphysical. It’s no wonder then, that the very same dismissal of brutalism is found with respect to permissibility (SCHOENFIELD, 2016), consciousness (GOFF, 2014, 184-187), and personal identity (PARFIT, 1984, 239)¹⁴

Note that the sharp-cut offs required for the epistemicist to avoid the GVAS are different from the typical cutoffs they posit for predicates like ‘bald’. The brutalism typically associated with epistemicism is one in which there are many equally eligible candidate meanings for our terms and it’s an arbitrary fact that our usage privilege one of these candidate meanings. But a brute cutoff in the application of a fundamental property requires a metaphysical arbitrariness, according to which one of the candidate meanings is more natural than the rest.¹⁵ It is this metaphysical arbitrariness that Sider and others find objectionable.¹⁶

Once this form of the GVAS is noticed, it is easy to replicate that argument in various domains. Consider Parfit’s argument for reductionism about personal identity. Parfit constructs a sorites sequence for personal identity using

¹⁴Other metaphysicians have found such metaphysically brute cutoffs unobjectionable: see MARKOSIAN (1998) and NOLAN (2006).

¹⁵C.f. SIDER (2001b, 131-132)

¹⁶*Even if* we were somehow able to accept the metaphysical arbitrariness required to posit a sharp cut-off, the standard epistemicist explanation for our ignorance in cases of vagueness, due to WILLIAMSON (1994), cannot be extended to vagueness in terms targeted by the GVAS. See WASSERMAN (2012) for brief discussion.

the base properties of psychological or bodily continuity and notes that this puts pressure on us to accept that personal identity can be vague (V1). He is focused on semantic conceptions of vagueness (V3). And he takes these two claims to lend support for a *Reductionist View* of personal identity, according to which personal identity is not a fundamental feature of reality (rejection of (V2)). What is supposed to be the source of the tension? Parfit writes:

It is hard to believe (1)...that, in one of these [intuitively borderline] cases, the resulting person would quite straightforwardly be me, and that, in the next case, he would quite straightforwardly be someone else. It is also hard to believe (2) that there must be a such a sharp borderline, somewhere in the Spectrum, though we could never have any evidence where the borderline would be...claims (1) and (2), taken together, are extremely implausible. I believe that they are even more implausible than the only other possible conclusion, which is the Reductionist View (1984, 239)

Yet, no argument explaining why the Reductionist View is the only other possible conclusion is given. In other words: why can't one accept Parfit's argument for the semantic vagueness of personal identity yet retain the view that personal identity is fundamental?¹⁷ The reference magnetic GVAS gives us a way of filling in the argument. We can understand the Reductionist View, on which personal identity is not fundamental, as the claim that *being a person* is not perfectly natural.¹⁸ On this understanding, the reference magnetic GVAS

¹⁷Here's the *closest* argument that I could construct from the text: Parfit seems to think that the only feasible non-reductionist view is one in which there is a separate existing object, like a soul, distinct from the body or brain (240). Perhaps he thinks that existence cannot be vague as a matter of semantic vagueness. Therefore, a non-reductionist view will yield determinate answers to intuitively borderline cases. If this is Parfit's argument, it is certainly never made explicit. Two further points are relevant. First: that argument merely moves the bump in the rug. The best argument against semantically vague existence is due to Sider and is itself an instance of GVAS (discussed in the main text above). Second: that argument relies on claiming that the only feasible non-reductionist position is one that posits a separate existing entity. But Parfit gives no argument for the infeasibility of a non-reductionist view that doesn't posit a separate existing entity (the view he calls *the Further Fact View* (210)).

¹⁸See SIDER (2011) for a guide to understanding fundamentality and reduction in terms of naturalness (or in his jargon, structure) (although see BENNETT (Forthcoming) for a broader discussion of ways of understanding fundamentality.) EKLUND (2004, 492) plays with the idea that being a person picks out a natural kind, and notes that a reference-magnetic metasemantics may then force determinacy in the term 'person' despite our facts about our use under-determining

justifies Parfit in taking the non-metaphysical vagueness in personal identity to support a reductionist view: on a non-reductionist view, personhood would be a reference magnet which prevents vagueness in claims of personal identity.

Similar remarks apply to Barnes (2014), who argues that there is vagueness in various fundamental properties (specifically morality and existence) and takes this to imply metaphysical vagueness. It's not clear, however, why the vagueness for a fundamental property cannot be semantic or epistemic rather than metaphysical. The reference-magnetic GVAS also gives us a way to fill in her argument for this conclusion.

3 Vague Naturalness as Ersatz Metaphysical Vagueness

With the GVAS fully documented and interpreted, several strategies for resisting it become immediately obvious. We can, for instance, give up a reference magnetic meta-semantics. I, however, want to explore a more subtle strategy for resisting the GVAS – one that doesn't give up on the reference magnetic meta-semantics many find promising.

That strategy takes the term 'perfectly natural' to be a predicate of properties¹⁹, but treats the term as vague, where the vagueness is not metaphysical.²⁰ In the material mode, the strategy claims that for some particular property, it's

meaning. (Eklund does not appeal to the GVAS; he merely claims that the naturalness of personal identity *may* lead to determinacy, not that it is inconsistent with indeterminacy.) See also MILLER (2010) for an argument that persons are a *sui generis* ontological kind. Although Miller is not claiming that *being a person* is perfectly natural, I think her arguments provide the materials to support such a claim. See SIDER (2001a) for a denial of the claim that *being a person* is a perfectly natural property.

¹⁹LEWIS (1983a) discusses this option for those that don't wish to posit universals in their ontology.

²⁰The position is thus importantly different from the one advocated in CAMERON (2010), according to which *relative* naturalness facts are *ontically* vague and they determinate a *strict* ordering of properties in terms of naturalness. I, on the other hand, am claiming that the *perfect* naturalness facts are *semantically* vague, even if they only determine a *weak* ordering of properties in terms of naturalness. (Cameron does not discuss the GVAS argument schema.)

indeterminate whether that property is perfectly natural. I call this position *ersatz metaphysical vagueness* because of its ability to mimic genuine metaphysical vagueness. I'll show that the mv-ersatzer can resolve the GVAS in three steps.

Step 1. I'll argue that the mv-ersatzer can acknowledge that terms like 'person', 'permissible', 'uranium', 'conscious', etc. are vague while declaring that personhood, permissibility, being uranium, consciousness, etc. are perfectly natural.

Step 2. I'll argue that taking 'natural' to be vague doesn't require metaphysical vagueness.

Step 3. I'll argue that taking terms like 'person', 'permissible', 'uranium', 'conscious', to be vague, while simultaneously picking out perfectly natural properties, doesn't require metaphysical naturalness.

The first step establishes (V₁)-(V₂). The last two steps establish that (V₃) is preserved.

4 Step One: Vague Natural Properties

The first step of my proposal is to argue that, by taking 'perfectly natural' to be vague, the mv-ersatzer can jointly endorse (V₁) and (V₂): they can acknowledge that terms like 'person', 'permissible', 'uranium', 'conscious', etc. are vague while simultaneously declaring that *personhood*, *permissibility*, *being uranium*, *consciousness*, etc. are perfectly natural.

Suppose *personhood* is perfectly natural. The predicate 'is a person' is vague: it's vague which extension it refers to. Thus, the singular term '*personhood*' is vague: there are several candidate properties p_1, p_2, p_3, \dots and it's vague which is the referent of the singular term. According to the mv-ersatzer, the predicate of properties 'is perfectly natural' is also vague. It includes only one of $p_1, p_2,$

p_3, \dots in its extension, but it's vague which. But, if there is a penumbral connection between 'personhood' and 'is perfectly natural' such that the referent of the term '*personhood*' is in the extension of 'is perfectly natural', then it's determinately the case that the property *personhood* is perfectly natural. We would expect exactly such a penumbral connection if, for example, it was part of the meaning of '*personhood*' that it pick out a metaphysically significant, perfectly natural property.

We can explicate this penumbral connection in precisificational terms. There are several properties p_1, p_2, p_3, \dots that are candidate referents for the term '*personhood*' and several candidate extensions for the predicate 'is perfectly natural'. On some precisifications, '*personhood*' refers to p_1 , on others it refers to p_2, p_3, \dots . On some precisifications, the extension of 'is perfectly natural' includes p_1 and not p_2, p_3, \dots . On others it includes exactly one of p_2, p_3, \dots . But on every precisification, the referent of '*personhood*' on that precisification is in the extension of 'is perfectly natural' on that precisification.

So, although it's vague what '*personhood*' refers to (it's unsettled which property is the property of *personhood*) and it's vague what extension 'perfectly natural' refers to (it's unsettled which properties are perfectly natural), it's not vague that the referent of '*personhood*' is in the extension of 'is perfectly natural' (it's settled that the property *personhood* is perfectly natural). And what goes for 'person' goes for other predicates targeted by the GVAS. Thus, the mv-ersatz can retain (V₁) and (V₂).

I argued that the mv-ersatz can accommodate (V₁)-(V₂). It remains to show that the mv-ersatz can retain (V₃) – to show that mv-ersatzism doesn't require metaphysical vagueness. I do this in the next two sections. First, I argue that the vagueness in 'natural' is not metaphysical. Then, I show that the vagueness in the term targeted by the GVAS ('person', 'uranium', etc.) is not metaphysical – despite the fact that the term picks out a perfectly natural property.

5 Step Two: Vague Naturalness is not Metaphysical Vagueness

It may be thought that in taking perfect naturalness as vague, the mv-ersatzer is positing a sort of metaphysical vagueness. Many metaphysicians (most prominently Sider (2011)) have understood claims of fundamentality in terms of naturalness: a property is fundamental iff it's perfectly natural. On a naturalness-based conception of fundamentality, the mv-ersatzer seems to be claiming that the fundamental is indeterminate. Is that not a sort of metaphysical vagueness?

It is not. I'll spot my opponent a natural-based conception of fundamentality. Even with this concession, the thought that vague perfect naturalness is a sort of metaphysical vagueness arises from a confusion between two theses. Consider the thesis that indeterminacy is fundamental:

Indeterminacy is Fundamental (IiF) The indeterminacy operator is fundamental.

Admittedly (IiF) is stated in obscure terms. But, let's simply reflect on (IiF) in its present gloss for now. In particular let's compare (IiF) with the claim that it's indeterminate what is fundamental:

Fundamentality is Indeterminate (FiI) For some property p , it's indeterminate whether p is fundamental

Distinguish (IiF) and (FiI). (IiF) says that facts about indeterminacy are fundamental; (FiI) says that it is indeterminate what counts as fundamental. The difference is significant: (IiF) is problematic whereas (FiI) is not.

From the looks of it, (IiF) entails (MV): if indeterminacy is fundamental or 'irreducible', then its source is not linguistic or epistemic. (IiF), then, inherits the implausibility of (MV). Indeed, one of the reasons we found (MV) so implausible was because its best developments adopt a thesis like (IiF) according to

which the indeterminacy operator is taken as an irreducible piece of ideology. But, like those (LEWIS, 1986) that find it strange to think of the world as irreducibly modal, it's strange to think of the world as irreducibly indeterminate. Indeed, both Barnes and Williams (2011, fn. 4) and Sider (2011, §7.12) take (IiF) as a sufficient condition for metaphysical vagueness.

What about (FiI)? Some writers seem to conflate (FiI) with (IiF). For instance, Sider (2011) claims that the “fundamental is determinate” and proposes that we understand this claim as requiring the rejection of (IiF), suggesting that (IiF) is a sufficient condition for metaphysical vagueness, which he rules out. But, in later arguments, he conflates (FiI) with (IiF), appealing to the rejection of the former rather than the later.²¹

Other authors distinguish (FiI) from (IiF), but take it as obvious that both are inconsistent with semantic or epistemic treatments of vagueness. For instance, Barnes (2014, §2) reflects on the distinction between (FiI) and (IiF) but takes both as sufficient for metaphysical vagueness.²² But she gives no argument for this claim. Surely we shouldn't take it as *obvious* that (FiI) is inconsistent with semantic treatments of vagueness.

Indeed, I claim that (FiI), understood as claiming that the term ‘perfectly natural’ is vague, is entirely consistent with there being no metaphysical vagueness – for the vagueness to be entirely semantic. I am not the first one to claim that the term ‘natural’ is semantically vague – or at least semantically indetermi-

²¹For instance, Sider considers the reason Dorr and Hawthorne adduce for taking ‘natural’ to be indeterminate (discussed below) and rejects it: “one might try to avoid the [hard] choice by saying that, although it's determinately the case that some logical connectives carve at the joints, it's indeterminate which ones do. But [earlier discussion] argued against indeterminacy at the fundamental level” (217).

²²Although she admits to being “tempted by the thought that a case for metaphysical indeterminacy – or at least an *interesting* case for metaphysical indeterminacy – requires a commitment to [(IiF)]” (347) she does think that (FiI) entails (MV) in that the indeterminacy involved in (FiI) would be non-semantic and non-epistemic. For instance, she treats (FiI) as a type of metaphysical indeterminacy and claims that metaphysical indeterminacy is “indeterminacy that would remain even if we spoke a perfect language and were omniscient” (339). And in her discussion of instances of (FiI) she claims that “the indeterminacy doesn't arise from imprecision in language” (358).

nate. Dorr and Hawthorne (2014) arrive at this conclusion, not by reflecting on the GVAS, but by considering perfectly natural asymmetric relations and their inverses. For instance: consider the set-theoretic relation *is a member of* borne by a member to the set that it is a member of. And consider the inverse relation *contains* borne by a set to one of its members. Which of the two relations *is a member of* or *contains* should be the perfectly natural relation? Given that the natural properties are supposed to form a minimal reductive basis, we shouldn't conclude that both relations are perfectly natural. Yet it seems objectionably arbitrary to take only one of the relations to be perfectly natural, as there doesn't seem to be anything about perfect naturalness that could distinguish between the relation and its inverse. The only solution seems to be to treat 'perfectly natural' as indeterminate in reference. Although it's determinate that exactly one of the two relations is perfectly natural, it's indeterminate which.

Of course, the above argument at best shows that 'perfectly natural' is indeterminate in these very particular cases. But the case of asymmetric relations discussed by Dorr and Hawthorne is importantly analogous to the situation faced by anyone who wishes to retain (V1) and (V2). On the one hand, we are under pressure to think that one and only one of various properties that are candidate referents for terms like 'is conscious', 'is a person', 'is uranium' etc., is perfectly natural, but for any such property, it's arbitrary to think that it is privileged.

Let's try to develop that proposal in more detail. If we are to show that 'perfectly natural' is semantically vague in the particular way required to dissolve the GVAS, we need to explain the source of this semantic vagueness. 'Perfectly natural' is a theoretical term, introduced in a theory of naturalness proposed by LEWIS (1983a). Presumably the term 'natural' gets its meaning in the way other theoretical terms get their meanings: in virtue of certain core theoretical roles specified by the theory in which the term is introduced.²³

²³C.f. LEWIS (1970). The insight to treat 'natural' in this framework is a key illuminating theme in DORR AND HAWTHORNE (2014).

At its heart, the idea that theoretical terms get their meaning from their theoretical role is simple. We propose a theory T that uses a new predicate ‘ F ’. By replacing that predicate with a variable, we get an open sentence that doesn’t include the new predicate. That open sentence defines the theoretical role of F -ness, giving us a natural way to specify the meaning of the new predicate. The predicate ‘ F ’ refers to the property that plays that theoretical role of F -ness.

Sometimes, however, there is no property that perfectly plays the entirety of the theoretical role specified by the term-introducing theory. Nevertheless, we want to say that the term refers. The simple idea is in need of at least two refinements. First: a property can be closer or farther from perfectly satisfying a particular theoretical role. Second: certain theoretical roles are more or less *negotiable* than others (LEWIS, 1997, 334). We might, for instance, take naturalness’s role in fixing reference to be a more provisional, less definitional part of the theory of naturalness, so a theorist can sensibly claim that some properties are natural but decline to endorse reference magnetism. The ‘semantic gods’ assign to the theoretical term that property which does best in terms of most closely satisfying the more definitional theoretical roles.²⁴

What if there is no unique property that does ‘best’? In such cases, we should say that the theoretical terms are vague. Consider, for instance, Field’s (1973) discussion of Newtonian mass. The theoretical role outlined for ‘mass’ in the laws of Newtonian mechanics is equally-well satisfied by two properties posited by relativistic physical theories: rest-mass and inertial-mass (and no other property satisfies that role better). The natural thing to say in this case is that ‘mass’ in Newton’s theory was indeterminate in reference between rest-mass and inertial-mass.

There is another, less discussed, way in which a theoretical term can be vague: if the theoretical roles that the term is supposed to play are specified in *vague* language. Consider a psychologist giving a theory that includes a new

²⁴If no property does well enough, the term fails to refer (in the way ‘phogiston’ fails to refer).

theoretical term, say, ‘flow’. A theory of flow might include various vague terms like ‘concentration’ or ‘interest’. When we replace ‘flow’ from the theory of flow with a variable, we are left with an open sentence that states the various theoretical roles in vague terms. Because the role of flow is itself vague, it will thus be vague which property best satisfies that role. The situation parallels that of a vague definite description stated in vague language, like ‘the first bald man to walk in the room’.

In both cases of vague theoretical terms, there is no spooky metaphysical vagueness: the vagueness is of a mundane semantic sort resulting from under-determination in the term-introducing theory. But are any of these stories of indeterminacy plausible for naturalness? Lewis introduced the term ‘perfectly natural’ via a theory in which naturalness was supposed to play a certain role in relation to resemblance, laws, reference, etc. But if our notions of resemblance, laws, and reference are vague, then the theoretical role of naturalness will be vague, in which case we should expect ‘perfectly natural’ to be vague.²⁵

Consider naturalness’s relation to resemblance. According to Lewis’ theory of naturalness: “Natural properties [are] the ones whose sharing makes for resemblance” (LEWIS, 1983a, 13). Indeed, Lewis seems to treat the resemblance role as a more definitional, less negotiable theoretical role, suggesting a nominalist “define natural properties in terms of the mutual resemblance of their members and their non-members” (14). So, if our notion of resemblance is vague, we would expect naturalness to be vague. And from the perspective of someone who takes properties like *personhood*, *consciousness*, *permissibility*, etc. to be perfectly natural, it’s very plausible that resemblance is vague.

Take the perspective of someone who think that *personhood* is perfectly natural. Consider a determinate person, Mary, and a sorites sequence of candidate-persons – Bob₁, Bob₂,...,Bob₁₀₀ – such that Bob₁ is determinately a person and

²⁵In a footnote, Dorr and Hawthorne (2014, fn. 89) suggest that one set of properties might uniquely satisfy the theoretical role of naturalness but it might be indeterminate which; this can be seen as an attempt to vindicate that suggestion by claiming that the theoretical role is specified in vague language.

Bob₁₀₀ is determinately not a person. Now, suppose we ask, for each n : do Mary and Bob _{n} share a feature of resemblance? From the perspective of someone who thinks that personhood is perfectly natural, plainly they will answer ‘yes’ for $n = 1$ (otherwise, personhood wouldn’t make for resemblance so personhood wouldn’t be perfectly natural.) And, plainly the answer is ‘no’ for $n = 100$ (suppose Mary doesn’t share any natural properties besides personhood with any of the Bobs). At what point – at which Bob – does Mary stop sharing a feature of resemblance? One answer would posit a determinate cut-off at which Mary stops sharing a feature of resemblance. A more plausible response, however, acknowledges that ‘resemblance’ is vague: for some middle Bob (say $n = 50$) it’s vague whether Mary shares a feature of resemblance with that Bob. (And because, determinately, personhood makes for resemblance, ‘person’ is also vague.) Because naturalness is introduced in terms of resemblance, we should expect ‘perfectly natural’ to be vague and for the vagueness to match the vagueness in ‘resemblance’ and ‘person’.

On the proposal we’re sketching, perfect naturalness is introduced in terms of resemblance. It’s definitely the case that personhood makes for resemblance although ‘person’ and ‘resembles’ are vague. In naturalness terms: it’s definitely the case that personhood is perfectly natural, although ‘person’ and ‘perfectly natural’ are vague.

This is only one avenue through which vagueness can be introduced in the term ‘perfectly natural’, without the vagueness being metaphysical in any problematic sense: other theoretical roles are also plausibly vague. Fundamental laws of nature are the axioms of the system stated in perfectly natural terms which best balances simplicity and strength. If we think that properties like *permissibility*, *consciousness*, *personhood*, *composition*, etc. figure in moral, phenomenal-physical, or metaphysical laws, we should admit that the notion of ‘law’ is vague. For example, it will be vague which of the following law-candidates is the metaphysical law of personhood:²⁶

²⁶SIDER (2011, §12) acknowledges that metaphysical laws might be vague.

Law Candidate 1 x is a person₁ iff $\phi_1(x)$

Law Candidate 2 x is a person₂ iff $\phi_2(x)$

Law Candidate 3 x is a person₃ iff $\phi_3(x)$

where $\phi_1(x), \phi_2(x), \dots$ state various criteria of bodily / psychological continuity in perfectly natural microphysical language. If our notion of lawhood is vague, and we define the predicate ‘perfectly natural’ in terms of lawhood, then we would expect ‘perfectly natural’ to also be vague.

Let’s recap. The mv-ersatzer claims that it’s vague which of a range of properties ‘person’ refers to. She admits that one of those properties plays the role of perfect naturalness – it makes for resemblance, laws, etc. – although it’s vague which. That’s because the vagueness in the theoretical role of naturalness creates vagueness in the term ‘perfectly natural’.

There is, of course, a serious gap in the argument. We haven’t discussed the role that naturalness plays in fixing reference. Even if reference magnetism isn’t a definitional theoretical role for naturalness – even if the meaning of ‘perfectly natural’ doesn’t come from naturalness’s role in fixing reference – many have found reference magnetism to be a plausible theory of reference. It remains to show how the mv-ersatzer can allow terms like ‘person’ to be semantically vague when *personhood* is perfectly natural – while retaining naturalness’s reference magnetic theoretical role. In the following section, I show how vagueness in naturalness together with reference magnetism *predicts* the vagueness in terms like ‘person’ which pick out a perfectly natural property.²⁷

²⁷In what follows, I will assume that reference magnetism isn’t a definitional theoretical role of naturalness – that it doesn’t fix the meaning of ‘perfectly natural’ – and thus will show how vague naturalness induces vagueness in reference via reference magnetism. Nothing hangs on this decision. If we prefer to treat reference magnetism as a definitional theoretical role of naturalness, then the results of the next section should be interpreted as showing how vagueness in reference induces vagueness in naturalness via reference magnetism in the following way:

As with ‘resembles’, ‘refers’ is vague. So we should expect a theoretical role specified in terms of reference to also be vague. The mv-ersatzer claimed that it’s vague what ‘person’ refers to. If it’s part of the theoretical role of naturalness that terms like ‘*personhood*’ refer to the

6 Step Three: Vague Natural Properties is Not Metaphysical Vagueness

6.1 Where the GVAS Went Wrong

Recall the argument from §2.1 that sought to show that terms that pick out perfectly natural properties cannot be semantically vague. Suppose personhood is perfectly natural. Semantic vagueness results when our linguistic dispositions fail to privilege one of a range of properties as the referent for a term, and none of those properties is more natural than the rest. But (unless there is a range of perfectly natural properties) then there will be a property, *personhood*, which is more natural than the other properties in the range, because *personhood* is perfectly natural and the other properties are not. Therefore: there is no semantic vagueness.

That argument is invalid. In particular, it straightforwardly commits the fallacy of denying the antecedent. In our earlier discussion of reference magnetism, we took Weatherson's (2003) discussion as a guide, and drew two conclusions. The first outlined the role naturalness plays in determining reference:

- (21) If there are a range of properties p_1, \dots, p_n that are consistent with our use of predicate ' F ', then if one of those properties is more natural than the rest, ' F ' refers to that property.

The second outlined how semantic vagueness might remain despite naturalness's role in determining reference:

- (22) If there are a range of properties p_1, \dots, p_n that are consistent with our use

perfectly natural property among the candidate meanings, this theoretical role will be vague. Because it's indeterminate which property '*personhood*' refers to, but it determinately refers to the perfectly natural one, it will be vague which property is perfectly natural. As with the rest of the theoretical role, the reference-role will be satisfied by one of the candidate-extensions for 'perfectly natural' – although it will be vague which.

of predicate ‘*F*’, and if none of those properties is more natural than the rest, then it’s vague which of these properties ‘*F*’ refers to.

Note the direction of the conditional in (22): it gives a *sufficient* condition for semantic vagueness. It does not give a necessary condition for semantic vagueness. The argument that sought to eliminate the possibility of semantic vagueness merely shows that one way a predicate may be semantically vague is not satisfied and fallaciously concludes that there is no semantic vagueness.

But, clearly there are other ways for semantic vagueness to crop up. Indeed, *even if* we thought that semantic vagueness must result from naturalness’s failure to privilege a particular candidate meaning (which I doubt), there are at least two ways in which naturalness might fail to privilege a particular candidate. One way is given by (22): no property is more natural than the rest. But there is a second way: it may be vague which property is more natural than the rest. If perfectly natural properties are reference magnets that anchor our terms, and it’s vague which property is the reference magnet, it will be vague which property our term is anchored to.

We can see this a bit more carefully by considering the implications of the theoretical role of naturalness specified by (21). Consider the range of properties $\text{personhood}_1, \dots, \text{personhood}_n$ that are consistent with our use of ‘person’ such that it’s indeterminate which of those properties is the perfectly natural one. The first antecedent of (21) is satisfied for this predicate, so we have:

(23) If one of the properties $\text{personhood}_1, \dots, \text{personhood}_n$ is more natural than the rest, ‘person’ refers to that property.

According to the mv-ersatzer, one of the properties $\text{personhood}_1, \dots, \text{personhood}_n$ is more natural than the rest – but it’s vague which of the properties that is. Thus, it’s vague which of these properties ‘person’ refers to. Thus, ‘person’ is vague.

The idea is a simple one. There is more than one way for semantic vagueness to arise; (22) only gives us one such way. But another way for there to be

vagueness in reference is for there to be vagueness in the ingredients that fix reference.

And the idea is a general one. It's one that should be endorsed by anyone who admits that there is vagueness in their meta-semantic theory – not just the *mv*-ersatzer. For instance, even if 'perfectly natural' is determinate, many philosophers are open to the idea that *relative naturalness* is vague: there is a pair of properties such that it's indeterminate whether one is at least as natural as the other.²⁸ Suppose our linguistic dispositions for a term fail to distinguish between these two properties. Because it's indeterminate which is more natural, naturalness has failed to privilege one candidate-meaning. By (21), we should claim that the term is vague between these two properties.

The underlying suggestion applies regardless of what we take to be the ingredients for reference, so long as those ingredients are vague. Set aside a reference-magnetic meta-semantic theory and consider, for instance, how a causal theory of reference should accommodate vagueness in causation. According to some causal theories of reference, my token of the word 'Mt. Kilimanjaro' refers to the object that stands in the right causal relation to my tokening.

(24) The name *n* refers to the object that stands in such-and-such causal relation to tokens of *n*.

But it might be vague which object stands in the right causal relation to my tokening and, by (24), this vagueness induces vagueness in which object *n* refers to. Suppose for the sake of argument that there is only one object – Mt. Kilimanjaro – that stands in the right causal relation to my tokening of the name 'Mt. Kilimanjaro', although it's vague which particular hunk of rock that one object is. It's natural to take the vagueness in which hunk of rock appropriately causes my tokening of 'Mt. Kilimanjaro' to induce vagueness in the reference of the name. Even if we object to the specifics of this case – perhaps because we think Mt. Kilimanjaro is distinct from any hunk of rock – it's undeniable that

²⁸C.f. CAMERON (2010)

the causal facts that fix reference can be vague, and it's natural in these cases to treat the vagueness in the causal facts as inducing vagueness in reference. On the causal meta-semantic theory, if it's vague which object my tokening is causally anchored to, then it's vague which object the name refers to.

None of this requires spooky metaphysical vagueness: causation is non-metaphysically vague, and if reference is fixed by causation, reference will be non-metaphysically vague. I've argued that naturalness might be non-metaphysically vague, and if reference is fixed by naturalness, reference will be non-metaphysically vague. Had we wanted to, we could have chosen to use the term 'person' a bit more carefully in order to pick out one of the range of candidate properties.

Return again to the reference-magnetic meta-semantic theory. We can embed the above observation into a supervaluationist theory. As Weatherson pointed out, semantic vagueness results when there is a range of candidate-meanings for a term, each of which is consistent with our use and none of which is more natural than the other. Assume the *at least as natural as* relation is total. Presumably, there will be ties between the relative naturalness of properties. In this case, there will be sets of properties in which no property is more natural than the rest. Thus, the strict ordering based on the *more natural than* relation will be partial: for some pair of properties, neither will be more natural than the other. In light of this observation, Weatherson offers the supervaluationist a natural theory of precisifications according to which:

Every precisification is a completion of the 'naturalness' partial order. That is, each precisification p defines a strict order on possible contents of terms, *more natural- p than*, such that o_1 is more natural- p than o_2 if (but not only if) o_1 is more natural than o_2 . The particular content of terms according to p is then defined by using the more-natural- p -than relation...(485)

As we've noted, there is another way, besides ties, that naturalness can fail to privilege a particular content, and therefore fail to resolve all semantic indeterminacy. By (21), vagueness concerning which property is more natural can also generate semantic indeterminacy. This suggests an obvious way to gen-

eralize Weatherson's suggestion. For some pairs of properties, it's settled that one property is more natural than the other. For other pairs of properties, it's not settled that one property is more natural than the other. That might be because the two properties are tied in naturalness (it's settled that neither property is more natural than the other) or because the *at least as natural as* relation is vague. We should generalize Weatherson's proposal as follows: a *precisification* is an assignment of content based on a way of ordering the naturalness of properties that respects all settled naturalness comparisons of the *more natural than* ordering. In the absence of vagueness in naturalness, this account has Weatherson's as an instance.²⁹

At this point, I can envision two reasons for hesitating. First: I've shown that vagueness in 'natural', and the ensuing vagueness in terms like 'person', is consistent with the reference magnetic role of naturalness as specified in (2I). But (2I) isn't a complete specification of that role. It's not clear how (or whether) vagueness in naturalness will fit into broader accounts of reference magnetism. Second, one might be worried that if vagueness in 'natural' can create vagueness in reference, then 'natural' can no longer play any role in resolving semantic indeterminacy.

In the following section, I'll address both of these hesitations. My response to the first is to simply generalize my discussion above to broader accounts of reference magnetism. We'll see that this generalization unearths additional ways vagueness in reference can be induced by vagueness in the metasemantic theory. My response to the second is to distinguish radical from moderate semantic indeterminacy and show that vague naturalness can resolve radical semantic indeterminacy.

²⁹This also gives us a theory of precisifications which vindicates Cameron's (2010) assumption that metaphysical vagueness in relative naturalness leads to vagueness in reference.

6.2 Generalized Reference Magnetism and Radical Indeterminacy

Reference isn't *radically* indeterminate. The predicate 'mountain' has an extension that includes hunks of rock; it does not have an extension that includes abstract objects like numbers. That reference isn't radically indeterminate should be treated as data. Treat *interpretations* as functions from basic terms in a language to their referents under that interpretation. Names are assigned to objects; *n*-place predicates are assigned to sets of *n*-tuples, etc. (Ignore context sensitivity and other complications.) The data can then be put as follows: some interpretations are clearly *incorrect* or *unintended*. For example, an interpretation that maps the predicate 'mountain' to a set of numbers is clearly incorrect.

As Putnam points out, this data cannot be explained if the standard of correctness for an interpretation merely amounts to making a privileged set of sentences in our language ('total theory') come out true under that interpretation. In general, if there is one interpretation that makes total theory come out true (that is, if the set of sentences are consistent) there will be a great many permutations of that interpretation that also preserve truth of total theory – including permutations that, say, assign the predicate mountain to a set of numbers. Some of those permutations will be intuitively incorrect even though they make the entirety of total theory true. So, the data cannot be explained by the truth of total theory alone. What, then, does correctness for an interpretation amount to?

In the course of giving his theory of naturalness, Lewis carves out a role for naturalness to play in determining correctness for an interpretation and uses that reference-fixing role to explain the fact that reference isn't radically indeterminate.³⁰ On Lewis's view:

³⁰More carefully, Lewis offers *several* possible roles that naturalness might play in order to explain this data: *reference magnetism* is best thought of as a group of meta-semantic theories rather than a single theory. For example, he suggests that naturalness might fix facts about causation and causation might resolve radical indeterminacy. Or, he suggests that naturalness facts might figure in assigning mental content which can then be used to rebut radical indeterminacy

Only an elite minority [of possible referents] are carved at the joints.... Only these elite things and classes are eligible to serve as referents....*Ceteris paribus*, an eligible interpretation is one that maximises the eligibility of referents overall....overall eligibility of referents is a matter of degree, making total theory come true is a matter of degree, the two desiderata trade off. The correct, ‘intended’ interpretations are the ones that strike the best balance. LEWIS (1984, 66)

According to Lewis, naturalness can explain the data by making some interpretations more or less *eligible* in virtue of assigning more or less natural referents to terms in the language. Then, the standard of correctness for an interpretation consists not only in truth of theory but also eligibility.

Interestingly, after sketching his theory of reference magnetism, Lewis claims that “The terms of trade [between truth of theory and eligibility] are vague; that will make for moderate indeterminacy of reference; but the sensible realist won’t demand perfect determinacy.” LEWIS (1984, 67) This last sentence is significant for two reasons.

First: Lewis makes clear that the data he seeks to explain is not perfect determinacy – our language is plainly not determinate. Rather the data to be explained is that our language is not *radically* indeterminate. But some minor vagueness in the *at least as natural than* relation allows us to avoid Putnam’s conclusion: it may be vague which of a limited range of properties is the most natural, but that doesn’t mean that naturalness cannot draw *any* determinate distinctions between properties.

Second: Lewis acknowledges that it might be vague what counts as a correct interpretation – that there might be vagueness in what we mean by ‘the best balance of eligibility and truth of theory’ – and that this vagueness will induce vagueness in reference. Reference is given by the interpretation that strikes the best balance between eligibility and truth of theory, but if it’s vague which

of reference. In what follows, I will focus on just one of the suggested ways in which naturalness might play a role in fixing reference. Although it does not appear to be Lewis’s favored view, and is surely over-simplistic, it is the meta-semantic theory that is typically associated with reference-magnetism. For important discussions by Lewis, see his (1983a; 1984; 1986; 1992; 1994). For various interpretations of Lewis’s discussion, see WILLIAMS (2007), WEATHERSON (2013), SCHWARZ (2014), and SIDER (2001a).

interpretation strikes that balance, reference will be vague.

Lewis points to one way in which it might be vague whether an interpretation strikes the best balance between eligibility and truth of theory: we haven't settled how much of a gain in eligibility is required to compensate for a loss of truth of total theory. But there are other sources of vagueness. For instance, we haven't settled which sentences count as 'total theory'. And we haven't settled on a function from the relative naturalness of referents to degree of eligibility of an interpretation. And, as I've already suggested, it's plausible that we haven't settled on a way to determine relative naturalness of referents from perfect naturalness. Each of these are ways in which we may not have settled the meaning of 'the best balance of eligibility and theory' – ways it may be indeterminate which interpretations strike the best balance – and can thus make for vagueness in reference.

The mv-ersatzer merely adds one more way in which it can be vague whether an interpretation strikes the best balance between eligibility and truth of total theory. If we haven't settled on which properties are perfectly natural then there is vagueness in what makes an interpretation more or less eligible and thus which interpretation strikes the best balance between eligibility and truth of theory.

Suppose '*personhood*' is vague between some properties p_1, p_2, \dots and 'perfectly natural' is vague between some predicates N_1, N_2, \dots such that determinately, personhood is perfectly natural. Because 'perfectly natural' is vague, 'eligibility' will be vague. If we measure eligibility in terms of N_1 , we'll get an interpretation that best balances truth of theory and N_1 -eligibility. This interpretation will assign 'person' to the set of things that instantiate p_1 . More generally, if we measure eligibility in terms of N_n , we'll get an interpretation that best balances truth of theory and N_n -eligibility. That interpretation will assign 'person' to the set of things that instantiate p_n . Because the reference of 'person' is given by the interpretation that strikes the best balance between eligibility and truth of theory, and it's vague which interpretation that is, reference

will be vague.³¹

7 Natural as Natural

When suggesting that ‘natural’ is vague, I am often met with the following reaction:

Naturalness doesn’t seem like the sort of thing that can be vague. Properties like baldness or tallness – properties that aren’t metaphysically deep – can be vague. But some properties, like negative charge, are too metaphysically deep to be vague. And naturalness, like negative charge, is also metaphysically deep – it’s a fundamental feature of reality. So, it cannot be vague.

We can translate the above reaction into a more rigorous objection. Treating a fundamental property as a natural property, the objector seems to be claiming that perfect naturalness is perfectly natural. There are technical complications with making sense of talk about the naturalness of naturalness. But let’s grant the advocate of this objection such talk.

My opponent seems to be constructing the following argument:

- A1. If ‘perfectly natural’ is vague, then perfect naturalness is not perfectly natural
- A2. Perfect naturalness is perfectly natural
- A3. So: ‘perfectly natural’ is not vague.

³¹Once we note the multitude of ways that it might be vague whether an interpretation strikes the best balance between eligibility and truth of theory – besides just vagueness in the more natural than ordering – we see that our Weatherson-inspired theory of precisifications is not general enough. We can further generalize our Weatherson-inspired theory of precisifications as follows. For some pairs of interpretations, it’s settled that one interpretation strikes a better balance between truth and eligibility than the other. For others pairs, it’s not settled that one interpretation strikes a better balance between truth and eligibility – either because the *at least as good a balance* relation is vague or because the two interpretations strike equally good balances. A precisification can be thought of as any interpretation for which it is not settled that there is an interpretation that strikes a better balance.

Before we examine this argument, it will be helpful to note a related argument, discussed in the extant literature, that may also be used to object to my proposal. According to my proposal, ‘perfectly natural’ is only *slightly* vague – there are enough clear cases and its theoretical role sufficiently rich to imbue the term with meaning. However, there is an argument against the hypothesis that ‘natural’ is only *slightly* vague which is inspired by Sider (2011) and is discussed by Dorr and Hawthorne (2014).

- B1. If ‘perfectly natural’ is vague, then perfect naturalness is not perfectly natural
- B2. If so, the simplest definitions of perfect naturalness in terms of the perfectly natural properties will consist of long lists of the perfectly natural properties, of the form ‘is identical to P_1 or is identical to P_2 or is identical to P_3 or...’
- B3. If so, perfect naturalness is extremely unnatural
- B4. If so, ‘perfectly natural’ is horribly vague

As I mentioned above, Dorr and Hawthorne also think that ‘perfectly natural’ is indeterminate, which is why they are concerned to respond to the above argument. They end up accepting premise (B1) (and A1) and instead end up focusing most of their attention on denying the step from premises (B3) to (B4).

I agree with Dorr and Hawthorne that the step from premise (B3) to (B4) is deniable: it may be that ‘perfectly natural’ is extremely unnatural without making the term horribly vague. But I think Dorr and Hawthorne focus their attention on the wrong premise: if the opponent of vague naturalness wins premises (B1)-(B3), the consequences would be disastrous.

Premise (B1) alone would do much to undermine the position that naturalness is vague because of its role in argument (A1)-(A3). (A2) seems quite plausible: it may be part of the *theoretical role* of naturalness that it is perfectly natural. Thus, if my opponent wins premise (A1) / (B1), there is considerable

pressure, from the very concept of naturalness, to admit that ‘perfectly natural’ cannot be vague.³²

Fortunately, we do not need to accept premise (A1) / (B1). The reason that Dorr and Hawthorne adduce for that premise is the following: if perfect naturalness is perfectly natural, then it is highly magnetic, which makes it precise (unless there is metaphysical vagueness). But this argument is simply an instance of the reference magnetic GVAS. And it can be resolved in exactly the same way that we resolved other instances of the GVAS. There can be several candidates N_1, N_2, \dots for the referent of ‘perfectly natural’ even though perfect naturalness is perfectly natural, so long as each of the candidates applies to itself: N_1 is N_1 and N_2 is N_2 , etc. And this vagueness in the meta-semantic facts would induce vagueness in the semantic facts. On one way of specifying what we mean by ‘eligibility’, the interpretation that assigns the predicate ‘perfectly natural’ to the set of properties that instantiate N_1 maximizes (the balance between truth of theory and) eligibility. On another, the interpretation that assigns the predicate ‘perfectly natural’ to the set of properties that instantiates N_2 maximizes (the balance between truth of theory and) eligibility. It’s vague which interpretation is the correct one, thus the term ‘perfectly natural’ is vague.

The previous paragraph was merely meant to show that it’s *consistent* to accept that ‘perfectly natural’ is vague while simultaneously accepting that perfect naturalness is perfectly natural. No doubt, it’s *also* consistent to accept that ‘perfectly natural’ is precise while simultaneously accepting that perfectly natural is perfectly natural. And those that are willing to give up one of (V1)-(V3) need not follow the mv-ersatzer in accepting that ‘perfectly natural’ is vague. But those that find (V1)-(V3) an attractive package of views should follow the

³²With premises (B2)-(B3), premise (B1) also has the result that perfect naturalness would be extremely unnatural. I think this would also be a significant blow to deal to a defender of vague naturalness. That’s because I agree with SIDER (2011) that taking naturalness to be highly unnatural would rob the notion of the explanatory significance required for the the notion to serve its role in reductive accounts of lawhood, duplication, etc. A defense of this claim is well outside the scope of the present paper, but I just note it as a further reason to focus our attention on premise (A1) / (B1).

mv-ersatz. The point of the previous paragraph is that both camps can accept that perfect naturalness is perfectly natural.

8 Conclusion

This paper has sought to accomplish two goals. The first goal is to unify various arguments under a single argument schema – the GVAS – and develop the most plausible version of that schema. Instances of the schema were discovered in a variety of philosophical disputes including debates over personal identity, consciousness, special science properties, ethics, composition, persistence, ontological realism, naturalness, and metaphysical vagueness. And the instances were used to establish rather shocking conclusions. The second goal of the paper is to develop a highly general strategy for avoiding the GVAS. That strategy was ersatz metaphysical vagueness. We demonstrated that, by taking the term ‘perfectly natural’ to be vague, we can mimic genuine metaphysical indeterminacy in a way that allows us to avoid the GVAS. Because the GVAS is a popular way to motivate genuine metaphysical vagueness, ersatz metaphysical vagueness provides us with an unmysterious alternative to objectionable metaphysical indeterminacy.

References

- Elizabeth Barnes. Ontic vagueness: A guide for the perplexed. *Noûs*, 44(4): 601–627, 2010.
- Elizabeth Barnes. Fundamental indeterminacy. *Analytic Philosophy*, 55(4):339–362, December 2014.
- Elizabeth Barnes and J.R.G. Williams. A theory of metaphysical indeterminacy. *Oxford Studies in Metaphysics*, 6, 2011.
- Karen Bennett. Making things up. Manuscript, Forthcoming.

- Ross P. Cameron. Vagueness and naturalness. *Erkenntnis*, 72(2):281–293, 2010.
- Cian Dorr. What we disagree about when we disagree about ontology. In Mark Eli Kalderon, editor, *Fictionalism in Metaphysics*. Oxford University Press, 2005.
- Cian Dorr and John Hawthorne. Naturalness. *Oxford Studies in Metaphysics*, 2014.
- Tom Dougherty. Vague value. *Philosophy and Phenomenological Research*, 2013.
- Matti Eklund. Personal identity, concerns, and indeterminacy. *The Monist*, 87(4):489–511, 2004.
- Hartry Field. Theory change and the indeterminacy of reference. *The Journal of Philosophy*, 70:462–81, 1973. Reprinted with Postscript in FIELD 2001, 177–198.
- Hartry Field. Disquotational truth and factually defective discourse. *The Philosophical Review*, 103:405–52, 1994. Reprinted in FIELD 2001, 222–258.
- Hartry Field. *Truth and the Absence of Truth*. Oxford University Press, 2001.
- Philip Goff. Orthodox property dualism + the linguistic theory of vagueness = panpsychism. In Richard Brown, editor, *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*. Springer, 2014.
- Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell Publishing, 1998.
- David Lewis. How to define theoretical terms. *The Journal of Philosophy*, 67:427–46, 1970. Reprinted in LEWIS 1983b, 78–96.
- David Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377, December 1983a. Reprinted in LEWIS 1999, 8–55.

- David Lewis. *Philosophical Papers, Volume 1*. Oxford University Press, Oxford, 1983b.
- David Lewis. Putnam's paradox. *Australasian Journal of Philosophy*, 62(3):221–236, 1984. Reprinted in LEWIS 1999, 56–77.
- David Lewis. *On the Plurality of Worlds*. Blackwell Publishing, 1986.
- David Lewis. Meaning without use: Reply to Hawthorne. *Australasian Journal of Philosophy*, 70:106–110, 1992. Reprinted in LEWIS 2000, 145–151.
- David Lewis. Reduction of mind. In Samuel Guttenplan, editor, *A Companion to Philosophy of Mind*. Blackwell, 1994. Reprinted in LEWIS 1999, 291–324.
- David Lewis. Naming the colours. *Australasian Journal of Philosophy*, 75:325–342, 1997.
- David Lewis. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 1999.
- David Lewis. *Papers in Ethics and Social Philosophy*. Cambridge University Press, Cambridge, 2000.
- Ned Markosian. Brutal composition. *Philosophical Studies*, 92(3):211–249, December 1998.
- Vann McGee and Brian McLaughlin. Distinctions without a difference. *The Southern Journal of Philosophy*, 33, 1994.
- Kristie Miller. Persons as sui generis ontological kinds: Advice to exceptionalists. *Philosophy and Phenomenological Research*, 81(3):567–593, November 2010.
- Daniel Nolan. Vagueness, multiplicity, and parts. *Noûs*, 40(4):716–737, 2006.
- Derek Parfit. *Reasons and Persons*. Oxford University Press, 1984.

- Miriam Schoenfield. Moral vagueness is ontic vagueness. *Ethics*, 126:257–282, January 2016.
- Wolfgang Schwarz. Against magnetism. *Australasian Journal of Philosophy*, 92 (1):17–36, 2014.
- Theodore Sider. Criteria of personal identity and the limits of conceptual analysis. *Philosophical Perspectives*, 15:189–209, 2001a.
- Theodore Sider. *Four-Dimensionalism*. Oxford University Press, 2001b.
- Theodore Sider. Against vague and unnatural existence: Reply to Liebesman and Eklund. *Noûs*, 43:557–67, 2009a.
- Theodore Sider. Ontological realism. In Ryan Wasserman David Chalmers, David Manley, editor, *Metametaphysics*, pages 384–423. Oxford University Press, 2009b.
- Theodore Sider. *Writing the Book of the World*. Oxford University Press, 2011.
- Ryan Wasserman. Personal identity, indeterminacy, and obligation. In Georg Gasser and Matthias Stefan, editors, *Personal Identity: Simple or Complex?* Cambridge University Press, 2012.
- Brian Weatherson. Many many problems. *Philosophical Quarterly*, 53(213), 2003.
- Brian Weatherson. The role of naturalness in Lewis’s theory of meaning. *Journal for the History of Analytical Philosophy*, 1(10), 2013.
- J.R.G. Williams. Eligibility and inscrutability. *The Philosophical Review*, 116(3): 361–399, 2007.
- J.R.G. Williams. Multiple actualities and ontically vague identity. *Philosophical Quarterly*, 58(230):134–154, January 2008a.

- J.R.G. Williams. Ontic vagueness and metaphysical indeterminacy. *Philosophy Compass*, 3(4), 2008b.
- J.R.G. Williams. Supervaluationism and logical revisionism. *The Journal of Philosophy*, pages 192–212, 2008c.
- Timothy Williamson. *Vagueness*. Routledge, 1994.
- Jessica Wilson. A determinable-based account of metaphysical indeterminacy. *Inquiry*, 56(4):359–385, 2013.
- Dean W. Zimmerman. From property dualism to substance dualism. *Proceedings of the Aristotelian Society Supplementary Volume*, 84:119–150, 2010.

Chapter 3

ONTOLOGICAL DEFLATIONISM, VAGUE EXISTENCE, AND METAPHYSICAL VAGUENESS[†]

The Carnapian campaign for ontological deflationism has been making a comeback as of late. This resurgence has been led by Hirsch and Thomasson, under the respective banners of quantifier variance and easy ontology. According to the quantifier variantist, there are a range of possible languages with quantifier expressions that, when completely unrestricted, have different semantic contributions to the sentences in which they occur.¹ According to the easy ontologist, certain ‘ontologically ampliative’ conditionals from application conditions to the existence of something not mentioned by those application conditions have the status of conceptual or analytic truths.

[†]Thanks to Mercy Corredor, Daniel Drucker, Umer Shakih and Brian Weatherson for discussion. I am especially grateful to David Manley for several discussions and for his help illuminating the metaphysical picture underlying deflationism.

¹A ‘quantifier expression’ is an expression that has a meaning that is similar to the English quantifier. For example, it must obey the introduction and elimination rules for English quantification. Sider adds an additional requirement for quantifier variance: none of the possible quantifier expressions can be *metaphysically distinguished*, but many self-avowed quantifier variantists will take this requirement to be empty (either because it’s trivial or nonsense).

With this (admittedly sketchy) understanding of ontological deflationism, it is a short step from the position to vague existence – that it can be vague what exists. On both views, the extension of the quantifier is highly sensitive to the way we decide to speak – either because there are a range of quantifier meanings that our use is settling as the meaning of the English ‘exists’ or because we could have treated slightly different ontologically ampliative claims as conceptual truths. But, if the quantifier is so sensitive to semantic decision, it is plausible that there will be cases in which we are semantically undecided – cases in which we haven’t settled how to use the word.

Indeed, deflationists frequently take this short step. This attitude is pervasive, for instance, in the work of Eli Hirsh²:

The possibility of quantifier variance brings with it another possibility: quantifier vagueness. If the quantifiers in different languages can have different semantic functions, then it may be indeterminate...precisely how the quantifier functions in a given language. (HIRSCH, 2008, 191)

Is there a problem with vague existence? Lewis (1986) thought so. Lewis distinguishes semantic indecision accounts of vagueness from metaphysical accounts, writing:

The only intelligible account of vagueness locates it in our thought and language. The reason it’s vague where the outback begins is not that there’s this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word ‘outback’. Vagueness is semantic indecision. (121)

He then argues that vague existence is inconsistent with a semantic indecision account of vagueness. Not surprisingly, ontological deflationists are unmoved by this last step in Lewis’s argument. Hirsch, for example, thinks that there are multiple candidate meanings for ‘existence’ over which a linguistic community can be semantically undecided:³

²See also HIRSCH (1999, 43), HIRSCH (2002, 89), HIRSCH (2014, 129).

³See also HIRSCH (1999, 43), HIRSCH (1999, 129), and HIRSCH (2008, 191 and fn.16). Hirsch claims quantifier variance provides a way out of the Lewis/Sider vagueness argument for unre-

I accept Lewis's assumption that vagueness is a matter of semantic indecision. ...[W]hat is the problem Lewis is raising about the vagueness of the quantifier? Since the meaning of a logical constant is given by its role in determining the truth-conditions of sentences, the vagueness of the quantifier would consist in our semantic indecision with respect to the truth-conditions of certain sentences. (2002, 89)

Narrowly construed, Hirsch is claiming that vague quantification is consistent with a semantic indecision model of vagueness. But, it's easy to confuse this narrow claim for a broader claim: that vague existence is not metaphysical vagueness. This broader reading is encouraged by ontological deflationists' tendency to distance their view from anti-realism and endorse flat-footed formulations of realism according to which the world and its objects are 'independent' of our language and thoughts. As Hirsch is careful to emphasize "Nothing is being said here to imply the idealist view that what exists in the world depends on our linguistic or conceptual decisions." (2002, 71). And Thomasson writes: "we should not suggest that the entities...are 'ontologically shallow' or that their existence is somehow to be understood in a deflationary manner." (2015, 146).

Is this broader claim correct? Or does the ontological deflationists' commitment to vague existence commit them to metaphysical vagueness? At first glance the answer seems obvious. And the apparent obviousness of the answer might explain why the question hasn't been examined. The problem is *both* answers look obvious.

On the one hand, the answer looks to be obviously 'yes': vague existence *just is* a form of metaphysical vagueness. Ontology is part of metaphysics: objects are in the world. So if our ontology is vague, then we have metaphysical vagueness: we have vagueness in the world.

On the other hand, the answer looks to be obviously 'no': metaphysical vagueness is not semantic vagueness but the source of vagueness in 'exists', according to the deflationist, is semantic. According to the ontological deflationist, the world is a particular way, but we are merely failing to describe that way

stricted mereological composition – even though this argument explicitly assumes a linguistic theory of vagueness.

in precise quantificational terms. The vagueness of our description need not be reflected in the world.

When two opposite answers to a question both look obvious, that's often a sign that parties are speaking past one another – that the terms of the question are unclear. The purpose of this paper is to (i) clarify the terms of the question and (ii) answer it under various clarifications. Both parts of the project are significant.

First: claims of vague existence and metaphysical vagueness are bandied about without sufficient clarity, so offering regimented interpretations of these claims allow us to press the claims with greater rigor in whatever dialectic they appear.

Second: the tension in the two obvious answers is generated by the deflationist's realist ambitions. On the one hand, we've been instructed that objects are not 'to be understood in a deflationary manner'. But on the other hand, we're told that vagueness in existence is not metaphysical because it's merely the result of our language. If vagueness in how we use the word 'exists' is leading to vagueness in the world, that seems to be a *sort* of anti-realism. For if vagueness in language leads to vagueness in the world, the world is not as independent of language as the deflationist would have us believe. Emphasizing the linguistic source of the worldly vagueness is the only way for the deflationist to soothe those of us who worry about unexplained radical metaphysical vagueness – but it undercuts the view's realism. Investigating the extent to which ontological deflationists must endorse theses of metaphysical vagueness allows us to give a rigorous statement to the extent of their anti-realism.

Two recurring themes will emerge in the course of our investigation.⁴ First: I apply insights from debates over modality to the present inquiry. A parallel is drawn between contingent existence and vague existence and that analogy is the source of many of the advances presented in this paper. Second and relatedly: I deploy higher-order logic to clarify and advance the debate.

⁴Both themes are inspired by the work of Timothy Williamson.

I Characterization of Vague Existence

Consider a standard case of vague existence. Suppose that composition is restricted to ordinary objects like people, tables, watches, etc. And suppose there are some X s on a mat that are in the process of being assembled into a watch, so that it isn't quite clear whether they compose anything. This is an intuitive case of vague existence: it's vague whether the X s compose something.

With this case in mind, we can ask: what is the thesis of vague existence? Here is an obvious response. We can regiment the claim that everything determinately exists as:

$$(25) \quad \forall x \Delta (\exists y (y = x))$$

And the thesis of vague existence is the negation of (25).⁵

$$(26) \quad \neg \forall x \Delta (\exists y (y = x))$$

(26) is equivalent to:⁶

$$\text{Immodest Vague Existence (IVE)} \quad \exists x \nabla (\exists y (y = x))$$

Hawley (2002) considers (IVE) but ultimately rejects it, claiming that (IVE) requires a commitment to non-existent objects. Hawley puts the worry as follows:

The thought is that to posit an object, to quantify over it or refer to it is already to be committed to its existence. Having committed oneself to the existence of an

⁵Strictly speaking, the theses I survey in this paper are about *indeterminate* existence and metaphysical *indeterminacy* rather than *vague* existence and metaphysical vagueness. I will follow other authors in ignoring the distinction between vagueness and indeterminacy. Although there may be special objections to vague existence and metaphysical vagueness that do not apply to indeterminate existence and metaphysical indeterminacy (c.f. EKLUND (2011)).

⁶Informal proof: (26) equivalent to the claim that something doesn't definitely exist: $\exists x \neg \Delta (\exists y (y = x))$. But whatever this thing is that doesn't definitely exist, that thing doesn't definitely fail to exist. (Nothing definitely fails to exist: $\neg \exists x \Delta \neg (\exists y (y = x))$). So, this thing doesn't definitely exist and it doesn't definitely fail to exist, which is the claim of (IVE).

object – ‘there is an object...’ – it is somehow wrong then to draw back from the commitment and claim that it is indeterminate whether the object exists...talk of vague existence conjures up the idea of an object which somehow straddles two domains, the existent and the non-existent. And, if there is no domain of non-existent objects, then such straddling looks impossible, and there cannot be vagueness in existence

Barnes (2010a) also worries about immodest vagueness, writing:

Indeterminate existence is deeply problematic when interpreted de re. The claim that there is some thing such that it’s indeterminate whether that thing exists is hard, if not impossible, to make sense of...the de dicto reading is much weaker: not that there is some thing such that it’s indeterminate whether that thing exists, but rather that it’s indeterminate what things exist.

The Hawley/Barnes worry comes in two strengths. Sometimes, they merely claim that we cannot accept the claim $\exists x \nabla (\exists y (y = x))$.⁷ Other times, they make the stronger claim that we must accept the negation of immodest vagueness: $\neg \exists x \nabla (\exists y (y = x))$.⁸

Presumably it’s this worry with (IVE) that leads Hawley and Barnes away from ‘de re’ characterizations of vague existence. Although Barnes does not say what she means by a ‘de dicto’ characterization of vague existence, she presumably means something like vagueness in a counting sentence, restricted by some appropriate non-vague predicate F in a suitable language. Inspired by VAN INWAGEN (1990), Hawley contrasts (IVE) with her preferred formulation, *modest vague existence*:

Modest Vague Existence (MVE) Existence is vague just in case $\nabla (\exists x F x) \wedge \neg \Delta (\exists x \nabla (F x))$ is true for some predicate F in a suitable language.

⁷In support of this weak reading, note that Hawley makes statements like: “If there are no non-existent objects for quantifiers to range over, then we cannot claim immodestly that there is an object which is such that it is indeterminate whether that object instantiates [the first order property] *existence*.” (138) and “to embrace immodest vagueness we must countenance non-existent objects” (138).

⁸In support of this stronger reading note that Hawley makes claims like: “Without non-existent objects, there are no borderline cases of *existence*, even if there is modest vague existence.” (138) and “if there is no domain of non-existent objects, then such straddling looks impossible, and there cannot be vagueness in existence” (135). It’s not clear which strength Barnes is endorsing.

These alternative characterizations of vague existence strike me as baroque and, in virtue of their appeal to predicates, objectionably language-dependent. But, like Hawley and Barnes, I also do not think that **(IVE)** is the best characterization of vague existence. Fortunately, I disagree with their reasons: de re vague existence neither requires non-existent objects nor is it nonsense. Instead, my preferred statement of vague existence is **(VE)**:

Vague Existence (VE) $\neg \Delta (\forall x \Delta (\exists y (y = x)))$

(VE) is weaker than **(IVE)**. It doesn't claim that there is an object that indeterminately exists; rather it claims that we cannot rule out such an object. Note, however, that **(VE)**, like **(IVE)**, involves the sort of de re vague existence that Hawley and Barnes reject. And the Hawley/Barnes worry for de re characterizations of vague existence like **(IVE)** would also infect my preferred characterization **(VE)**. For suppose the strong worry with **(IVE)** is correct. That is, suppose we know the negation of **(IVE)**: $\neg \exists x \nabla (\exists y (y = x))$. Then we would know that everything definitely exists: $\forall x \Delta (\exists y (y = x))$.⁹ If we know p , then it's definitely the case that p . So, we have $\Delta (\forall x \Delta (\exists y (y = x)))$ which contradicts **(VE)**. So it's important for our purposes to see why their criticism is off-track.

It's helpful at this point to turn to a modal analogue. Note how similar **(VE)** and **(IVE)** are to statements of contingentism in a language with modal operators. (I take this as further reason to prefer **(VE)** over **(MVE)**.) And just as contingentism does not involve positing merely possible objects, immodest vague existence involves no commitment to non-existent objects.

My watch exists. And it has several properties, including the property of *being a watch*. But my watch could have failed to exist. So, it has the property of possibly not existing: $\lambda x (\Diamond \neg \exists y (y = x))$. Consider the modal parallel to the Hawley/Barnes complaint against immodest vague existence applied to my

⁹Informal proof: If nothing indeterminately exists, then everything is such that it either determinately exists or determinately fails to exist. But nothing determinately fails to exist. So, everything determinately exists.

watch: *to posit an object, my watch, is already to be committed to the watch's existence. Having committed to the watch's existence, it is somehow wrong then to draw back and claim that it is contingent whether the watch exists; such talk conjures up the idea of an object which somehow straddles two domains, the actual existent and the merely-possible existent.*

This complaint is simply wrongheaded – contingentists can retain their serious actualist credentials, at least in the face of the present criticism. The watch (actually) exists, so we can quantify over it. And we can predicate all sorts of properties to the watch. Just as we can predicate the property of *being a watch* to the watch, we can predicate the property of *possible non-existence*. At no point are we quantifying over non-existent objects or attributing properties to something that does not exist. And that's true regardless of which property – *being a watch* or *possibly not existing* – we predicate of the watch.

Compare our guiding case of vague existence in which some watch parts *X* have come together to an intermediate degree, such that it's indeterminate whether they compose a watch. Consider the following sentence:

(27) There is a watch composed of the watch parts

Now plainly, (27) is not assertable. And its unassertability isn't mysterious. The following conditionals are definitely true:

(28) If there is something which the watch parts compose, then there is a watch composed of the watch parts

(29) If there is nothing which the watch parts compose, then it's not the case that there is a watch composed of the watch parts.

But it's indeterminate which antecedent of the above conditionals is true. So it's indeterminate whether there is a watch composed of the watch parts. So, (27) is indeterminate. Simply put: since (27) depends on whether the watch parts compose, and it's indeterminate whether the watch parts compose, (27) will be indeterminate.

Of course, then, (27) is unassertable: sentences that are indeterminate are not assertable – just like the sentence ‘Harry is bald’ is unassertable when Harry is a borderline case of baldness. And it’s natural to treat the negation of an indeterminate sentence as indeterminate – this, for example, is true for supervaluationists and allows them to explain why the sentence “Harry is not bald” is also unassertable. So, the sentence

(30) It’s not the case that there is a watch composed of the watch parts

may also be indeterminate. The same story applies to any predicate F for which the conditionals:

(31) If there is something which the watch parts compose, then there is an F composed of the watch parts

(32) If there is nothing which the watch parts compose, then it’s not the case that there is an F composed of the watch parts.

are assertable. In such a case, the sentence

(33) There is an F composed of the watch parts

will be indeterminate. And, plausibly, so will its negation:

(34) It’s not the case that there is an F composed of the watch parts

The *same* is true of a predicate of indeterminate existence. Hawley and Barnes are right that we cannot say that there is an indeterminately existing object composed of the watch parts. But this has nothing to do with the fact that we are predicating indeterminate existence rather than some other property. It’s not that indeterminate existence requires non-existent objects. Rather it’s because it’s indeterminate whether there is an object composed of the watch parts to instantiate the property of indeterminate existence. If the watch parts

do compose, there is an object composed of the watch parts that indeterminately exists. If the watch parts do not compose, there is not an object composed of the watch parts that indeterminately exists. Since whether there is an object composed of the watch parts that indeterminately exists depends on whether the watch parts compose, and it's indeterminate whether the watch parts compose, it's indeterminate whether there is an object composed of the watch parts that indeterminately exists.

Even the weak version of the Hawley/Barnes complaint fails. We can see this by reflecting on what I will call *Joint Vague Existence* (JVE) cases. JVE cases can be described roughly as ones in which we have two cases of vague existence but they are joined so that existence occurs in at least one of the two cases.

Imagine the following case: there are some *X*s and some *Y*s such that (i) it's indeterminate whether the *X*s compose and (ii) it's indeterminate whether the *Y*s compose but (iii) determinately either the *X*s compose or the *Y*s compose.¹⁰ A bit less abstractly: imagine God is creating the world and arranges some *X*s catwise and, some distance away, some *Y*s catwise and declares that there is one cat but says nothing as to whether or not the *X*s or the *Y*s form the cat. Even less abstractly, some three-dimensionalist believes that when caterpillars go out of existence a new object – a butterfly – comes into existence (and the two objects never exist simultaneously). In these cases, there will be a time where (if we are presentists), it may be vague which of the objects (caterpillar or butterfly) exists.

In each of these cases, (IVE) is assertable. Consider the final case: it's indeterminate whether the caterpillar exists or whether the butterfly exists. But, because it's determinate that one of the objects exists, it's determinate that one of the objects will have the property of indeterminate existence (even though it's indeterminate which has this property). I take this to undermine even the weak form of the Hawley/Barnes worry: we can assert (IVE) without positing

¹⁰And, it's not the case that there is something that definitely exists such that's indeterminate whether it is composed of the *X*s or the *Y*s.

non-existent objects.

Thus, in some cases of vague existence, (IVE) is assertable and determinate. In other cases, (IVE) is indeterminate. What these cases have in common, however, is that we cannot *rule out* (IVE). More carefully: in all of these cases, it's not determinate that there is not something that indeterminately exists:

$$(35) \neg \Delta (\neg \exists x \nabla (\exists y (y = x)))$$

We can put the point more simply in terms of determinate existence. In special cases of vague existence – like JVE cases – it's not the case that everything definitely exists. But in *all* cases of vague existence, it's not definitely the case that everything definitely exists, which is my preferred formulation, (VE).

2 Negative Characterizations of Metaphysical Vagueness

We opened by asking whether vague existence requires metaphysical vagueness but noted that the question was not well-posed. With our statement of vague existence as (VE), we are now in a position to ask our question in clearer terms: does (VE) require metaphysical vagueness?

As we saw, if ontological deflationism is correct, it's tempting to conclude that the answer is obviously 'no'. (VE) might be true as a result of semantic vagueness. But semantic vagueness is not metaphysical vagueness. In this section, I'll argue that the motivating question is not so easily dismissed.

Proponents of the 'obviously no' answer to the opening question seem to be gesturing at a negative characterization of metaphysical vagueness made popular by BARNES (2010b). Her negative characterization of metaphysical vagueness is motivated by the simple idea that if there is vagueness and it's neither semantic nor epistemic, then it must be metaphysical:

vagueness has three potential sources – how we represent the world (representational or semantic vagueness), the limits of our knowledge of the world (epistemic

vagueness), or the way the world is in and of itself (ontic vagueness). [...] So if we know that there's (non-epistemic) indeterminacy and we know that our representations are wholly blameless, then we can conclude that the source of the indeterminacy is the world itself. (603-4)

Call this the *exclusion test*. Barnes takes the exclusion test to justify her official characterization of metaphysical vagueness:

(MV) Sentence *S* is metaphysically vague iff: were all representational content precisified, there is an admissible precisification of *S* such that according to that precisification the sentence would still be non-epistemically indeterminate in a way that is Sorites-susceptible.¹¹

If (MV) is true, then vague existence does not entail metaphysical vagueness, given the possibility of ontological deflationism. If existence is vague because we haven't settled on one among many candidate meanings for the quantifier, then if we were to precisify the 'representational content' of the quantifier – if we were to settle on one of the candidate meanings – then there would be no more indeterminacy in the quantifier. Thus, by (MV), (VE) is compatible with the absence of metaphysical vagueness.

The tempting conclusion, however, relies on (MV) being an adequate statement of metaphysical vagueness. But it is not. Modulo concerns with our grip of Barnes' talk of 'sources' of vagueness, the exclusion test is plausible: if we can rule out non-metaphysical sources of vagueness, we are left with metaphysical vagueness. Unfortunately, the exclusion test only gives us a *sufficient* condition for metaphysical vagueness. (MV), on the other hand, attempts to give us necessary *and* sufficient conditions for metaphysical vagueness.

The exclusion test cannot serve as a necessary condition for metaphysical vagueness because a sentence can have sources of vagueness which are both metaphysical and non-metaphysical. It may seem as though (MV) can accommodate such cases, and thus can extend the exclusion test to a full analysis of

¹¹Barnes uses the word 'ontic vagueness' rather than 'metaphysical vagueness' but I don't think the difference is relevant. BARNES (2010b, 604)

metaphysical vagueness. The strategy underlying (MV) is to first eliminate non-metaphysical sources of vagueness and check if any vagueness remain. According to this underlying strategy: vagueness remains just in case the original sentence was metaphysically vague to begin with.

This strategy, however, is misguided. It is neither a necessary nor sufficient condition of metaphysical vagueness that vagueness remains after semantic vagueness is removed. It not sufficient: removing semantic vagueness may *generate* new metaphysical vagueness. More problematic for our purposes, it is not necessary. If semantic vagueness is generating metaphysical vagueness, then removing the semantic vagueness may also remove metaphysical vagueness. If metaphysical vagueness is generating semantic vagueness, removing semantic vagueness may require us to remove its metaphysical source. Or metaphysical and semantic vagueness may have a common source that must be removed in order to remove the semantic vagueness. Or perhaps a single source can both be classified as semantic and metaphysical. This is not merely a complaint about Barnes' use of counterfactuals to regiment talk of 'sources' of vagueness. Counterfactual analyses of dependence relations are notoriously fraught, but my complaint is targeting the strategy that is underlying the counterfactual analysis: that strategy assumes that metaphysical and non-metaphysical sources of vagueness are largely independent of one another.

To make the point less abstract, consider Merricks' (2001) argument that semantic vagueness generates metaphysical vagueness.¹² Suppose our term 'bald' suffers from semantic indecision so that we haven't decided whether Harry falls under its extension. Then the sentences (i) 'Harry is bald' and (ii) "Bald' has the property *describing Harry*' are both vague. And, according to Merricks, the second sentence is *metaphysically* vague because it's vague whether an object (the word 'bald') has a property (the property of *describing Harry*). This would be a case in which semantic indecision is generating metaphysical vagueness but

¹²See also TAYLOR AND BURGESS (2015) for the claim that semantic indeterminacy brings metaphysical indeterminacy.

(MV) fails to count this as a case of metaphysical vagueness. Were all representational content precisified, the sentence “Bald’ has the property *describing Harry*’ would no longer be vague. I don’t wish to suggest that Merricks’ argument is a good one: ultimately I think Merricks’ argument fails¹³. And perhaps we can modify (MV) in a way that accommodates the metaphysical vagueness he envisions. Rather the suggestion is more general: metaphysical and non-metaphysical sources of vagueness may not be as independent as (MV) would have us imagine and we should not rule out this possibility by fiat.

The worry I am raising is particularly relevant for the investigation of this paper. We are seeking to understand whether or not ontological deflationists’ commitment to vague existence requires them to accept metaphysical vagueness. But ontological deflationists tend to put pressure on the independence of the metaphysical from the semantic. It’s hard to make that last claim precise and any deflationist worth their salt will try to accommodate the claim of independence on some readings (for example, deflationists deny that what exists modally depends (in a radical way) on our language). Nevertheless, we should be suspicious of characterizations of metaphysical vagueness that ignore the possibility of such dependencies.

We cannot so quickly conclude that (VE) is consistent with the absence of metaphysical vagueness because (MV) is a flawed analysis of metaphysical vagueness, especially in the context of ontologically deflationary views. Unfortunately, I have no alternative analysis of ‘metaphysical vagueness’ to offer.

I think the term ‘metaphysical vagueness’ is a bit like the term ‘realism’. Some choose to regiment claims of realism in terms of excluded middle and bivalence. Others regiment such claims in terms of fundamentality, grounding, structure, or truth-making. Still others understand realism in modal terms. Perhaps there is an important debate about which characterization *really* captures

¹³See WEATHERSON (2003, 500-501).

the claim of realism (although I am skeptical) and perhaps that debate will eventually be settled. Until then, however, we are best off just understanding the regimentations in their own terms. Then, instead of asking whether a position is realist *simpliciter*, we can ask whether a position conforms to a particular regimented thesis that is supposed to capture a particular *sort* of realism.

I adopt the same methodology with respect to metaphysical vagueness. I'll survey a range of theses that seem to capture some sense of metaphysical vagueness and the slogan 'vagueness in the world'.¹⁴ We'll try to understand those theses on their own terms. And then we can ask whether or not (VE) entails or violates those theses. If we can come to understand the metaphysical import of the thesis under discussion, we can explore the consequences of vague existence on our picture of reality, without delving into verbal disputes about applications of the label 'metaphysical vagueness'.

Of course, many will treat (VE) itself as an interesting sort of metaphysical vagueness – which is why some found the answer to the opening question so obvious. Of course, the deflationist will explain (VE) as resulting from our language. But, as I've pointed out, this is in tension with her insistence that what exists in the world is 'independent' of language. In order to see the full extent to which the deflationist must admit that vagueness in language results in worldly vagueness, we'll continue to explore alternative ways of capturing the notion of metaphysical vagueness.

3 Fuzzy Set Theory

It might be thought that *fuzzy sets* are a way that the world can be vague:

¹⁴One sort of metaphysical vagueness that I won't explore here can be stated in terms of metaphysical fundamentality. Roughly put: fundamental vagueness is a sort of metaphysical vagueness. Borrowing the apparatus of *SIDER* (2011) in which we can talk about a fundamental language, we can put the thesis as follows: if the determinacy operator is part of the fundamental language we have a sort of metaphysical vagueness. It's not clear, however, how vague existence would require fundamental vagueness. See Chapter 2 for relevant discussion.

Fuzzy Sets (FS) $\exists x \exists y \nabla (x \in y)$

Perhaps we can construct an argument that vague existence entails (FS). Hawley (2002) attempts to show just this:

If there is vagueness in what things the world contains, this might be thought to entail vague identity among sets. For example, consider the set of the simples, P [P names the plurality of simples], which are the candidate-composers of the unhealthy [vaguely existing] hamster: call this set ' SP '. Now consider the set of the things on the mat, SH . If the hamster were determinately dead, then, according to van Inwagen, P would be the only things on the mat. If the hamster were determinately alive, then the things on the mat would be P plus the hamster. This suggests that in the indeterminate situation, it is indeterminate whether SP is identical to SH . (132-133)

Hawley worries that this sort of indeterminacy would violate the conclusion of the Evans-Salmon argument against vague identity, and so suggests a modification of the Extensionality Principle of set identity according to which "Sets can be determinately distinct without determinately differing in their membership, just so long as it is indeterminate whether they differ in their membership." (133) On this modification, SP and SH are (determinately) distinct.¹⁵

(FS) is a radical claim. But, it's not clear that it does much to capture a sense of metaphysical vagueness. That's because it's not clear that the term ' \in ' isn't vague in the way terms like 'smarter than' are vague. Admittedly, ' \in ' is a different sort of term from 'smarter than': it's a quasi-logical theoretical term. Perhaps this can form the basis for thinking that the vagueness is different in kind from the more mundane sorts of vagueness in terms like 'smarter than'.¹⁶

Fortunately, we don't need to answer these questions in order to understand the ramification of (VE) and ontological deflationism. That's because Hawley's argument that vague existence requires vague set membership is fallacious – even spotting her the cogency of the Evans-Salmon argument.

¹⁵Hawley does not endorse vague existence or (FS), only the claim that *if* there is vague existence then there will be vagueness in set membership.

¹⁶For instance, it might be subject to a reference magnetic GVAS of the sort examined in Chapter 2.

Note first that not all vague identity statements run afoul of the Evans-Salmon argument. As Thomason (1982) and Lewis (1988) have pointed out, the Evans-Salmon argument rules out:

$$(36) \exists x \exists y \nabla(x = y)$$

It does not, however, rule out:

$$(37) \nabla(a = b)$$

so long as (37) doesn't entail the problematic (36). And that entailment is deniable if the singular terms '*a*' and '*b*' are vague.

Hawley thinks that the conclusion $\nabla(SP = SH)$ runs afoul of the Evans-Salmon argument and, for this reason, gives up Extensionality. She is, however, aware of the Lewis/Thomason observation that, if the singular terms flanking the identity statement are vague, statements of vague identity don't run afoul of the Evans-Salmon argument. Presumably, she suspects that in the case of vague existence described, the terms '*SP*' and '*SH*' are not vague. Let's explain that suspicion by considering a modal analogue to the case.

Suppose with van Inwagen that composition only occurs when simples form a living organism. Suppose further that there are some simples *G* on the mat and those simples are arranged hamster-wise and compose a hamster. It's also possible that those simples are arranged otherwise on the mat and do not compose anything. Let ' $\Theta(S)$ ' be short for ' $\neg \Box(S) \wedge \neg \Box(\neg S)$ '; so $\Theta(S)$ says that it's contingent whether *S*. In this case, we can say:

$$(38) \Theta(\text{the set containing the } G\text{s} = \text{the set containing everything on the mat})$$

Clearly, we cannot infer that there is contingent identity: $\exists x \exists y \Theta(x = y)$. That's because it is contingent what the definite descriptions refer to. It's possible that they refer to the same thing; it's also possible that they refer to distinct things. Moreover, because the existence of a hamster composed of *G* is contingent, the existence of any set containing such a hamster is also contingent

(assuming that set contain their members essentially). Contingent existence begets contingent existence.

A parallel response is available to the proponent of vague existence. We cannot infer the problematic (36) from $\nabla(SP = SH)$. The name ‘*SP*’ was introduced as the set of simples on the mat and the name ‘*SH*’ was introduced as the set of everything on the mat. But it’s indeterminate what the definite description ‘the set containing everything on the mat’ refers to: it may refer to the set of simples on the mat, or, if such a set exists, it may refer to a set containing some simples and a hamster. It’s not determinately the case that it and the definite description ‘the set of simples on the mat’ refer to distinct things. Nor is it determinately the case that they refer to the same thing. Because it’s vague whether there exists something other than the simples on the mat, it’s vague whether ‘*SP*’ and ‘*SH*’ refer to the same object or not. Moreover, because the existence of a hamster is vague, the existence of any set containing such a hamster is also vague. Vague existence of members begets vague existence of sets.

4 The Modal Analogy and De Re Vagueness

In order to focus our efforts on extracting senses of metaphysical vagueness, we can continue to exploit the modal analogy. The world and its properties are infused with modality in a way that (if there is no metaphysical vagueness) the world and its properties are not infused with vagueness. What, however, is the way that the world and its properties are infused with modality?

One way in which the world and its properties are infused with modality is that objects are subject to *de re* modal properties. In Quine’s (1953) terms, objects are subject to the third grade of modal involvement. On this proposal, the world is infused with contingency in case we can find some formula ϕ free in z such that:

De Re Contingency $\lceil \exists x(\lambda z(\diamond\phi z)x \rceil$

The analogous suggestion is that the world is infused with indeterminacy in case:

De Re Indeterminacy (DRI) $\lceil \exists x(\lambda z(\nabla\phi z)x) \rceil$

Sainsbury (1994) considers this proposal for capturing the claim that the world is vague and correctly rejects it:

It is hard to see how one could prevent a thesis of ontic vagueness based on [(DRI)] entailing that every borderline case is a vague object. Suppose this ripening tomato is a borderline for 'red'. Then it is vague whether or not this is red. Why is not the corresponding statement, in which the vagueness operator is imported so as to have narrow scope, also true? That is, why is it not also true that $\lambda z(\nabla(\text{red}z))$ this? If it is, then this tomato counts as a vague object merely because it is a borderline case for 'red'; a result which...ensures that we do not have a thesis of ontic vagueness.

Sainsbury's point here is simple. Suppose it's indeterminate whether Harry is bald. Then, supposing that the name 'Harry' isn't vague, Harry will be a borderline case of baldness. That is, Harry is indeterminately bald. So, there is something that is indeterminately bald – namely Harry. Thus, any view that accepts the existence of borderline cases must admit (DRI), but surely we can admit borderline cases without being committed to worldly indeterminacy! (DRI) doesn't capture an interesting sense of worldly indeterminacy.

Sainsbury uses this observation to question another prominent characterization of metaphysical indeterminacy, put forth by Michael Tye. According to Tye, there is metaphysical indeterminacy if

$$(39) \exists x \lambda y \exists z (\nabla (y \text{ is part of } z))x$$

More simply, there is metaphysical indeterminacy if there is an object that is indeterminately part of another object – a borderline case of parthood. But Sainsbury's observation undermines the *prima facie* plausibility of this definition: just as we can give non-metaphysical accounts of borderline cases of baldness according to which (let's say) we haven't settled the meaning of 'bald', we can

give a non-metaphysical account of borderline cases of parthood according to which (let's say) we haven't settled the meaning of 'part'. Absent some reason to think that such an account won't apply to 'part', we are left wondering why (39) would give us a case of metaphysical vagueness.

Is there some reason to think that whatever account applies to 'bald' cannot apply to 'part'? Is there some reason to think that (39) is a special sort of metaphysical vagueness? In response to Sainsbury, Tye (2000) gives two such reasons, neither of which is persuasive. First, he argues:

Having a possible borderline case...is not a sufficient condition for a predicate's being vague. So long as the predicate does not have a meaning that precludes its being satisfied by precise objects, what is needed, I suggest, for a necessary *and* sufficient condition is the possibility of a borderline case involving precise objects. (200)

Because 'part' cannot have borderline cases involving precise objects, the predicate 'part' cannot be vague according to Tye's criterion. So, if there is vagueness in parthood, it must not be from the vagueness of the predicate but must instead be metaphysical.

It's not clear that metaphysical accounts of vague parthood are consistent with the predicate 'part' not being vague – the predicate will be vague, even if the vagueness is metaphysical. But even setting this worry aside, Tye's argument is clearly question begging. Tye takes a precise object to be one that has only determinate parts. But in so far as one thinks that 'part' can be non-metaphysically vague, they will clearly reject Tye's criterion, which has no independent plausibility. Tye's second argument is no better.

...every object is a part of itself. If this is so, then there cannot be an object such that it is indeterminate whether it is a part of itself, even if some objects, actual or possible, are vague. Where a relational term is vague, it has possible borderline cases in which a single object is related to itself – provided that the term's meaning permits it to have non-borderline cases of application of this sort....It follows again that 'part of' is not vague. (200)

This second proposed criterion for a predicate's being vague not only lacks independent plausibility: it is positively implausible. Consider the predicate *is*

at least as smart as or *or is smarter than* or *identical to*. Clearly, that predicate is vague, yet there are no possible borderline cases in which a single object is related to itself: everyone is at least as smart as themselves.

The best proposal for arguing against non-metaphysical accounts of vagueness in ‘part’ is to link vague parthood to vague existence and argue that vague existence is a sort of metaphysical vagueness.¹⁷ The second step of that argument is, of course, exactly what is being investigated in this paper.¹⁸

5 Vague Properties and the Problem of the Many Properties

After noticing that (DRI) fails to capture an interesting sense of metaphysical vagueness, Sainsbury gives up the modal analogue for formulating theses of metaphysical vagueness.¹⁹ He continues his exploration for a thesis of metaphysical vagueness without this analogy, exploring in particular the thesis that there are vague properties. Unfortunately, the project of articulating the thesis of vague properties ultimately frustrates him.

Can we do better? The slogan that we are trying to capture is: if there is no metaphysical vagueness, then properties themselves aren’t vague, even if our predicates are. But what is a precise property? A misguided attempt can make the slogan look as hopeless as (DRI).²⁰ It might be thought:

A property like *having at least 50,000 hairs* is (or approaches something like) a precise property, while a property like *baldness* is a vague property. We can see this because, in mundane cases of vagueness without e.g. vague existence, the

¹⁷For more on the link between vagueness in parthood and vagueness in existence, see SIDER (2001, §4.9)

¹⁸Another proposal treats parthood as a reference magnet and thus not apt for semantic indecision accounts. Applying the exclusion test, we can conclude that vagueness in ‘part’ is metaphysical. I discuss this style of argument in Chapter 2.

¹⁹He goes on to explore a refinement of the test according to which the world is vague in case we can find a *sharp* predicate ϕ for which (DRI) is true but worries about finding a way of defining ‘sharp predicate’ in a way that does not trivialize the thesis.

²⁰SAINSBURY (1994) also expresses the concern below, but does not offer the solution I do.

predicate ‘bald’ is a vague one in that we can say: $\exists x \nabla (\text{Bald}x)$ – or at least there are possible worlds in which this is true. But $\exists x \nabla (\text{having less than } 50,000 \text{ hairs}x)$ isn’t true at any possible world.²¹ So, if there is a property *baldness*, there is a vague property.

But, this way of capturing the slogan looks hopeless, because *of course* there is a property of baldness. Surely there are people who are bald. And from the obviously true conditional ‘if someone is bald, then someone has the property baldness’, we can conclude that there is a property of baldness. So, there are vague properties.

The only option for avoiding this trivialization is to deny the ‘obviously true’ conditional. And indeed, many ‘hardcore’ ontologists who find the existence of properties a substantive question will deny that the conditional is obviously true. But, it looks quite plausible that the conditional is true – even if not obviously so. Even if the proposed test doesn’t entirely trivialize the thesis of metaphysical vagueness, it makes the thesis as plausible as the overwhelmingly plausible conditional. And, presumably, opponents of this instance of the de-nominalizing conditional will oppose all instances of the conditional. For these nominalists the fact that there are no vague properties is a uninteresting consequence of the more general fact that there are no properties. But moving to nominalism is patently dodging the question of metaphysically vague properties.

Reflecting on the above failure to cash out this understanding of vague properties yields two lessons. First: the issue of metaphysically vague properties seems orthogonal to the realism/nominalism debate over property ontology. Speaking roughly, the claim that there are metaphysically vague properties has two parts: (i) that there are properties and (ii) that these properties are metaphysically vague. It’s the second part that we are interested in. The claim that there are metaphysically vague properties should have a nominalist analogue

²¹As SAINSBURY (1994) points out, defining a predicate as precise or vague will require modal notions. One problem with this use of modal notions in an attempt to define metaphysical vagueness is that it makes the thesis of metaphysical vagueness *necessary*. But one may think that metaphysical vagueness is merely possible. I won’t discuss this problem further.

that doesn't presuppose the existence of properties and that can be understood and debated by realists and nominalists alike.

Second: the strategy for distinguishing precise and vague properties ties predicates too closely with the properties they stand for. This should raise alarms: whether a predicate is vague is a linguistic matter whereas whether a property is vague is a metaphysical matter. Recall the slogan we are trying to capture: properties aren't vague, even if our predicates are. From the perspective of a theorist that endorses this slogan, a predicate can be vague even if there are only precise properties if it's vague which precise property that predicate picks out. Suppose borderline cases for the predicate 'bald' are limited to those with either 50,000 or 50,001 hairs on their head. From this perspective, there are precise properties like *having less than 50,000 hairs on your head*, *having less than 50,001 hairs on your head*, etc.²² And the property of *baldness* isn't some distinct property over and above any of these precise properties. Rather, it's vague which of these precise properties is picked out by the name 'baldness'. Similarly for the complex properties of **(DRI)** defined by λ -abstraction using the indeterminacy operators. The property of *being indeterminately bald* isn't some metaphysically strange property – it's just the property of *having either 50,000 or 50,001 hairs*.

This approach parallels a popular solution to the Problem of the Many that denies metaphysically vague objects.²³ On this solution, Kilimanjaro isn't some distinct object over and above the various candidate-mountains with precise borders. Rather the term 'Kilimanjaro' vaguely refers to one of these precise candidate-mountains. According to this solution: $\neg\exists x \Delta (x = k)$ (there isn't something that is definitely Kilimanjaro). We can express the analogous idea for the property of *baldness* (symbolized as B). And, we can nominalize the idea by using second-order logic with primitive quantification into predicate position. The claim that there isn't some property that is definitely baldness can be writ-

²²For exposition purposes, I am assuming that there is no vagueness in the attribution of this property.

²³See LEWIS (1993).

ten as: $\neg\exists X \Delta (X \equiv B)$ (where ‘ \equiv ’ expresses (something like) identity for the nominalized properties).

6 Vague Properties and the Return of the Modal Analogy

We’ve seen that the strategy of defining vague properties as the properties referred to by vague predicates fails, because a vague predicate can vaguely refer to a property we intuitively think of as precise. How then do we define a vague property? Sainsbury gives up the modal analogy prematurely: the modal analogy is fruitful in formulating a thesis of metaphysically vague properties. Our properties are infused with modality in a way that, if there is no metaphysical vagueness, they are not infused with vagueness. Properties have a modal dimension to them: it’s contingent which objects instantiate these properties. We can put that idea more carefully as follows:

Modality of Properties (MP) $\Diamond(\exists X \exists x((Xx) \wedge \Diamond(\exists Y((Y \equiv X) \wedge (\neg Yx)))))$

Indeed, we may want to go so far as to include the following comprehension schema for (monadic) properties:²⁴

Strong Modal Comprehension (SMC) $\exists X \Box \forall x (Xx \leftrightarrow \phi)$

where ϕ can be substituted for any formula not free in X and the resulting instance of the schema can be prefixed with any number of universal quantifiers and necessity operators. (SMC) ensures that properties have a modal dimension to them. For example, we can substitute the formula ‘is five feet from a banana’ for ϕ ; the resulting instance would guarantee the existence of a property that applies at world w to all and only objects that are five feet from a banana in world

²⁴This is a major piece of Williamson’s (2013) argument for necessitism.

w. Given that different things can be five feet from a banana, the property's intension will vary and **(MP)** will be satisfied.²⁵

(MP) captures the idea that properties are infused with modality – that properties have a modal dimension. Following the modal analogy reveals an obvious way to distinguish precise and vague properties. A property is precise if it definitely applies to the same things – if, wherever the property exists, it has the same extension. Properties infused with vagueness – properties with a determinacy dimension to them – make for a sort of metaphysical vagueness. The claim that definitely there are no metaphysically vague properties can be expressed as follows:

No Indeterminacy of Properties (NIP) $\Delta (\forall X \forall x (Xx \rightarrow \Delta ((\exists Y (Y \equiv X)) \rightarrow Xx)))$

Now, a proponent of vague existence that is attempting to avoid metaphysical vagueness need not accept the vagueness analogue of **(SMC)**

Strong Vague Comprehension (SVC) $\exists X \Delta \forall x (Xx \leftrightarrow \phi)$

Instead, they can accept a weaker comprehension schema

Weak Vague Comprehension (WVC) $\exists X \forall x (Xx \leftrightarrow \phi)$

Unlike **(SVC)**, **(WVC)** does not guarantee violations of **(NIP)**. Rather, **(WVC)** only guarantees that there is a property for every extension, but is silent on whether that property has that extension definitely or indefinitely. Thus, it is

²⁵Williamson also claims that Strong Modal Comprehension establishes higher-order necessitism – that is, necessarily all properties necessarily exist. I disagree: if properties are individuated more fine-grainedly than their intensions, **(SMC)** does not establish higher-order necessitism. This is my preferred response to Williamson's argument for (first-order) necessitism. We can accept that necessarily there is a property that has the intension of my haecceity without admitting that my haecceity necessarily exists. See WILLIAMSON (2013, 264-266) and FRITZ AND GOODMAN (2016).

consistent with (WVC) that properties, where they exist, have their extensions definitely.²⁶

With yet another clear statement of one sort of metaphysical vagueness, we can ask: does vague existence violate (NIP)? It does not, but avoiding violations of (NIP) reveals some interesting consequences for vague existence.

The following constraint is intuitive:²⁷

Being (B) $\forall X \forall x \Delta (Xx \rightarrow \exists y (y = x))$

In English: every property and everything is such that it's determinate that, if that thing has that property then that thing is something. I don't have an argument, per se, for this principle. Yet, I take this principle to be about as obvious as any metaphysical principle: I simply can't see how something can have a property – can be a certain way – without existing. Similarly for the modal version of the thesis, which is sometimes called the claim of Serious Actualism.

Now suppose it's vague whether some X s on a mat compose something. Determinately, if the X s compose some y then y will have a great number of properties. But those properties cannot definitely apply to the same things. By (B), the property cannot definitely apply to y because y doesn't definitely exist. So, we have an apparent violation of (NIP).

Recall, (NIP) only demanded that properties have the same extension *if they exist*. The above argument rests on the assumption that (definitely) all

²⁶Note that the argument that Williamson (2013) puts forth in favor of (SMC) does not apply to (SVC). Roughly: he argues that (SMC) is the best explanation for the fact that there is a property for every extension *and some* varying intensions. But in the vagueness analog, we need not accept that any varying 'intensions' have corresponding properties. Although I don't explore it here, Williamson's argument may be exploited for an interesting conclusion. For those that admit some indeterminacy in properties, in the sense that some properties have a determinacy dimension to them, there is abductive pressure for them to accept (SVC). And (SVC) might be used to argue against (VE) in the same way Williamson uses (SMC) to argue against contingent existence. Thus violations of (NIP) may rule out (VE).

²⁷See WILLIAMSON (2013) and FINE (1985) for a discussion of the modal version of this principle.

properties definitely exist. A proponent of vague existence need not accept this claim. They may simply deny that properties definitely exist. Then: definitely for any object that indeterminately exists, if there are any properties that apply to that object, those properties also indeterminately exist.

This is yet another case in which vague existence begets vague existence. In order to avoid metaphysical vagueness, deflationists that posit vagueness in first-order quantification must admit that higher-order quantifiers can also be vague. This is a more significant result than it may first appear. Consider the sorts of properties that would apply to the composite of the X s, if that composite exists. Presumably, the property of *being self-identical* is one such property. And if the X s are arranged as a hamster in the process of dying, the property of *being a hamster* is another. In order to dodge violations of (NIP), the proponent of (VE) must accept that, of the properties *being self-identical* and *being a hamster*, those properties don't determinately exist. Thus, deflationists must not only admit vagueness in existence; they must also either accept metaphysically vague properties or vagueness in the existence of a great number of properties. Both options, however, look like radical results. The deflationist must try to explain this apparently radical consequence of their view by pointing to its source in the vagueness of language. That might help us swallow the results, but undermines the extent of the ontological deflationists' realist ambitions.

7 Vague States of Affairs

Return once more to the modal analogy. As explained above, modality infuses properties in a way that vagueness does not. Modality not only infuses properties; modality also infuses *states of affairs*. Consider the state of affairs of *Anisha's wearing a purple shirt*. Let $[\phi(x)]$ name the state of affairs of x 's ϕ -ing. So $[Pa]$ names the state of affairs of Anisha wearing a purple shirt. That state of affairs is contingent: although it obtains, it isn't necessary that it obtains. More formally, we have:

$$(40) \exists x(x = [Pa] \wedge (Ox \wedge \neg \Box(Ox)))$$

where ‘O’ predicates obtainment of states of affairs. More generally, we can state the claim that there are possible contingent states of affairs as:

$$(41) \Diamond(\exists x(Ox \wedge \neg \Box(Ox)))$$

Contingency infuses states of affairs, but if there is no metaphysical vagueness, indeterminacy does not infuse states of affairs. States of affairs are determinate:

$$\text{Determinate States of Affairs (DSA)} \quad \Delta(\forall x(Ox \rightarrow \Delta(Ox)))$$

Prima facie, (DSA) does a nice job capturing the idea that reality is not vague even if our language is. Even though the language we use to describe states of affairs is vague, the states of affairs *themselves* are not vague. We might worry that, like prior attempts to state theses of metaphysical vagueness, (DSA) makes metaphysical vagueness trivial. Any theory of vagueness should allow that it’s indeterminate whether Harry is bald. But then, they should allow that it’s indeterminate whether the state of affairs of *Harry’s being bald* obtains. Doesn’t this violate (DSA)?

It does not. It may be indeterminate whether the state of affairs of Harry’s being bald obtains, but, as I’ve been stressing, that doesn’t entail that there is something such that it’s indeterminate whether that thing obtains. An opponent of metaphysical vagueness should say that the singular phrase “the state of affairs of *Harry’s being bald*” is a vague singular term for states of affairs that definitely obtain or definitely do not obtain – states of affairs like *Harry’s having less than 50,000 hairs*. More formally, the claim that it’s indeterminate whether the state of affairs of *Harry’s being bald* obtains is:

$$(42) \nabla O[Bh]$$

But (42) does not entail the de re claim:

(43) $\exists x \nabla O x$

And it is this de re claim that would violate (DSA).

(DSA) puts pressure on proponents of deflationary vague existence to admit metaphysical indeterminacy. The intuitive source of the pressure is that states of affairs are structured. That is, they involve various objects and properties. For example, the state of affairs of *Obama's being president* involves the object Obama and perhaps the property *being president*. In general, the state of affairs picked out by $[\phi(a)]$ where 'a' is a proper name, involves *a*.

The claim that states of affairs are structured in this way is supposed to be quite weak. I don't mean to be making any controversial claims about the state of affairs being *grounded* in the objects they involve. Nor do I mean to be making the claim that states of affairs that involve objects have *haecceitistic obtaining-conditions* unique to states of affairs that involve that object.²⁸ All I mean to claim is that states of affairs have objects as constituents.

Plausibly, the existence of an object comes with an obtaining state of affairs which involves that object. So, if we cannot rule out some *X*s composing a table, we cannot rule out an obtaining state of affairs involving that table. But if (DSA) is right, then it's determinate that every obtaining state of affairs determinately obtains. But that obtaining state of affairs, which involves the table, cannot obtain without the table that it involves. Thus, that table definitely exists.

Let's try to formalize the above reasoning. Introduce the two place predicate '*I*' so that ' xIy ' says that an obtaining state of affairs *x* involves the object *y* in the way the state of affairs of *Obama's being president* involves Obama. It's definitely the case that, for every object, there is a state of affairs involving that object:

(44) $\Delta (\forall x \exists y (yIx))$

²⁸Although, the argument that I make below can be given in these terms as well. All that is required is a conception of a state of affairs according to which the existence of an object brings with it a new state of affairs. See HAWTHORNE (2009, §2) for some relevant considerations.

From **(DSA)** we know that every state of affairs definitely obtains; combined with (44) we can conclude that for any object, there is a state of affairs involving that object which definitely obtains:

$$(45) \Delta (\forall x \exists y (yIx \wedge \Delta (Oy)))$$

And, states of affairs determinately depend on the objects they involve in order to obtain:

$$(46) \Delta (\forall y \forall x (yIx \rightarrow \Delta (Oy \rightarrow \exists z (z = x))))$$

But (46) and (45) jointly contradict **(VE)**: determinately everything determinately exists. According to (45), (determinately) each object comes with a state of affairs that determinately obtains, but by (46), in order to obtain, that state of affairs must ‘carry’ the object along with it. So, determinately every object must determinately exist.

8 Two Deflationist Replies

The above argument conceives of states of affairs as having a certain structure to them. It also claims that if a state of affairs involves an object, then that state of affairs definitely does not obtain without that object. Presumably, this last claim is motivated by the assumption that if a state of affairs has a certain structure to it, it definitely has that structure. Thus, the conception of states of affairs underlying the above argument is one on which they have definite structure.

And it’s no wonder such a conception of states of affairs causes trouble for the deflationist. Indeed, outside of the context of vagueness, Hirsch (2002) has explicitly distanced himself from such conceptions of states of affairs:

...statements involving different kinds of quantifiers can be equally true by virtue of the same (unstructured) facts of the world.

The notion of a structured fact does, however, raise certain problems for quantifier variance – but...[i]f I change what I mean by ‘a thing’ then I must also change

what I mean by ‘the way a fact is built up out of things and properties’; hence I must change what I mean by “a structured fact”....

A problem arises, however, if one wants to say...that each true sentence states one structured fact. How can we say which one that is...when I am able to ‘translate’ that sentence into a variety of structurally different sentences in my language having the same truth conditions?...I am inclined to agree with Putnam that, once we’ve accepted quantifier variance, there is no point in trying to hold onto language-shaped facts that are in the world independent of language. (78-79)

Hirsch’s discussion inspires two strategies the deflationist may adopt in order to avoid violations of (DSA), each questioning the conception of states of affairs with definite structure.

- **Strategy One:** Accept that there are structured states of affairs, but deny that the states of affairs have their structure determinately.
- **Strategy Two:** Deny that there are structured states of affairs.

Strategy one accepts that there are states of affairs with structure – that is, there are states of affairs that involve objects. But, it can be vague what structure it has – it can be vague which objects a particular state of affairs involves.

Given that we’re allowing structured states of affairs, presumably (44) will be unproblematic: every object has a state of affair that involves it. Instead, the strategy motivates a rejection of (46) – the claim that (determinately) if a structured state of affairs involves an object, then determinately if that structured state of affairs obtains, the object exists. This strategy treats the ‘*I*’ predicate as vague. Even if the structured state of affairs cannot obtain without the object it involves, if it’s vague what object it involves then it need not determinately ‘carry’ the involved object with it in order to obtain. Without (46), (VE) is consistent with (DSA) – vague existence does not entail vagueness in the obtainment of states of affairs.

I am skeptical that this strategy can be made to work. My reason is closely related to the reason Hirsch himself isn’t optimistic of this strategy: which new description will we give to the object-involving state of affairs that we want to

re-describe as involving a different object? The existence of an object brings with it an object-involving state of affairs that is *distinct* and *in addition* to any other state of affairs. In some sense of ‘more’, there will be strictly more states of affairs if some *X*s compose than if they don’t compose. So, without the thing that the *X*s compose, there won’t be enough descriptions to accommodate this additional state of affairs. Unless we think that distinct states of affairs can have the same description, or we think that identity can be vague, I am doubtful that this strategy will work.²⁹ But, let’s set my skepticism aside and turn our attention to Hirsch’s second strategy.

According to strategy two, states of affairs simply don’t have structure. Thus, there are no states of affairs that involve objects. So, (44) is false. When we implicitly quantify over states of affairs with imperfect nominals like ‘Caesar’s dying’, we are quantifying over *unstructured* states of affairs.³⁰ Without (44), (VE) is consistent with (DSA) – vague existence does not require vagueness in the obtainment of states of affairs.

9 Vagueness in Worldly Structure

On both of Hirsch’s strategies, the argument that (VE) violates (DSA) fails. Either states of affairs don’t have structure or they have their structure indeterminately. This allows the deflationist to retain the claim that (determinately) every state of affairs that obtains determinately obtains. As far as I can tell, Hirsch’s strategy – especially his second strategy – succeeds in dodging violations of (DSA).

In picturesque terms, the core of Hirsch’s strategies is to deny that the world has determinate structure. That’s either because the world is unstructured

²⁹Much more would have to be said in order to make this argument rigorous – in particular we would have to get clearer on what we mean by there are ‘more’ states of affairs if some *X*s compose than if they don’t compose.

³⁰Contrast imperfect nominals with perfect nominals like ‘Caesar’s death’ which plausibly pick out events (CASATI AND VARZI, 2014).

(strategy two) or the world has structure but has that structure indeterminately (strategy one). Note, however, that denying that the world has determinate structure is *itself* a sort of metaphysical vagueness, even if it's not vagueness in which states of affairs obtains.

It was this indeterminacy that motivated the argument for violations of **(DSA)**. We treated states of affairs as determinately structured and then claimed that these structured states of affairs definitely obtain. But talk of states of affairs was nonessential to expressing the real source of indeterminacy that worried us: indeterminacy in the structure of the world. We can nominalize talk of structured states of affairs – an object's having of a property – by quantifying into both predicate and object position. Returning to the modal analogy one last time: contingency infuses these nominalized structured states of affairs – an object's instantiation of a property can be contingent. In higher-order modal logic we can express this as: $\Diamond(\exists X \exists x((\neg \Box X x) \wedge (\neg \Box \neg X x)))$. We can nominalize the claim that (determinately) the world has determinate structure as³¹

Determinate Worldly Structure (DS) $\Delta(\forall X \forall x(X x \rightarrow \Delta(X x)))$

Again, **(DS)** does not trivialize the claim of metaphysical vagueness. We can accept that it's indeterminate whether Harry is bald:

$$(47) \nabla(Bh)$$

without violating **(DS)** because we cannot quantify out (47) to deduce

$$(48) \exists X \exists x \nabla(X x)$$

In particular, there need not be a property that 'B' determinately refers to.³²

But **(DS)** entails that there is no vague existence:

$$i. \Delta(\forall x \exists X(X x)) \text{ (Assumption / } \mathbf{WVC})$$

³¹The idea of nominalizing states of affairs in this way in order to test for metaphysical vagueness comes from WILLIAMSON (2003).

³²C.f. WILLIAMSON (2003)

- ii. $\Delta (\forall X \forall x (Xx \rightarrow \Delta (Xx)))$ ((**DS**))
- iii. $\Delta (\forall x \exists X \Delta (Xx))$ (from (i) and (ii))
- iv. $\Delta (\forall X \forall x \Delta (Xx \rightarrow \exists y (y = x)))$ (Assumption / (**B**))
- v. $\Delta (\forall x \Delta (\exists y (y = x)))$ (from (iii) and (iv))

(i) claims that (determinately) every object instantiates some property. This follows from the (**WVC**) schema above and is independently plausible. (ii) is the statement of (**DS**): that (determinately) if something has a property, it's determinate that it has that property. While higher-order logic is controversial (even without indeterminacy operators), (iii) straightforwardly follows from (i) and (ii), so we can conclude that (determinately) everything has some property that it determinately instantiates. (iv) is the **Being** constraint which claims that definitely something must exist in order to instantiate a property. Together with (iii), the **Being** constraint straightforwardly entails (v), which contradicts (**VE**).

I see no way that the deflationist can avoid violations of (**DS**): they must admit that there is vagueness in the structure of the world. Indeed, I think many deflationists should *embrace* this result, as capturing the metaphysical heart of their position.³³ In his discussion of structured facts, Hirsch (2002) makes a gesture towards this picture of the world: “we can retain the notion of an unstructured fact. I think this is indeed our most basic notion of “reality,” “the world,” “the way it is,” and this notion can remain invariant through any changes in our concept of “the things that exist”” (79). While Hirsch admits that the world lacks determinate structure, he thinks that the *unstructured* nature of the world is determinate. We can express this unstructured determinacy by quantifying directly into sentence position.

Determinate Unstructured World (DUW) $\Delta (\forall S (S \rightarrow \Delta (S)))$

³³Many thanks to David Manley for pushing this line with me (and eventually convincing me that violations of (**DS**) is of a piece with their metaphysical view).

While deflationists must deny (DS), they can retain (DUW). This captures their commitment to the determinacy of the world's unstructured nature.

10 Conclusion

Ontological deflationists often market themselves as defenders of 'common-sense' metaphysics. It's thus important that their views vindicate a common-sense realism according to which the world is independent of the way we speak. Because deflationists tie questions of existence to our language and thought, opponents worry that the tie between the mind-independent world is *too close* – that it makes the world objectionably dependent on our language and thought. Some old-school deflationists, like Putnam, were happy to wear their anti-realism on their sleeves. More contemporary deflationists are not willing to give up the realism battle as easily as their predecessors. They are eager to show that their views on ontology are consistent with any world-independence thesis.

Flat-footed statements of realism are often phrased as modal independence theses: if we had spoken differently, such-and-such objects would have still existed. And, as ontological deflationists have correctly pointed out, their views are consistent with these modal independence theses. Other statements of realism appeal to controversial metaphysical gizmos like 'structure' or 'truth-making' which the deflationist can reject. Ontological deflationists are too slippery to stick them with an anti-realist thesis. For many though, a suspicion remains that ontological deflationists still tie the world too closely to our language.

I've suggested that one way to develop this suspicion is to see the extent to which the deflationist is committed to metaphysical vagueness. If their position leads to vagueness in the world as a result of vagueness in the word 'exists' that is a sort of anti-realism. Our investigation revealed the extent to which the deflationist must admit that there is vagueness in the world. They are committed to vague existence (regimented as (VE)) and this vague existence requires them

to accept vagueness in the structure of the world (regimented as (DS)). But, the extent of the metaphysical vagueness has a limit: the *unstructured* world remains determinate (regimented as (DUW)).

We can put the deflationist position loosely as: the *real* world is unstructured. Our results thus converge with the results of other studies of deflationism, like that of SIDER (2009, 2011), according to which deflationists claim that fundamentally the world doesn't have any quantificational structure. And we've done so without having to appeal to controversial metaphysical gizmos like structure or grounding.

Is the deflationist picture of an unstructured world one that we can ultimately accept? I'm not sure.³⁴ But as far as I can tell, evaluating this picture requires us to do more metaphysics. Thus, far from deflating metaphysics, ontological deflationism is just more metaphysics.

References

- Elizabeth Barnes. Arguments against metaphysical indeterminacy and vagueness. *Philosophy Compass*, 5(11), 2010a.
- Elizabeth Barnes. Ontic vagueness: A guide for the perplexed. *Noûs*, 44(4): 601–627, 2010b.
- Roberto Casati and Achille C. Varzi. Events. *Stanford Encyclopedia of Philosophy*, 2014.
- Matti Eklund. Being metaphysically unsettled. *Oxford Studies in Metaphysics*, 6, 2011.
- Kit Fine. Plantinga on the reduction of possibilist discourse. In James E Tomberlin and Peter van Inwagen, editors, *Alvin Plantinga*, pages 145–86. Springer, 1985. Reprinted in FINE 2005, 176–213.

³⁴See TURNER (Forthcoming) and SIDER (2011, §9.6.2 and §9.6.4) for some relevant considerations.

- Kit Fine. *Modality and Tense: Philosophical Papers*. Oxford University Press, 2005.
- Peter Fritz and Jeremy Goodman. Higher-order contingentism, part I: Closure and generation. *Journal of Philosophical Logic*, 2016.
- Katherine Hawley. Vagueness and existence. *Proceedings of the Aristotelian Society*, 102(1):125–140, June 2002.
- John Hawthorne. Superficialism in ontology. In Ryan Wasserman David Chalmers, David Manley, editor, *Metametaphysics*. Oxford University Press, 2009.
- Eli Hirsch. Objectivity without objects. *The Proceedings of the Twentieth World Congress of Philosophy*, 5, 1999. Reprinted in HIRSCH 2011, 36-44.
- Eli Hirsch. Quantifier variance and realism. *Philosophical Issues*, 12:51–73, 2002. Reprinted in HIRSCH 2011, 68-95.
- Eli Hirsch. Ontological arguments: Interpretive charity and quantifier variance. In *Metaphysical Debates*. Oxford University Press, 2008. Reprinted in HIRSCH 2011, 178-196.
- Eli Hirsch. *Quantifier Variance and Realism: Essays in Metaontology*. Oxford University Press, 2011.
- Eli Hirsch. Comments on Theodore Sider's *Writing the Book of the World*. *Philosophy and Phenomenological Research*, 87:733–754, 2014.
- Rosanna Keefe and Peter Smith, editors. *Vagueness: A Reader*. MIT Press, 1997.
- David Lewis. *On the Plurality of Worlds*. Blackwell Publishing, 1986.
- David Lewis. Evans misunderstood. *Analysis*, 48(3):128–130, June 1988. Reprinted in KEEFE AND SMITH 1997, 318-320.

- David Lewis. Many but almost one. In John Bacon Keith Campbell and Lloyd Reinhardt, editors, *Ontology, Causality, and Mind: Essays on the Philosophy of D.M. Armstrong*. Cambridge University Press, 1993. Reprinted in LEWIS 1999, 164-182.
- David Lewis. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 1999.
- Trenton Merricks. Varieties of vagueness. *Philosophy and Phenomenological Research*, 62(1):145-157, January 2001.
- W.V. Quine. Three grades of modal involvement. In *Proceedings of the XIth International Congress of Philosophy*, volume 14. North-Holland Publishing Co., 1953. Reprinted in QUINE 1966 (1976, 158-176).
- W.V. Quine. *The Ways of Paradox and Other Essays*. Harvard University Press, revised edition edition, 1966 (1976).
- R.M. Sainsbury. Why the world cannot be vague. *The Southern Journal of Philosophy*, 33:63-81, 1994.
- Theodore Sider. *Four-Dimensionalism*. Oxford University Press, 2001.
- Theodore Sider. Ontological realism. In Ryan Wasserman David Chalmers, David Manley, editor, *Metametaphysics*, pages 384-423. Oxford University Press, 2009.
- Theodore Sider. *Writing the Book of the World*. Oxford University Press, 2011.
- David E. Taylor and Alexis Burgess. What in the world is semantic indeterminacy? *Analytic Philosophy*, 56(4):298-317, December 2015.
- Richmond H. Thomason. Identity and vagueness. *Philosophical Studies*, 42: 329-332, 1982.
- Amie Thomasson. *Ontology Made Easy*. Oxford University Press, 2015.

- Jason Turner. Can we do without fundamental individuals. In Elizabeth Barnes, editor, *Current Controversies in Metaphysics*. Routledge, Forthcoming.
- Michael Tye. Vagueness and reality. *Philosophical Topics*, 28(1):195–209, 2000.
- Peter van Inwagen. *Material Beings*. Cornell University Press, 1990.
- Brian Weatherson. Many many problems. *Philosophical Quarterly*, 53(213), 2003.
- Timothy Williamson. Vagueness in reality. In *The Oxford Handbook of Metaphysics*, pages 690–716. Oxford University Press, 2003.
- Timothy Williamson. *Modal Logic as Metaphysics*. Oxford University Press, 2013.