

Heavy Ontology, Light Ideology

Jonathan Schaffer, Rutgers University
Draft of February 6th, 2018

[T]here is nothing in the logic of existential and universal quantification to tell us whether we should say that mereological sums exist or don't exist; nor is there some other science that answers this question. I suggest that we can decide to say either. (Hilary Putnam 2004: 38)

The plum survives its poems. (Wallace Stevens, *The Comedian as the Letter C*)

I think that there is a serious question as to whether the world hosts things such as protons, plums, and planets, but that there is no serious question as to whether the world prefers us to talk about such things using conjunction and negation, or disjunction and negation, or the Sheffer stroke. I thus uphold *heavy ontology* plus *light ideology*.

There is an ongoing debate between Hirsch (1993, 2002, 2005, 2008, 2009, 2011) and Sider (2001, 2007, 2009, 2011, 2014) over *quantifier variance*. Hirsch says that heavy ontology falls to quantifier variance, and Sider replies that heavy ontology can be propped up by heavy ideology (in particular, by a metaphysically distinguished quantifier). Both agree that heavy ontology plus light ideology is unstable.

I argue—*pace* both Hirsch and Sider—that heavy ontology plus light ideology is stable. In a sentence: I argue that quantifier variance is best understood as *domain variance*, that heavy ontology can be propped up by a metaphysically distinguished domain, and that positing a metaphysically distinguished domain puts no pressure on ideology since the domain is not part of the language (rather it is an object posited in the metalanguage to model the language). In a slogan: *heavy ontology needs a distinguished domain not heavy ideology*.

I am not arguing for heavy ontology. I consider it a plausible starting point, but its defense must be left for another day. Likewise I am not arguing for light ideology. I consider it plausible that there is no serious question as to which package of logical connectives the world prefers (an intuition that Sider (2011: 217) acknowledges), but have only an incredulous stare for those who say otherwise. I am just arguing that heavy ontology plus light ideology is stable.


My argument relies on two main semantic assumptions, which I should flag. First, I assume an orthodox semantics for quantification, on which the semantic contribution of quantificational expressions such as '∃' stems from two sources: *the clause* specifying their semantic values as quantifiers, and *the domain* over which the quantifiers range. Secondly, I assume an orthodox model-theoretic framework for semantics generally, in which the interpretation of a given object language takes place in a metalanguage via a *model*, which is a (*Domain, Interpretation*) pair. I may be understood as saying that heavy ontology plus light ideology is stable, given these orthodox semantic assumptions.

Roadmap: In §1 I review the Hirsch-Sider debate over quantifier variance, and clarify the positions on the table. In §2 I discuss the semantics of quantification, argue that quantifier variance should be understood as variance in the domain invoked in the metalanguage, and thus conclude that heavy ontology plus light ideology is stable. And in §3 I consider objections, touching on the Eklund-Hawthorne argument against quantifier variance, and the Lewisian posit of natural properties.

1. Quantifier Variance and the Threat to Stability: the Hirsch-Sider Debate

1.1 Heavy ontology meets quantifier variance

One of the main debates in metaphysics over the past few decades is the debate over mereology, spurred especially by van Inwagen (1990). As the textbook example goes, we are to imagine an otherwise empty box, containing three mereologically simple atoms (things with no proper parts):



We then ask: “How many things are in the box?” The nihilist only counts simples and answers: “*Three* things are in the box,” while the universalist also counts all composites and answers: “*Seven* things are in the box.” (There are further characters to consider, but these two will suffice for our purposes.)

Correspondingly, one of the main debates in meta-metaphysics is whether this debate over mereology is a “good debate.” Broadly speaking, *heavy ontology* is the meta-metaphysical view that this debate over mereology is a good debate, as are other similar first-order metaphysical debates. Heavy ontology is intended to correspond to a familiar pro-attitude towards these debates, and is the implicit meta-view of all those who naively engage in these debates, and the explicit view of metaphysicians including not just van Inwagen (1990), but also Bennett (2009), Eklund (2009), Hawthorne (2009), Sider (2009), and myself (Schaffer 2009, 2017), *inter alia*.

Light ontology is the opposing view that this and other similar debates are not good debates but in some way defective. It is intended to correspond to a familiar anti-attitude towards these debates, and stands as the explicit view of anti-metaphysicians such as Carnap (1950), Putnam (1994, 2004), and Hirsch (2009), as well as Thomasson (2007, 2009) and Chalmers (2009), *inter alia*.¹

I leave the notion of a “good debate” intentionally vague, since heavy ontology and light ontology are each best understood as broad clusters of views, encompassing different philosophers with different background views. I place no weight on the notion. Rather I am interested in a specific Carnap-Putnam-Hirsch *quantifier variance* tradition within light ontology, so the vagaries of “good debate” may be set aside.

The quantifier variance tradition is rooted in Carnap’s (1950; for discussion see Eklund 2009) dismissal of metaphysical debates as mere confusions about language. For Carnap (1950: §3), there is a “practical” question of whether adopting a certain linguistic framework is “more or less expedient, fruitful, conducive to the aim for which the language is intended.” But there is no “theoretical” question of whether the entities the framework involves have any external reality, since any “alleged statement of the reality of the system of entities is a pseudo-statement without cognitive content.”

Putnam (1994, 2004; for discussion see Eklund 2008) takes up Carnap’s dismissal of metaphysical debates as confusions about language, drawing specific connections to quantification. In particular, Putnam (2004: 37) says that “to ask whether mereological sums *really exist* would be stupid” since “it is a matter of convention whether we say that mereological sums exist or not.” He (2004: 38; cf. 1987: 71) continues:

[I]here is nothing in the logic of existential and universal quantification to *tell us* whether we should say that mereological sums exist or don’t exist; nor is there some other science that answers this question. I suggest that we can *decide* to say either.

Thus, for Putnam, there is the “nihilist use” and the “universalist use” of quantification, each of which is consistent with the introduction and elimination rules that constitute the logic of quantification. There is then a merely conventional decision to make, as to which manner of speaking to adopt. Putnam thus delivers

¹ There is no fixed terminology in the literature. Heavy ontology is intended to correspond to Chalmers (2009) notion of heavyweight realism, and Sider’s (2009) notion of ontological realism. Light ontology is intended to cover what Eklund’s (2009: 137) calls “the Carnapian dismissal of ontological questions as shallow,” Hawthorne’s (2009) notion of superficialism, and Hirsch’s (2009: 231) claim that these debates are not substantive but “merely verbal.” The heavy/light opposition is also intended to correspond to what Chalmers, Manley, and Wasserman (2009: 4) label “mainstream metaphysics” versus “deflationism.”

(2004: ch. 4) “an obituary” for ontology, asserting that the Quinean (Quine 1948) who posits a single fixed sense of ‘exists’ is “already wandering in Cloud Cuckoo Land.”

Hirsch (1993, 2002, 2005, 2008, 2009, 2011) further develops and now champions the quantifier variance view, taking (2002: 51) Putnam’s insight to be that “[t]he quantificational apparatus in our language and thought—such expressions as ‘thing’, ‘object’, ‘something’, ‘(there) exists’—has a certain variability or plasticity.” Hirsch allows that there may be a fact of the matter as to “which sums exist” *given a fixed language* (e.g., English), but says that this at most reflects a shallow feature of the given language, not a deep feature of reality. One can imagine two tribes who speak homophonic languages, with one speaking like nihilists and the other like universalists. Hirsch (2011: 86) says that principles of charity would require us to interpret both as speaking truly in their respective languages, with different semantic rules for quantification:

The different semantic rules that would have the effect of rendering the sentence true in one language and false in the other must in some sense provide different rules for “counting what things there are in the world.” If there could be these two languages they would have to embody in some sense different concepts of what it is “for there to exist something.”

And so Hirsch says that the debate over mereology is either shallow (for a fixed language) or verbal, in that (2009: 239) “each party ought to agree that the other speaks the truth in its own language.” The nihilist and universalist tribes are both right.²

Textual matters aside, the view I am interested in—the quantifier variance view—says that there is no serious question as to whether the world hosts composite things like protons, plums, and planets, because there are many different “quantifier meanings,” including the nihilist and universalist meanings. It is at most a shallow question of language which quantifier meaning we invoke. And so it is at most a shallow question of language as to whether the nihilist or the universalist (or neither) is correct, not a deep question about reality.

So far I am merely reporting the relevant core of the quantifier variance view. I am not endorsing this view (indeed I think it misunderstands quantification). I am just setting the stage.

1.2 *Quantifier variance meets heavy ideology*

Enter Sider: Sider is a heavy ontologist who thinks that quantifier variance is a compelling challenge, and who argues that the threat of quantifier variance can only be averted by heavy ideology. While “ontology” concerns what there is, “ideology” concerns which terms we use.³ *Heavy ideology* is the view that debates over which terms to use are good debates; while *light ideology* is the view that debates over terms are not good debates but in some way defective. While there has been much less discussion of ideology, Sider (2011: vii) is the champion of the heavy view that “ideology matters” and questions of “whether causal predicates, quantifiers (or names), and modal operators carve at the joints... lie at the center of meta-metaphysics.”

Light ideology may take many forms, but the form I prefer is a kind of *ideological pragmatism*, on which—to echo Putnam’s language for ontology (§1.1)—it is “a matter of convention” which terms we use, and a matter that “we can decide” at our convenience. There are metaphysical constraints on overall expressive power. For instance, it would be inadequate to speak a language with just conjunction. And there are psychological reasons for us to prefer certain terms. For instance, it seems more natural for us to use conjunction and negation than to use the Sheffer stroke. But—the ideological pragmatist continues—

² While Carnap and Putnam cast themselves as anti-metaphysicians, Hirsch takes a more nuanced view in several respects. As Hirsch (2002) points out, his position is in many ways “realist.” Also Hirsch (2009) thinks that light ontology is appropriate for metaphysical debates *under the substantive further condition that each side can plausibly interpret the other as speaking truly*. Hirsch thinks this further condition is met in the debates over mereology and temporal ontology, but he is open to heavy ontology for Gods and numbers.

³ The distinction between ontology and ideology comes from Quine (1951: 13–14) who says that ontology “asks what there is” while ideology concerns, for a given theory, “what ideas can be expressed...” As Quine (1983: 501) later put it, ideology is “one’s stock of simple and complex terms or predicates.”

the world does not care as between logically equivalent packages of expressions, such as conjunction and negation, or disjunction and negation, or the Sheffer stroke. Likewise, the world does not care whether we use existential or universal quantification, or necessity or possibility operators (both pairs are pairs of inter-definable duals). As McSweeney (*forthcoming*) puts the point:

The logical realist must, it seems, distinguish between two otherwise equivalent theories, T, which employs ‘ \forall ’, ‘ $\&$ ’, and ‘ \sim ’, and T’, which employs ‘ \exists ’, ‘ \vee ’, and ‘ \sim ’. But something has gone wrong if we are in a position in which we are forced to ask and answer which of these theories is joint-carving, indeed if we are forced to think there is any worldly difference between them at all. The theories seem to be paradigmatic *mere notational variants*...

Again I am using “good debate” in an intentionally vague way, and place no weight on the notion. Rather I am interested in Sider’s specific reply to quantifier variance. (And ultimately I am interested in exhibiting an alternative reply on which quantifier variance can be averted, with no pressure on ideology.)

So Sider’s (2001, 2007, 2009, 2011, 2014) idea is to answer quantifier variance by extending the Lewisian notion of *naturalness* “beyond the predicate” into logical terms, and then to speculate that there is a distinguished “quantifier meaning,” more natural than any other by a wide enough margin.⁴ This distinguished quantifier meaning serves as an external constraint on semantic interpretation, putting overriding pressure on us to interpret metaphysical disputants (and members of rival tribes) not just charitably but *eligibly*, for using their idioms of quantification in the metaphysically distinguished way. The most natural meaning is said to serve as a “reference magnet.” As Sider (2014: 565) usefully summarizes:

i) [T]here are joint-carving meanings that are suitable to be meant by quantifiers; ii) Lewisian reference magnetism is true; and iii) charity is trumped by the eligibility of an interpretation that assigns the joint-carving meanings to the quantifiers.

Though it should be noted that Sider (2014: 565) allows a second reply (“the Ontologese gambit”) which sacrifices to Hirsch any claims about natural language, but takes the line that metaphysicians can (and perhaps already have) stipulated that their claims are couched in a language (“Ontologese”) exactly like English save that the quantifiers are stipulated to have their metaphysically distinguished meanings. Both of these replies require that quantifiers can have metaphysically distinguished meanings (they merely differ over whether natural language quantifiers bear them already). This reply leads Sider (2011: ch. 10) to think that there is a worldly preference for distinguished ideology generally, and so a serious question as to whether the world prefers conjunction and negation, or disjunction and negation, or the Sheffer stroke, and a serious question as to whether the world prefers existential or universal quantification, etc.⁵

I am not here concerned with the viability of Sider’s reply to Hirsch, or the twists and turns of their ongoing debate. (I think their debate set off in the wrong direction.) I am just reporting their starting point.

1.3 The quest for stability

What is especially noteworthy—and troubling to those like myself whose starting point is that of heavy ontology plus light ideology—is that Hirsch and Sider both think that “quantifier meanings” are the crux of the dispute. Indeed Sider (2009: 397) calls this “the central question of metaontology,” saying:

⁴ Lewis (1983, 1984) invokes naturalness for properties, in part to answer Putnam’s (1981) attack on external realism. See Williams (2007) for further discussion of reference magnetism. See §3.2 for further discussion on Lewisian natural properties, given light ideology.

⁵ McDaniel (2009, 2017; see also Turner 2010) builds on Sider’s ideological realism to claim that there is a substantive debate as to whether there are multiple ways of being, in terms of whether the world prefers a single existential quantifier ‘ \exists ’ or some plurality of restricted existential quantifiers ‘ \exists_1 ’-‘ \exists_n ’ instead, given the truth of ‘ $(\exists x)\Phi \equiv ((\exists_1 x)\Phi \vee \dots \vee (\exists_n x)\Phi)$.’ Thus he (2017: 37) says: “[O]ne believes in ways of being if one believes that there is more than one *relatively fundamental* meaning for an existential quantifier.”

[Q]uantifier variance remains the crux. The central question of metaontology is that of whether there are many equally good quantifier meanings, or whether there is a single best quantifier meaning. It is a question about nature’s joints; it is a question of *how much quantificational structure the world contains*.

Thus both Hirsch and Sider deny:

Stable: Heavy ontology plus light ideology is stable

Hirsch and Sider merely draw different conclusions from the denial of *Stable*. Hirsch rejects heavy ontology, while Sider rejects light ontology. In other words, both Hirsch and Sider agree that, if heavy ontology can be saved at all, it is via heavy ideology. They merely disagree on the further question of whether heavy ontology can be saved at all, even in this way.

I aim to uphold *Stable*. Thus I am not arguing for heavy ontology or for light ideology, but only for the claim that they can be stably conjoined. My stance to both Hirsch and Sider is in this respect akin to Sider’s stance to Hirsch, which Sider (2014: 566) explains “isn’t meant to convince Hirsch or anyone else with neo-Carnapian tendencies.” Instead:

It’s rather supposed to be a stable position from which one can resist neo-Carnapian arguments. It’s supposed to have some independent appeal; and it’s supposed to undermine the arguments *if its metaphysical assumption of realism about joint-carving is true*. (If this assumption is false then some form of neo-Carnapianism might well be correct.)

I too am trying to state a stable position with independent appeal—a position from which one can not only resist Hirsch’s neo-Carnapian arguments, but also Sider’s “neo-Lewisian” arguments.

I am also not arguing that my position is the *only* way in which heavy ontology and light ideology can be stably conjoined. (I will mention an alternative in §2.3.) I only mean to say that it is *one* way.

Putting the positions on a map:

	<i>Light Ontology</i>	<i>Heavy Ontology</i>
<i>Light Ideology</i>	Hirsch	Schaffer
<i>Heavy Ideology</i>	X	Sider

The bottom-left box—LO-HI—is to my knowledge unoccupied: graduate students, attend! I stand in the top-right box: I defend the HO-LI ground!

2. Quantifier Variance as Domain Variance: Stability Regained

2.1 Two sources of “quantifier meanings”

Hirsch and Sider both speak of variation in *quantifier meanings*, such that “Three things are in the box” may be true in a nihilist language, and “Seven things are in the box” may be true in a universalist language (§1). But it is not clear what a “quantifier meaning” is, and how these may vary. Quantifiers are not isolated atoms of meaning. Rather—as per the first of my two main assumptions flagged in the introduction—the semantic contribution of quantificational expressions stems from two sources: *the clause* specifying their semantic values as quantifiers (which is a thin logical meaning), and *the domain* over which the quantifiers range (which is not specific to the quantifier or even an element of the language at all, but an object posited in the model used to interpret the whole language). Fix the clause and the domain, and no room for variation remains. Thus I say:

Sources: Any quantifier variance must stem from variance in the clause or the domain

As an initial rough illustration of *Sources*, when the existential quantifier closes a formula, the result is true if and only if something in the domain satisfies the formula. The “something in the domain” condition is the thin logical meaning of the clause, drawing on the background domain. Fix that clause and the background domain being drawn on, and no room for variation remains.

The rationale for *Sources* can be made more precise by considering Tarski’s early conception of quantification (Tarski 1983 [1933]), and at modern generalized quantifiers.⁶ In Tarski’s early conception, we start from a domain, and an assignment mapping variables to individuals in the domain. Tarski then states a recursive definition of when an assignment satisfies a formula, including the following clause for existentially quantified formulae:

Existential-Tarski: Assignment a satisfies ‘ $(\exists x)\Phi$ ’ iff: there is an assignment β that is an x -variant of a , such that β satisfies ‘ Φ ’

(Assignment β is an x -variant of a iff: a and β map all variables other than x to the same individuals.)

What is relevant about Tarskian semantics is that, given a domain and the clause *Existential-Tarski* saying when an assignment satisfies an existentially quantified formula, there is no wiggle room left for variation. To illustrate, consider ‘ $(\exists x) Fx$.’ Suppose that the domain contains three elements: c , d , and e . And suppose that the extension of the predicate ‘ F ’ is such that ‘ Fc ’ is true but ‘ Fd ’ and ‘ Fe ’ are false.⁷ Now we can prove that assignment a satisfies ‘ $(\exists x) Fx$ ’ by considering the assignment β that is an x -variant of a and maps x to c . Thus Tarskian semantics upholds *Sources*.

In the modern conception of generalized quantifiers, we see at least three modifications of the Tarskian machinery.⁸ First, the background machinery now is the model-theoretic machinery of a domain plus an interpretation function mapping predicates to extensions from the domain. Schematically:

Model = (*Domain*, *Interpretation*)

I pause on this point, since—as per the second of my two main assumptions flagged in the introduction—in model-theoretic semantics the domain is not part of the object language under study. It is rather an object posited in the metalanguage by an interpreter modeling the object language. This will matter later (§2.3). Secondly, the structure and range of quantifiers is richer. Instead of Tarski’s pair of binary constructions ‘ $(\exists x)\Phi$ ’ and ‘ $(\forall x)\Phi$ ’, we have a full range of ternary constructions of the schematic form:

[[*Quantifier*] *Restrictor*] *Scope*

That is, all quantifiers have a slot for a restrictor argument (which can but need not be filled by an indiscriminating condition, such as *being self-identical*). And we do not merely treat the two quantificational expressions ‘some’ and ‘all,’ but also ‘most,’ ‘few,’ and ‘between ten and seventeen’ etc., in a unified framework. Thirdly, the inputs to quantifiers—for compositional reasons—are a pair of properties: the restrictor and scope argument slots are both saturated by properties.

⁶ I work with Tarski’s early conception of quantification since it is usually invoked in the philosophical discussion (indeed van Inwagen (2014: 1–3) dubs the language of ontology “Tarskian;” see also Hirsch & Warren *forthcoming*: 1), and with generalized quantifiers since they are standard in natural language semantics.

⁷ In Tarski’s 1933 conception we do not yet see the model-theoretic idea of an interpretation function specifying the extensions of predicates, but rather this is treated as part of the fixed background. In any case these matters concern the meanings of predicates not quantifiers. For more on predicates see §3.2.

⁸ The fundamental idea of generalized quantifiers comes from Mostowski (1957) and Lindström (1966), and its application to natural language semantics comes through Barwise & Cooper (1981).

Against these background modifications, generalized quantifiers have the thin logical meaning of expressing *comparisons*. Specifically they compare the extensions of the restrictor and scope arguments, and ask whether *Quantifier*-many members of the restrictor set are members of the scope set. Schematically:

Template-Generalized: $[[Q(R)(S)]]_M = \text{True}$ iff Q -many members of R are members of S

(‘ $[[\dots]]$ ’ is the semantic value operator, subscripted to a model.) One can then further characterize specific quantifiers in terms of set-theoretic requirements. For instance:

$$\begin{array}{lll} [[\textit{all} (R)(S)]]_M = \text{True} & \text{iff} & R \subseteq S \\ [[\textit{some} (R)(S)]]_M = \text{True} & \text{iff} & R \cap S \neq \emptyset \\ [[\textit{ten} (R)(S)]]_M = \text{True} & \text{iff} & |R \cap S| = 10 \\ [[\textit{most} (R)(S)]]_M = \text{True} & \text{iff} & |R \cap S| > |R - S| \end{array}$$

(‘ $|\dots|$ ’ is a cardinality operator.)

So a fixed domain (alongside a fixed interpretation, to set the meanings of the relevant predicates) fixes the input sets. And a fixed clause fixes the set-theoretic comparison. No wiggle room remains. Thus generalized quantifiers uphold the claim that fixing the clause and the background domain fixes the semantic contribution of the quantificational expression. To illustrate, consider ‘Some student smiles’ in a given model. One first looks to the domain (given an interpretation) for the sets *Smilers* and *Students*. Then ‘some’ expresses the thin logical requirement that the intersection of these sets be non-empty, thereby requiring that there be an individual in the domain that is both a student and a smiler. Fixing the domain (and interpretation) fixes the sets *Smilers* and *Students*. Fixing the clause fixes that ‘some’ requires non-empty intersection. It simply remains to ask whether or not $\textit{Smilers} \cap \textit{Students} = \emptyset$.

So I have clarified *Sources*, and shown how it fits with both the Tarskian conception of quantification, and the modern conception of generalized quantifiers (within a model-theoretic backdrop). I pause to address three issues arising. The first issue concerns metasemantics. It is sometimes said that the distinction between atomistic (or “bottom-up”) and holistic (or “top-down”) metasemantics is relevant to quantifier variance. Indeed Warren (2017: 88) says that “Some bottom-up approaches will simply take the facts about things like quantifier domains as given,” but:

In contrast, top-down theorists think that facts about a quantifier’s domain are explained by facts about the semantic properties of whole sentences involving the quantifier, typically the truth conditions of these sentences. This is where Putnam and Hirsch step in.

Warren (2017: 91) goes on to claim that Putnam and Hirsch’s “quantifier deflationism” follows from “top-down” plus “quantifier inferentialism” (the view that what makes a particular expression take the meaning of a quantifier is the inferential role it plays). So it might be thought that I have in some way assumed an atomistic (or non-inferentialist) metasemantics and thus *begged the question* against Putnam and Hirsch. Or in general one might wonder how to think about the connections between *Sources* and metasemantical issues.

I say that, while issues of metasemantics may be relevant to some aspects of quantifier variance, they are not relevant to *Sources*. All that matters for my purposes is the semantics of quantifiers, not their metasemantics. Everyone should agree (as a semantic fact) that various expressions, including the English expressions ‘some,’ ‘most,’ and ‘all,’ have the meanings of quantifiers. I have argued that this supports *Sources*. There is a further question as to how and why expressions come to be infused with their semantic values, and so a further question as to how and why particular expressions like the English expressions ‘some,’ ‘most,’

and ‘all’ come to have the meanings of quantifiers. But with respect to supporting *Sources*, all that matters is the meanings, not the meta-question of how and why expressions come to have these meanings.⁹

A second issue arising is that the semantics for quantification—in both the Tarskian conception, and with generalized quantifiers—uses quantification in the metalanguage. Thus Tarski quantifies over assignments (*Existential-Tarski* uses ‘there is an assignment β ’), and generalized quantifiers are treated using metalinguistic quantification as well (*Template-Generalized* uses ‘ Q -many,’ and the set-theoretic expressions use the quantificationally-loaded ideology of set theory). This is unavoidable for two convergent reasons. The first reason is that the metalanguage is a language itself, and needs the resources of quantification. (Indeed the metalanguage is supposed to faithfully embed an image of the object language, and just to add the additional resources needed to talk about semantic values of object language expressions.) The second reason is that semantics is generally not in the business of reductions. Instead one typically sees homophonic clauses like:

$$[[\textit{smokes}]]_M = \lambda x. x \textit{ smokes}^{10}$$

So it should come as no surprise to see homophonic clauses for quantifiers such as:

$$[[\textit{some}]]_M = \lambda R S. \textit{ some } (R, S)$$

This captures the comparison ‘some’ draws between properties, which is simply that it requires that some elements of the restrictor set be in the scope set (in accord with *Template-Generalized*).

So it might be thought that I have simply *postponed* the real issue of quantifier variance, which just re-arises anew in the metalanguage. Or in general one might wonder how to think about issues of quantifier variance in light of the fact that the semantics of quantifiers is treated via metalinguistic quantification.

I say that every claim made about the semantics iterates into the metalanguage (and into the meta-metalanguage, etc.) If one wants to understand the quantificational expressions in the metalanguage, one need only ascend to the meta-metalanguage. The quantificational expressions in the metalanguage will then be treated via quantification in the meta-metalanguage, these quantificational expressions will in turn be treated via quantification in the meta-meta-metalanguage, *ad infinitum*. Crucially, *Sources* holds at every level. (There is no getting underneath quantification.) So either domain variance shows up at some level, or clausal variance shows up at some level, or else *quantifier meanings are fixed all the way up the hierarchy of metalanguages*. (In my preferred view there is metaphysically distinguished domain, and so domain realism holds sway all the way up the hierarchy of metalanguages: §2.3.)

2.2 Why domain variance is the relevant sort of variance

So far I have argued—as per *Sources* (§2.1)—that any quantifier variance must stem from domain variance or clausal variance. I now argue that any clausal variance is irrelevant to heavy ontology, and that the relevant sort of variance is domain variance. Starting with the irrelevance of clausal variance, the core idea is that the clause links an expression to its semantic value. So varying the clause of a given quantificational expression does not vary the meaning of any given quantifier, but merely varies which quantifier (if any) the expression denotes. To illustrate, one can of course consider swapping the semantic clause of ‘some’ with that of ‘many,’ ‘snow,’ or ‘guacamole.’ But that mundane fact—which none would contest—only shows that the word ‘some’ need not have the meaning of the existential quantifier at all. It does not show any possible variation *within* the meaning of something that is an existential quantifier (for that we will need domain variance).

⁹ So does a holistic metasemantics plus quantifier inferentialism lead to quantifier variance? I have no idea. But either this combination of views gets it right that expressions like the English expressions ‘some,’ ‘most,’ and ‘all’ are quantifiers, or not. If so then it leads to *Sources* and I may proceed. If not then it fails to yield correct semantic values and stands refuted.

¹⁰ Object language expressions are italicized to distinguish them from their meta-language homophones. So the left side of the equality is to be read: “The semantic value of ‘smokes’ in the object language, according to model *M*...” The right side of the equality then reads “Being an *x* such that *x* smokes.”

As a first confirmation that clausal variance is not at issue, Putnam and Hirsch are explicit that the logic of quantification is to be held fixed. Thus Putnam (2004: 37) clarifies: “These uses are not totally different; for example, in all of its uses the existential quantifier obeys the same logical laws...” Likewise Hirsch (2002: 53) says: “The purely syntactic and formal logical properties of [the existential quantifier] will not be changed at all (the formal principles of quantificational logic will be unaltered).” Presumably Putnam and Hirsch say this because they think—in at least partial accord with quantifier inferentialism—that it is at least a necessary condition on a term’s being a given quantifier that it obeys the logic of that quantifier.

But clausal variance would not preserve the logic. The thin logical meaning of ‘some’ is merely some, or (more formally) in generalized quantifiers (§2.1):

$$[[some]]_M = \lambda R.S. \text{ some } (R, S)$$

Anything that is not comparing two properties is not a generalized quantifier in structure, and anything that does not compare them by requiring that some R be an S (equivalently, that $R \cap S \neq \emptyset$) is not the existential quantifier and so not bound by the existential introduction and elimination rules. (If the requirement allows satisfaction in a case when no R is an S , the elimination rule will not be truth-preserving. If the requirement denies satisfaction in a case when some R is an S , the introduction rule will not be truth-preserving.)

A second way to confirm that clausal variance is not at issue is to imagine a language just like English except that ‘some’ is associated with a different semantic clause, and then to imagine augmenting this language with a new term ‘schmome’ taking the old clause:

$$[[schmome]]_M = \lambda R.S. \text{ some } (R, S)$$

It should then be clear that it is their ‘schmome’ (not their ‘some’) that means what our ‘some’ means. No clausal variance remains: ‘schmome’ has the same semantic clause as ‘some’ by construction.

A third way to see that clausal variance is not at issue is to return to the mereology debate we began with (§1.1). If quantifier variance were mere clausal variance, then the quantifier variantist position would not speak against a fixed domain, which—let us say—includes composites. But then the nihilist who denies that there are seven things would be interpreted on a fixed domain of seven things, but with blinkered idioms of quantification unable to see the whole domain. She would look like someone who only had a “restricted quantifier.”¹¹ This image of the nihilist as recognizing seven things in the box, but merely lacking the quantificational resources to speak of the four composites, is not a faithful image of the nihilist. We were supposed to find a straightforward nihilist using the existential quantifier in standard ways. Instead we have found a twisted nihilist who is just like a universalist, save that she cannot express unrestricted quantification.

Such a twisted nihilist will, moreover, have trouble upholding the logical rules for her quantifiers. For if she has the predicate ‘Composite’ (interpreted as intended), and has names—say ‘Ann’—for one, then ‘Composite(Ann)’ will be true in her language. But the associated existential generalization ‘ $(\exists x)\text{Composite } x$ ’ will be false, because her existential quantifier cannot see the composites. (The underlying problem is that the constants are still getting their denotations from the domain, and the predicates are still getting their interpretations drawn from the domain, so clausal variance for the quantifiers is going to do nothing to the semantics of names or predicates. Variation in the quantifiers must be matched with variations in names and predicates, and so it is really the underlying domain that needs to vary.)

¹¹ With generalized quantifiers, there is no difference between a restricted and unrestricted quantifier, but only a difference in what goes in the restrictor slot. So to operate without an “unrestricted quantifier” would be to operate without a means to express an indiscriminating restrictor (such as *being self-identical* or *being such that $2+2=4$*). This would be a radically impoverished language, whose defect has nothing to do with quantifiers, but rather traces to the absence of any indiscriminating predicate.

So I conclude that quantifier variance, in the sense relevant to heavy ontology, cannot be clausal variance. Given *Sources*, it follows that the relevant sort of quantifier variance must be domain variance. For instance, in the mereology debate, the quantifier variantist should view the nihilist and the universalist as best interpreted as using an existential quantifier with the same thin logical clause, but each drawing on different domains. A fitting model for the nihilist's language has a domain without composites, while a fitting model for the universalist's language has a domain with composites.¹²

I offer three points of confirmation for the idea that quantifier variance is best understood as domain variance, the first of which is that it fits the quantifier variantist's core idea that the nihilist and the universalist are each speaking truly in their respective languages. 'There are no composites' is true relative to a model whose domain has no composites, and 'There are composites' is true relative to a model whose domain has composites. Secondly, domain variance does not affect the logical laws, and so fits the quantifier variantist's idea (Hirsch 2002: 53) that "the formal principles of quantificational logic will be unaltered." Thirdly, and perhaps most importantly, only domain variance captures the quantifier variantist's idea that the semantics of names and predicates is going to have vary alongside the variance with the quantifier. The domain is a kind of semantic "master switch"—names take their denotations from the domain, predicates take their extensions from the domain, and quantifiers take their ranges from the domain. So shifting the domain manages to keep names, predicates, and quantifiers properly coordinated. Here at last is what seems like a viable way to understand how the universalist and nihilist could both be using existential quantifiers and speaking truly in their respective and otherwise so similar languages.

2.3 How to avert quantifier variance without heavy ideology

I have argued that quantifier variance is best understood as domain variance. This allows me to defend *Stable* by arguing that heavy ontology can be propped up by a metaphysically distinguished domain, and that positing a metaphysically distinguished domain puts no pressure on ideology.

Working backwards, I begin with the point that domain variance has nothing to do with ideology. Indeed domain variance has nothing specifically to do with quantification. Rather in orthodox model-theoretic semantics we assign semantic values to expressions in a given "object language" by moving to a metalanguage, and interpreting the object language via a model. Schematically, as we have seen:

$$\text{Model} = (\text{Domain}, \text{Interpretation})$$

Every bit of the object language—not just its quantificational expressions, but its names and predicates equally—is interpreted with reference to the model. So the label 'quantifier variance' proves misleading. Hirsch and Sider should have spoken, not of variation in quantifier meanings, but of variation in the model. The model is not part of the object language being interpreted, but part of the metalanguage in which the interpretation is given. So in that sense 'quantifier variance' is doubly misleading, for pointing at variation in the meaning of the *quantifiers* in the *object* language, rather than variation in the *model* in the *metalanguage*.¹³

¹² This perspective is anticipated by Koslicki (2003: 118):

The universalist and the holder of the intermediary position can, I think, agree on the meaning of the existential quantifier, in the sense that they can agree on which logical operation is denoted by the symbol '∃'... But settling on the meaning of the existential quantifier by itself does not settle what its range is: two philosophers can perfectly well agree on what the symbol '∃' means, while still carrying on a thoroughly sensible dispute over the size and the nature of the domain of quantification...

¹³ On this point I agree with Finn & Bueno (*manuscript*: §2.1), who write: "[D]omain shifting is insufficient for quantifier variance, given that it only marks a difference in quantificational scope, whilst in each domain the quantifiers still operate in the same way: the universal quantifier ranges over all objects of the domain, the particular over some." I take this to show that 'quantifier variance' is a poor label. (Thanks to Suki Finn for discussion of these points.)

This understanding of the domain as a part of the model posited in the metalanguage is the second of my two main assumptions flagged in the introduction. There are many reasons for separating the object language and the metalanguage, ranging from Tarski's (1983) idea of avoiding liar paradoxes via a hierarchy of 'truth' predicates, to the more pedestrian reason that a single given object language can be interpreted via multiple models. Thus contemporary semantics textbooks usually invoke notational conventions such as that of writing object language expressions in italics, to separate them from their metalanguage homophones, and relativizing semantic values to models, as we have already seen in:

$$[[\textit{some}]]_M = \lambda R S. \textit{some} (R, S)$$

My point is not merely to re-label quantifier variance as domain variance, but to use this better understanding (drawn from *Sources* and the denial of clausal variance: §§2.1–2.2) to show how heavy ontology plus light ideology is stable (as per *Stable*: §1.3). The key point is two-fold. First, *the dispute over domain variance is not a dispute over ideology*. The domain is not part of the object language, but rather lives in the metalanguage (as an entity: a set). So whatever one might think of heavy ideology (worldly preference for terms), it simply has *no bearing* on domain variance.

Secondly, the dispute over domain variance is naturally understood as a dispute about the ontology of the theorist interpreting the language. But at this stage a straightforward heavyweight stance is (not mandatory but) available:

Domain realism: There is a metaphysically distinguished domain

I cannot specify what this domain is without resorting to quantification (one should never expect to get under quantification), but I do not myself need to be able to specify what this domain is for the thesis to be true. Of course the metaphysically distinguished domain is intended to contain *everything*. That is, it is supposed to correspond to a distinguished totality of things, which may or may not contain composites.¹⁴

What it means for there to be a *metaphysically distinguished* domain is for there to be worldly pressure on interpreters to adopt the right model. *Domain realism* says that a theorist interpreting a given object language must specify a (*Domain, Interpretation*) pair *with the distinguished domain*. The theorist who adopts a metalanguage with a non-distinguished domain misinterprets the object language(s) she studies.

To put this point another way: models are mathematical representations. Whenever we indulge in representation, there is always a question of whether the representation is *apt*. For instance, we can model English by taking the domain to be the set of bananas, just so long as there are enough bananas to go around. But this presumably would not make for an apt representation of English. *Domain realism* says that it is a necessary condition on a model's being apt that it invokes the distinguished domain.

¹⁴ Here paradoxes lurk. If sets are in the domain, and the domain itself is a set, it must contain itself and it must contain all sets (both of which are barred in orthodox set theory). There are technical fixes including shifting to less orthodox set theories that consistently allow for such behavior, or shifting away to pluralities. Obviously I cannot engage with these issues here (see the papers collected in Rayo & Uzquiano 2006 for further discussion). But there is one approach which directly intersects the current discussion, which is that of *indefinite extensibility* (Dummett 1991: 316–9), on which it is said that there is no ultimate totality but always the opportunity to enlarge any given domain. Warren (2017) offers a model of indefinite extensibility in terms of domain variance (he calls it “quantifier variance” but makes clear (2017: 83) that it is the domain that matters). Given indefinite extensibility, I think that one should modify *Domain realism* to say that, for every stage *s*, there is a worldly preference for a distinguished domain *D_s*. After all, indefinite extensibility could still be thought of along broadly heavyweight lines, if the “avenues of extensibility” are thought to be fixed and objective. Also one could endorse indefinite extensibility in the set theoretic case while thinking that there is a heavyweight fact of the matter in the mereological debate. But this is not a matter I can pursue further here, so in the main text I adopt the orthodox expedient of supposing that there is a total domain.

To see how this works, suppose that the metaphysically distinguished domain in fact includes all the composites, and consider a tribe whose members speak like nihilists, saying things like ‘Three things are in the box.’ We then come to interpret them. *Domain realism* compels us to interpret them via a model with the distinguished domain. By the thin logical meaning of ‘three’ (and the intended interpretations of ‘things’ and ‘in the box’), their sentence is true if and only if the cardinality of the intersection of the set of things and the set of box-dwellers is three:

$$[[\textit{Three things are in the box}]]_M = \text{True} \quad \text{iff} \quad |\{\text{Things}\} \cap \{\text{Box-dwellers}\}| = 3$$

But given that the distinguished domain includes the composites, the cardinality of this set is not three but seven. So according to *Domain realism*, nihilists—even whole tribes of them—speak falsely if the world hosts composites. (Likewise, universalists—even whole tribes of them—speak falsely if the world hosts no composites.) Thus *Domain realism* befits heavy ontology.

Four clarifications: First, domain realism works differently from Sider’s quantifier realism, which seeks a privileged object language (such as English or “Ontologese”: §1.2). What it means for there to be a metaphysically distinguished domain is not for the world to privilege certain terms for speakers, but rather for the world to privilege certain models for interpreters (strictly speaking: to privilege those (*Domain, Interpretation*) pairs in the metalanguage that use the distinguished domain). In short: there is no distinguished “language of ontology” but rather a privileged interpretive stance on languages.

Secondly, the issue is not whether to interpret the nihilists as using a “restricted quantifier” but rather whether to aptly model them using one domain or another. Restricted quantification is a local object language phenomenon, occurring when a discriminating predicate sits in the restrictor slot of a given occurrence of a generalized quantifier, as in ‘Three simples are in the box’ (which is uncontroversially true, but a different sentence from ‘Three things are in the box’). Apt modeling is a metalanguage phenomena concerning how best to interpret a given object language *in toto*.

I am saying that the quantifier variantist should be a domain variantist, and should say that it is better to model the nihilists by stating an interpretation in the metalanguage using the domain of only simples (that makes the nihilists speak truly, even in a world of composites). And I am saying that the heavy ontologist should be a domain realist, and say that it is better to model the nihilists by stating an interpretation in the metalanguage using the metaphysically distinguished domain (that makes the nihilists speak falsely, given a world of composites).

Thirdly, from the perspective of the hierarchy of metalanguages (§2), *Domain realism* should be understood to apply at every level up the hierarchy. One interprets object language quantification by working in a metalanguage, invoking a model with the distinguished domain, and using quantificational resources. One then interprets metalanguage quantification by working in a meta-metalanguage, invoking a model with *the same distinguished domain* and using quantificational resources, and so on.

Fourth and finally, *Domain realism* represents an exclusively *worldly* and strongly *non-conciliatory* stance. Given *Domain realism*, considerations of usage such as charity are wholly irrelevant to the aptness of the interpreter’s choice of domain. Only the world counts. If the world has composites, those who speak as nihilists speak falsely (no matter how frequently or fervently they do so). Likewise if the world lacks composites, those who speak as universalists speak falsely (no matter how frequently or fervently they do so).

There is an alternative position to consider, on which a theorist interpreting a given object language who specifies a (*Domain, Interpretation*) pair with the distinguished domain gets bonus points, to be weighed against potentially competing factors of usage.¹⁵ Part of the interest of this alternative is that it shows that the heavyweight ontologist can tolerate some limited domain variance, and leads to an alternative stable ways

¹⁵ My thanks to Ezra Rubinstein and Ted Sider for suggesting this alternative.

to combine heavy ontology with light ideology (as hinted at in §1.3). But since I am only claiming to provide one stable way to combine heavy ontology with light ideology, I set alternatives aside.

Clarifications aside, the core idea of *Domain realism* harkens back to Frege's (1879) transformative treatment of quantification, which does not explicitly reference a domain but merely operates against an implicit fixed and total collection of all things. As Dummett (1981: 529) reports: "When these individual variables are those of Frege's symbolic language, then, ... their domain is to be taken simply as the totality of all objects... [E]very object automatically belongs to the domain." Likewise Westerståhl (2011) says:

[T]he only difference between Frege's notion of a second-level concept and the modern notion of a generalized quantifier is that Frege did not have the idea of an *interpretation* or *model*, ... Frege's symbols all had fixed meanings, and the only universe he considered was the totality of everything.

Model theory adds an explicit representation of the domain. I am saying that this technical development is not of much ontological moment vis-a-vis heavy ontology. We may understand the explicit domain in model theory just as Frege understood the implicit domain for his quantifiers, as an objectively fixed totality.

Indeed Sider (2001: xxii) himself speaks as a Frege-style domain realist in the following passage: "[T]he world comes 'ready-made' with a single domain **D** of objects: the class of all the objects there are." So far, so good. But he immediately continues: "This class is the most eligible meaning possible for any symbol playing the inferential role of the unrestricted existential quantifier." I am saying that this continuation is a confusion, reflecting Putnam's and Hirsch's loose talk of "quantifier variance." The domain is not part of the meaning of the quantifier (or any bit of object language ideology), but rather an entity posited in the metalanguage. Here Sider's push for heavy ideology can be seen as stemming from a conflation of object language quantifiers and metalanguage domains.

I have argued that quantifier variance is best understood as domain variance, and that domain variance has nothing to do with the ideology of the object language but rather can be blocked by domain realism. So I conclude that *heavy ontology requires domain realism not heavy ideology*. And I conclude—in favor of *Stable*—that domain realism reconciles heavy ontology with light ideology.

Again, I have not argued for heavy ontology or for light ideology, but merely tried to exhibit a stable—and I think plausible—position that reconciles these views, via domain realism plus ideological pragmatism. To put this point a different way: I say that if there is anything unstable in heavy ontology plus light ideology, the Hirsch-Sider debate has not exposed it.

3. Objections Considered

If I am right, then much of the traditional discussion of quantifier variance and ideology has been mistaken. So it is natural to think "Surely Putnam, Hirsch, Sider, and other such excellent philosophers have not been mistaken!" I acknowledge that it is much more likely that I am the one who is mistaken, I just would like to know how! Fortunately my interlocutors have been generous. I discuss the three objections I have met that I consider the most interesting. (Your mileage may vary.)

3.1 Surrender to the Eklund-Hawthorne semantic argument?

Some may object to my idea (§2.2) that quantifier variance must give way to domain variance, on grounds that it is *uncharitable* to the quantifier variantist, as her view would then be refuted by *the semantic argument* from Eklund (2007: 386–7, 2009: 145) and Hawthorne (2006: 59–60).¹⁶ Sider (2011: 181; cf. Eklund 2007: 386–7, 2009: 145; Hawthorne 2006: 59–60) offers a perspicuous presentation of the argument:

Consider two characters, Big and Small. Big speaks an "expansive" language, Biglish, in which speakers freely quantify over tables. Big introduces a name, '*a*,' for a table, and thus accepts 'Table(*a*).'

¹⁶ I owe the idea behind this objection to Ted Sider (*personal communication*), though I do not know if he would uphold the objection as I phrase it or otherwise. Thanks also to Matti Eklund for helpful discussion here.

Small speaks a “smaller” language, Smallish, in which speakers refuse to quantify over tables. But Small is a quantifier variantist, and thinks that he does not genuinely disagree with Big. So Small says to himself, speaking in Smallish: “Even though there are no tables, the sentence ‘Table(*a*)’ is true in Biglish.” But this commits Small—and all quantifier variantists, who must accept the scenario as described—to rejecting familiar Tarskian ideas about semantics... In order for “‘Table(*a*)’ is true in Biglish’ to be true in Smallish, ‘There is something that “*a*” denotes in Biglish’ must be true in Smallish. But what would this object—in the Smallish sense of “object”—be, if not a table?

This looks like surrender for Small. Indeed, it looks the quantifier variantist has been unmasked as siding with Big, and in general as having a serious and fixed ontology, namely that of the permissive heavy ontologist who permits the objects posited by both sides of the debate.¹⁷

In terms of domain variance, the problem Small has is that there is no consistent interpretation of Smallish that is both nihilistic and variantist. To be nihilistic, Smallish needs to be interpreted via a small domain (§4), but to be variantist and stage a (natural, homophonic) semantics for Biglish, Smallish needs to be interpreted via a big domain.

So is my guiding idea that “quantifier variance” must give way to domain variance *uncharitable* to the quantifier variantist, for dooming her to refutation by the semantic argument? First—just on the interpretive question—I think that it is not uncharitable but fitting for the quantifier variantist to be exposed to the semantic argument. Textually, both Eklund and Hawthorne think that the quantifier variantist is so exposed, and Hirsch & Warren (*forthcoming*) agree—they do not cry “misinterpretation” but rather attempt to face the argument.¹⁸ And indeed I have given an argument that quantifier variance can only be clausal variance or domain variance (§2.1), and that only domain variance is the relevant sort of variance (§2.2). So there is no “more charitable” interpretation.

Does the semantic argument then refute quantifier variantism? I wish! If so then I would have at hand an argument for heavy ontology, and indeed for precisely the sort of permissive heavy ontology I favor (or at least a refutation of a main objection). But alas! I think that the semantic argument fails.

The reason why I think that the semantic argument fails is that Small should not have stated his quantifier variantism in Smallish, but rather—following the lead of orthodox model-theoretic semantics—should have first ascended to a metalanguage—called it *Metalish*—adequate for both Smallish and Biglish (and any other object languages under discussion). This is because quantifier variantism involves semantic claims about these object languages, and in orthodox model-theoretic semantics we are careful to couch semantic claims about object languages in the metalanguage. (This is again the second of my main assumptions flagged in the introduction.) So I say that Small should continue to state his nihilism in Smallish, but should state his variantism in Metalish and not (*pave Sider*) while “speaking in Smallish.”¹⁹

¹⁷ Such a view fits the views of the “maximalist” (Eklund 2007), the “Plenitude Lover” (Hawthorne 2006), and my own “permissivist” (Schaffer 2009), for whom the main questions of existence are to be answered with a dismissive ‘Yes of course!’ (The more interesting questions are the grounding questions.)

¹⁸ See Sider (2011: 182) for another attempt to face the semantic argument. But Hirsch & Warren and Sider share the idea that Small should not interpret Biglish homophonically, or even as speaking a referential language, or having a quantifier. Thus Hirsch & Warren (2017: 13) say: “The semantic shape of an alien language need not fit easily against the shape of our language.” My domain variantist approach allows Small to interpret Biglish homophonically, as a natural human language with standard referential and quantificational resources. This strikes me as a more psychologically plausible interpretive stance for Small to take, with respect to the Biglish speakers.

¹⁹ Likewise Eklund (2009: 145), considering a Carnapian nominalist interpreting ‘2 is prime’ in a platonistic language L_P , says: “[T]he nominalist must concede defeat! For then it can be concluded, in L_N [the nominalist language], that ‘2’ refers and that there are numbers.” I am saying that the nominalist should not be saying this in L_N , but rather in a metalanguage M adequate to interpret both L_N and L_P .

If Small states his nihilism down in Smallish, but states his variantism up in Metalish, he can then consistently hold both. When counting objects, Small can continue to speak Smallish and say nihilistic things like ‘Three things are in the box.’ When doing semantics, Small can ascend to Metalish and interpret Smallish via a small domain and Biglish via a big domain (as per domain variantism), and then say variantist things like ‘“Three things are in the box” is true in Smallish’ and ‘“Seven things are in the box” is true in Biglish.’ (If Small wishes to ascend to *Meta-Metalish*, he will interpret Metalish with a big domain, to handle the semantic values of Smallish and Biglish, etc.) All of this looks stable. So it seems to me that Small can persist in speaking Smallish as his home object language, and can ascend to Metalish (and Meta-Metalish) for his quantifier variantist semantics.

Indeed, this approach strikes me as especially plausible for the quantifier variantist, since it allows for homophonic interpretation. On this approach, Small can see that both Small and Big express exactly the same meaning by ‘There are three things in the box,’ namely the one with the truth-conditions of $|\{\text{Things}\} \cap \{\text{Box-dwellers}\}| = 3$. These truth-conditions will be satisfied in a model with the three element domain, but not in a model with the seven element domain. Domain variance is the key to allowing one and the same meaning to come out true in Smallish and false in Biglish, since this meaning is assessed in different models.

The quantifier variantist will hardly think that any one of these models is metaphysically privileged. (That is the view of her heavy ontology rival.) Rather the domain variantist should say—charting by Hirsch’s idea that the questions of ontology are superficial questions of language—that if one asks “Do composites exist?” then the true answer depends, not on deep features of reality, but just on whether one was asking the question in Smallish (*no*) or Biglish (*yes*), or Metalish (*yes*), etc.

It is easy to feel that Small is cheating if he uses Metalish to interpret Biglish, given that Small is a nihilist, and that in using Metalish to interpret Biglish he uses a domain of composites. How can Small invoke composites?²⁰ I reply that Small’s nihilism always and still consists in the fact that he truly says ‘There are no composites’ in Smallish (a fact which Small can indeed verify from Metalish). This does not change. Small’s relationship to Metalish is not relevantly different, for the quantifier variantist, from Small’s relationship to Biglish. As a quantifier variantist, Small all along thought ‘There are composites’ is true in Biglish, and there should be no barrier to Small speaking Biglish if he wishes. (Small may be a nihilist but he need not be a monoglot.) I am just adding that there is one more language of interest alongside Biglish—Metalish—in which ‘There are composites’ is also true, and which Small can also speak if he wishes. This is nothing new.

So I conclude that, if there is anything unstable in domain variance, the semantic argument has not exposed it. But likewise if there is anything uncharitable in my claim that quantifier variance must give way to domain variance, the objector has not exposed it.

3.2 Loss of Lewisian natural predicates?

Some may object to light ideology before quantifier variance even comes up for discussion, on grounds that Lewis (1983) showed that we needed naturalness for predicates all along, and so needed heavy ideology all along. Indeed Sider (2011: ch. 6) speaks of extending naturalness “beyond the predicate,” and once we admit naturalness for some terms (e.g. predicates), is it so much worse to spread it a bit further? Relatedly, since light ideology doesn’t even grant naturalness for predicates, have I gone too far in the other direction and missed out on all of the uses of Lewisian natural properties? In general one might wonder what implications light ideology has for the Lewisian idea of natural properties.²¹

²⁰ I owe this concern to Eli Hirsch.

²¹ My thanks to Kris McDaniel for helpful discussion on this question.

I say that light ideology is consistent with Lewisian natural properties (if such there be—I need take no stand on that here).²² Indeed, just as it emerged that all of the relevant variability in quantifiers concerns the domain, which is not part of the ideology of the object language but rather a posit in the metalanguage, so I say that all of the relevant variability in predicates concerns the interpretation, which is also not part of the ideology of the object language but equally a posit in the metalanguage. (Domains and interpretations are the two elements of models.) Heavy ideology does nothing for predicates either.

So as a representative example, imagine that Emma is talking about the color of emeralds. The friend of natural properties wants Emma to come out talking about the property *being green* rather than *being grue*. I say that Emma can pick any predicate she likes. She can say that all emeralds are ‘green’ or that all emeralds are ‘grue.’ What matters is not which predicate Emma picks, but rather which extension the interpretation in the metalanguage assigns to Emma’s chosen predicate. (It is only through an interpretation that a predicate gets a semantic value at all.) The issue is whether the interpretation of Emma’s chosen predicate—be it ‘green’ or ‘grue’—maps it to one extension (e.g., the one that includes all the emeralds) or another (e.g., the one that includes only the emeralds unobserved by 2222). In short: Whatever one might think of heavy ideology it has *no bearing* on whether Emma is talking about *being green* or *being grue*.

If one wants Emma to come out talking about the property *being green* rather than *being grue*, what is needed is privileged interpretation in the metalanguage. That is, if one wants Emma’s predicate—be it ‘green’ or ‘grue’—to get the “natural” extension with all the emeralds, what is needed is a constraint on *interpretation* that privileges interpreting Emma’s language via a model whose interpretation maps her predicate to the natural extension rather than the gruesome extension.

So the way to have Lewisian natural properties—and the associated idea of reference magnetism—alongside light ideology is to say that there is pressure on interpreters to prefer interpretations that map predicates to natural extensions, as per:

Interpretation magnetism. There is pressure on interpreters to prefer a (*Domain, Interpretation*) pair that maps predicates to natural extensions over rivals that maps predicates to gruesome extensions.

Interpretation magnetism leaves many details open (as is appropriate), including what makes for a natural extension, and the dynamics of interpretive pressures. As to what makes for a natural extension, perhaps a natural extension corresponds to the instantiation profile of a universal, or a resemblance class of tropes, or just a privileged nominalistic set. (Lewis (1986: 63–4) was himself “undecided” between these options). As to the dynamics of interpretive pressure, it is left open how strong this pressure is, how different “vectors” of such pressure compose, etc. What is relevant is that light ideology permits one to posit natural properties in the world, and to use them as magnets for predicate interpretations, just as Lewis wanted.²³

²² See Dorr & Hawthorne (2013) for a useful overview of the various roles Lewis cast naturalness to play, and their interrelations. For skeptical stances towards the value of Lewisian naturalness see Loewer (2007) and Dasgupta (*forthcoming*). See also Hicks & Schaffer (2017) for the idea that the predicates in scientific laws need only bear a loose connection to metaphysical natural properties.

²³ Indeed *Interpretation magnetism* is not very far from Lewis’s (1983: 372) own first-pass phrasing: Say that a set from the domain is *eligible* to be the extension of a one-place predicate iff its members are just those things in the domain that share some natural property; and likewise for many-place predicates and natural relations. An *eligible interpretation* is one that assigns none but eligible extensions to the predicates.

Here Lewis is supposing all-or-nothing naturalness for properties, and (as he notes) the proposal needs to be loosened to allow for “gradations” (as *Interpretation magnetism* allows). But the relevant point is that natural properties are coming into the story as a constraint on the interpretation of predicates, which is not part of the object language but the interpreting metalanguage.

It is crucial to distinguish between an *uninterpreted predicate* p of a given language L , and its *interpreted counterpart* p of L in model M . An uninterpreted predicate of L is just a bare syntactic string, coupled with the minimal information “I am a predicate” and “I belong to language L .” Uninterpreted predicates can be interpreted in different ways in different models (and famously we can worry about unintended interpretations and various gruesome but theory-preserving global permutations of extensions: Putnam 1981). Reference magnetism is a claim about *uninterpreted predicates*, concerning how the worldly naturalness of certain extensions may exert a magnetic pull on terms whose usage pattern puts them in the neighborhood.²⁴

If we treat ‘green’ and ‘grue’ as uninterpreted predicates (as I have been doing above, since part of the work of naturalness is to guide predicates to the right extensions in the first place) then it should be clear that the world does not care which bare syntactic string Emma uses.

But we might also treat ‘green’ and ‘grue’ as interpreted predicates (and/or fix the language as English, under a fixed pre-conception of the proper interpretation for English). Then—on the Lewisian view—it may well matter which predicate Emma uses. In particular, there may well be worldly pressure on Emma to prefer ‘green’ over ‘grue’ *given that ‘green’ has the extension with all the emeralds, and ‘grue’ has the extension that includes only the emeralds unobserved by 2222*. But this is only because the interpretation has already been fixed, and the magnetic work of naturalness bypassed. The world may well care which interpreted terms Emma uses, but that is fully consistent with light ideology plus *Interpretation magnetism*. On this view, all the worldly pressure is coming via the ontology, through the idea that certain extensions (sets posited in the metalanguage) are specially privileged. By fixing an interpretation of a predicate as having such a privileged extension, the predicate comes to inherit this privilege. But the real action was getting the predicate onto the privileged extension in the first place, and in this interaction it is the status of the extension and not of the predicate that is doing all the pulling.

Indeed *Domain realism* and *Interpretation magnetism* can be thought of as independent but thematically connected ways for the world to “stabilize” domains and predicates respectively. But the important point is that all of this stabilizing action is happening (/failing to happen) via objects posited in the metalanguage, and nothing in either *Domain realism* or *Interpretation magnetism* cares about the ideology of the object language.

3.3 But aren’t quantifiers just second-order properties?

I have made two claims that a careful reader might think are in tension. In discussing Lewisian naturalness (§3.2), I have allowed that there may well be natural properties and relations. But in sketching the orthodox semantics for generalized quantifiers (§2.1), I have said that quantifiers express comparisons via properties, as encoded in the schema:

Template-Generalized: $[[Q(R)(S)]]_M = \text{True}$ iff Q -many members of R are members of S

This tells us that quantified expressions express *second-order relations*. For instance, ‘some’ expresses the second-order relation of *having a non-empty intersection* as between the first-order properties it compares. So does that not provide a path for natural quantifier meanings, as natural second-order relations?²⁵

²⁴ Lewis (1983: 370–1) is very clear on this, describing Putnam’s argument as concerning “constraints on interpretation” and as being driven by the concern that “intended interpretations [that satisfy a given theory] are surprisingly abundant,” and describing his reply to Putnam as: “[W]e need natural properties to explain determinacy of interpretation.”

²⁵ I am grateful to both Kris McDaniel and Jonathan Shaheen for suggesting this style of objection, though I do not know if either would uphold the objection as I phrase it or otherwise.

I do allow that there are natural properties and relations, and do accept that quantified expressions express second-order relations.²⁶ But I reply that the *way* that quantified expressions express relations works differently from the way that predicates express relations, such that metaphysical privileges matter only for predicates. First some background: Model-theoretic semantics takes for granted a given object language under study with a fixed vocabulary, whose terms comes divided into kinds, including individual terms (variables and constants), predicates, and logical terms (such as connectives and operators). The semantics treatments of these kinds of terms are wholly distinct. Individual terms are assigned to elements of the domain (constants are assigned directly, while variables are assigned through the assignment function). Predicates are assigned extensions from the domain (via the interpretation). And logical terms are simply stipulated to have thin and fixed logical meanings from the start, independent of the domain or the interpretation.

Indeed it is orthodox—following on Tarski’s conception of a logical truth—to define a logical truth as a sentence true in all models (whether “metaphysically distinguished” or not). This presupposes that we have the right range of models, namely those that vary the semantic values of constants and predicates, but do not vary the semantic values of logical terms. To illustrate, so that (i) neither ‘ Fa ’ nor ‘ $\neg Fa$ ’ counts as a logical truth, we need to look at models with different semantic values for ‘ F ’ and/or ‘ a ’; and so that (ii) ‘ $Fa \vee \neg Fa$ ’ counts as a logical truth, we must fix that ‘ \vee ’ has the thin logical meaning of disjunction, and that ‘ \neg ’ has the thin logical meaning of negation. So the orthodox definition of logical truth presupposes a sharp distinction between the variability of predicates and the fixity of logical terms across models.

Thus model-theoretic semantics treats quantificational expressions not as relational terms but as logical terms. In that sense, it actually draws a sharp distinction between ‘some’ and ‘having a non-empty intersection.’²⁷ The former is born logical and the latter is born relational. The former is stipulated to have a thin and fixed logical meaning, the latter is left to the interpretation function to handle, in variable and potentially gruesome ways.

This sharp distinction matters to naturalness. Naturalness was coming in as a constraint on interpretations, to block gruesome interpretations. That is perfectly appropriate *for the interpretation of predicates*. But it has no bearing on the interpretation of quantificational expressions, because quantificational expressions are simply given a thin and fixed logical meaning at birth, with no room for re-interpretation, gruesome or otherwise. So there is no path to more and less natural quantifier meanings (which *Interpretation magnetism* might then operate over), because there is no prospect of difference in quantifier meanings at all.

Indeed, since the fact that ‘some’ expresses the second-order relation of *having a non-empty intersection* between first-order properties is set by the semantic clause, the idea of variance in the second-order relations is just clausal variance by another name (which I argued was not relevant in §2.2).

The most we can do is to state equivalencies between quantificational expressions and predicates:

$$[[\textit{some} (R)(S)]]_M = \text{True} \quad \text{iff} \quad [[\textit{having a non-empty intersection} (R)(S)]]_M = \text{True}$$

But we must be careful about the status of this equivalence. It does *not* mean that quantificational expressions are open to re-interpretation, and it is not itself a logical truth (it does not hold in all models). Indeed, continuing the point about logical truth, we want ‘if Fa then $[\exists x: x=x] Fx$ ’ to be a logical truth, but ‘if Fa

²⁶ Sider (2011: 90) says that the linguistic treatment of quantifiers as second-order relations is metaphysically wrong: “For the treatment of quantifiers as expressing second-order properties or relations, however appropriate in linguistic theory, does not ring true at the metaphysical level.” If he is right then the objection under discussion falls. But I do not take this route because I do not share his intuitions, and in any case prefer to be guided by the orthodox semantics over any metaphysical intuitions in this case.

²⁷ Here and in what follows I adopt the simplifying expedient of ignoring the internal structure of ‘having a non-empty intersection’ and treating it as a simple, unstructured predicate.

then F and F have a non-empty intersection' to be non-logical, since 'having a non-empty intersection' might get interpreted in any of many ways, including *having an empty intersection*.

Rather I suggest that it is best to think of the equivalence just stated as holding in virtue of meta-semantic constraints on the intended interpretation of the predicate 'having a non-empty intersection.' Just as it is a meta-semantic constraint (perhaps due to usage patterns) that guides us in mapping 'remembers' and 'knows' onto extensions such that the extension of the former is a subset of the extension of the other, so we may be guided to mapping 'having a non-empty intersection' onto an extension that relates property pairs just in case the associated quantified claim hold. In short, interpreters should prefer models with an interpretation function that validates this equivalence.

Perhaps model theory's sharp distinction between logical terms like 'some' and relational terms like 'having a non-empty intersection' is artificial. If so then orthodox model theory may not be the best framework for our meta-metaphysical debate. The objector is invited to exhibit a better alternative. But I must conclude that heavy ontology plus light ideology is stable, given orthodox semantic assumptions.²⁸

References

- Barwise, John and Robin Cooper 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4: 159–219.
- Bennett, Karen 2009. Composition, Colocation, and Metaontology. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 38–76. Oxford: Oxford University Press.
- Carnap, Rudolf 1950. Empiricism, Semantics, and Ontology. *Revue Internationale de Philosophie* 4: 20–40.
- Chalmers, David 2009. Ontological Anti-Realism. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 77–129. Oxford: Oxford University Press.
- Chalmers, David, David Manley, and Ryan Wasserman. Introduction: A Guided Tour of Metametaphysics. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 1–37. Oxford: Oxford University Press.
- Dasgupta, Shamik *forthcoming*. Realism and the Absence of Value. *The Philosophical Review*.
- Dorr, Cian and John Hawthorne 2013. Naturalness. *Oxford Studies in Metaphysics* 8: 3–78.
- Dummett, Michael 1981. *Frege: Philosophy of Language (Second Edition)*. Cambridge, MA: Harvard University Press.
- 1991. *Frege: Philosophy of Mathematics*. Cambridge, MA: Harvard University Press.
- Eklund, Matti 2007. The Picture of Reality as an Amorphous Lump. *Contemporary Debates in Metaphysics*, eds. Theodore Sider, John Hawthorne, and Dean Zimmerman: 382–96. Oxford: Basil Blackwell.
- 2008. Putnam on Ontology. *Following Putnam's Trail: On Realism and Other Issues*, eds. María Uxía Rivas Monroy, Celeste Cancela Silva, and Concha Martínez Vidal: 203–22. Amsterdam: Rodopi.
- 2009. Carnap and Ontological Pluralism. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 130–56. Oxford: Oxford University Press.
- Finn, Suki and Otavio Bueno *manuscript*. Quantifier Variance Dissolved.
- Frege, Gottlob 1879. *Begriffsschrift*. Halle: Louis Nebert.
- Hawthorne, John 2006. Plenitude, Convention, and Ontology. *Metaphysical Essays*: 53–70. Oxford: Oxford University Press.
- 2009. Superficialism in Ontology. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 213–30. Oxford: Oxford University Press.
- Hicks, Michael Townsen and Jonathan Schaffer 2017. Derivative Properties in Fundamental Laws. *The British Journal for the Philosophy of Science* 68: 411–50.
- Hirsch, Eli 1993. *Dividing Reality*. Oxford: Oxford University Press.
- 2002. Quantifier Variance and Realism. *Philosophical Issues* 12: 51–73.

²⁸ Thanks to Chiara Brozzo, Janelle Derstine, Matti Eklund, Suki Finn, Eli Hirsch, James Miller, Thomas Hofweber, Kris McDaniel, Ezra Rubinstein, Jonathan Shaheen, Ted Sider, Tatjana von Solodkoff, the members of the Rutgers Metaphysics Reading Group, and the participants at *The Language of Ontology* (Trinity College), and *The Uppsala Metaontology Workshop* (Uppsala University).

- 2005. Physical-Object Ontology, Verbal Disputes, and Common Sense. *Philosophy and Phenomenological Research* 70: 67–97.
- 2008. Language, Ontology, and Structure. *Noûs* 42: 509–28.
- 2009. Ontology and Alternative Languages. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 231–59. Oxford: Oxford University Press.
- 2011. *Quantifier Variance and Realism*. Oxford: Oxford University Press.
- Hirsch, Eli and Jared Warren *forthcoming*. Quantifier Variance and the Demand for a Semantics. *Philosophy and Phenomenological Research*.
- Koslicki, Kathrin 2003. The Crooked Path from Vagueness to Four-Dimensionalism. *Philosophical Studies* 114: 107–34.
- Lewis, David 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61: 343–77.
- 1984. Putnam's Paradox. *Australasian Journal of Philosophy* 62: 221–36.
- 1986. *On the Plurality of Worlds*. Oxford: Basil Blackwell.
- Lindström, Per 1966. First-Order Predicate Logic with Generalized Quantifiers. *Theoria* 32: 186–95.
- Loewer, Barry 2007. Laws and Natural Properties. *Philosophical Topics* 35: 313–28.
- McDaniel, Kris 2009. Ways of Being. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 290–319. Oxford: Oxford University Press.
- 2017. *The Fragmentation of Being*. Oxford: Oxford University Press.
- McSweeney, Michaela Markham *forthcoming*. Following Logical Realism Where it Leads. *Philosophical Studies*.
- Mostowski, Andrzej 1957. On a Generalization of Quantifiers. *Fundamenta Mathematicae* 44: 12–36.
- Putnam, Hilary 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- 1987. Truth and Convention: On Davidson's Refutation of Conceptual Relativism. *Dialectica* 41: 69–77.
- 1994. The Question of Realism. *Words and Life*: 295–312. Cambridge, MA: Harvard University Press.
- 2004. *Ethics without Ontology*. Cambridge, MA: Harvard University Press.
- Quine, W. V. O. 1948. On What There Is. *Review of Metaphysics* 2: 21–38.
- 1951. Ontology and Ideology. *Philosophical Studies* 2: 11–15.
- 1983. Ontology and Ideology Revisited. *The Journal of Philosophy* 80: 499–502.
- Rayo, Agustin and Gabriel Uzquiano (eds.) 2006. *Absolute Generality*. Oxford: Oxford University Press.
- Schaffer, Jonathan 2009. On What Grounds What. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 347–83. Oxford: Oxford University Press.
- 2017. The Ground Between the Gaps. *Philosophers' Imprint* 17: 1–26.
- Sider, Theodore 2001. *Four Dimensionalism*. Oxford: Oxford University Press.
- 2007. Neo-Fregeanism and Quantifier Variance. *Proceedings of the Aristotelian Society* Sup. 81: 201–32.
- 2009. Ontological Realism. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 384–423. Oxford: Oxford University Press.
- 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- 2014. Hirsch's Attack on Ontologese. *Noûs* 48: 565–72.
- Tarski, Alfred 1983 [1933]. The Concept of Truth in Formalized Languages. *Logic, Semantics, Metamathematics; Papers from 1923 to 1938 (Second Edition)*, trans. J. H. Woodger: 152–278. Indianapolis, IN: Hackett Publishing.
- Thomasson, Amie 2007. *Ordinary Objects*. Oxford: Oxford University Press.
- 2009. Answerable and Unanswerable Questions. *Metametaphysics*, eds. David Chalmers, David Manley, and Ryan Wasserman: 444–71. Oxford: Oxford University Press.
- Turner, Jason 2010. Ontological Pluralism. *The Journal of Philosophy* 107: 5–34.
- van Inwagen, Peter 1990. *Material Beings*. Ithaca, NY: Cornell University Press.
- 2014. Introduction: Inside and Outside the Ontology Room. *Existence: Essays in Ontology*: 1–14. Cambridge: Cambridge University Press.
- Warren, Jared 2017. Quantifier Variance and Indefinite Extensibility. *The Philosophical Review* 126: 81–122.
- Westerståhl, Dag 2011. Generalized Quantifiers. *Stanford Encyclopedia of Philosophy*: <<<https://plato.stanford.edu/entries/generalized-quantifiers/>>>, last accessed 11/20/17.
- Williams, R. J. G. 2007. Eligibility and Inscrutability. *The Philosophical Review* 116: 361–99.