Higher-order metametaphysics*

Theodore Sider July 20, 2022

The recent turn to higher-order languages—languages with quantification into predicate, sentence, and other nonnominal positions—promises elegant and more accurate modes of expression, new solutions to old problems, transformation of problem spaces, and generation of new questions: a paradigm shift. The excitement in the peroration of Cian Dorr's agenda-setting paper "To be *F* is to be *G*" is typical of the spirit of this movement, and is undeniably infectious:

> And the exploration has barely begun: there is a whole continent of views waiting to be mapped out, and at this point we can only guess which of them will look most believable in the long run. Onwards!

As a community, the best way to handle such new, all-encompassing, and programmatic ideas is to run with them. Many of us should embrace the new framework, explore it from the inside, and see where that leads. Setting aside the inevitable nay-sayers, that's what we did with logic and linguistic analysis in the early twentieth century, and with possible worlds and modal logic in the 1970s (to take just two examples), in each case with great success. In-depth exploration is needed to tell whether ideas are on the right track; we know them by their fruit. Onwards indeed!

But nay-saying has its place too. Strawson and the other ordinary-language philosophers provided a corrective to Russell and his heirs, important parts of which were eventually assimilated into the mainstream. Quine (at the very least) forced modal enthusiasts to clearly articulate and embrace their metaphysical commitments.[1]

And sometimes nay-sayers are right. Despite its appeal and promise, there are important metaphysical questions about the foundations of the higher-order approach. Are higher-order languages in metaphysically good standing? That is, do such languages succeed in latching onto reality; is reality such as to be well-represented by them? If so, then it would indeed make sense to stay up

*A more accurate title would be "Meta-(higher-order metaphysics)", but.... Thanks to Daniel Berntson, Dan Marshall, Jeff Russell, and Timothy Williamson.

[1]Hirsch (2011), Thomasson (2007, 2015), and others have played a similar role in another context, forcing the dominant "Quinean" tradition in ontology to articulate and defend its foundational assumptions.

nights wondering whether, for instance, $p = (p \,\&\, p)$, for all p. Such questions would concern reality's higher-order structure. If not, the questions might not have answers at all (if higher-order sentences fail to be truth-apt), or might fail to have determinate or objective answers, or might have "unwanted" answers (if, say, higher-order sentences have first-order, set-theoretic truth conditions), or might suffer some other sort of "discourse failure".[2]

1. Higher-order languages[3]

In the late nineteenth century, the concept of set—the concept of a collection conceived as an individual thing—became central to the foundations of mathematics. But around the turn of the century, apparent contradictions in this idea were discovered, the simplest and most famous of which was Russell's Paradox. Define r as the set of all and only those sets that are not members of themselves. Thus a set is to be a member of r if and only if it is not a member of itself:

$$\forall x(x \in r \leftrightarrow \sim x \in x) \tag{R}$$

As Russell observed, if we substitute 'r' for the universally quantified variable 'x', we obtain a contradiction:

$$r \in r \leftrightarrow \sim r \in r$$

Two parts of this reasoning bear emphasis. First, we are treating expressions for sets as being grammatically like expressions for their members. Thus in addition to formulas like $x \in r$, formulas like $x \in x$ and $r \in r$ (and thus their negations $\sim x \in x$ and $\sim r \in r$) are also grammatical. In modern terms, we are speaking of sets using a first-order language; we refer both to sets and their

[2] One might stay up nights even given discourse failure. Addressing a question can have value beyond the value of answering it. In any inquiry into metaphysics (or anything in philosophy, for that matter), there is also the value of cartography of logical space, including the creative/expansive invention of new positions, and the dialectical/contractive search for reasons against positions. Because higher-order metaphysics is so formal, cartography in this domain has a certain kind of known value, the same value that attaches to mathematical investigations independent of their truth. But cartography has value even in less formal metaphysics, where we rely less on the kinds of clear-cut considerations logic studies, and more on the woollier, less well-understood (but no less essential) considerations we call philosophical. (The latter are still needed in higher-order metaphysics, for instance in informing the ubiquitous judgments about which hypotheses are worth exploring and which are not.)

[3] For further background on these issues, see Bacon (2022); Sider (2020*a*).

members using singular terms; and we ascribe membership using a two-place predicate $\in$.

Second, the existence of the set r is simply assumed. Without that assumption there is no paradox, just as the "paradox" of the barber who shaves all and only those who don't shave themselves is easily dissolved by denying the existence of the barber. The assumption that r exists is based on the assumption that every formula determines a set, or in modern terms, that every instance of the Naive Comprehension schema is true:

$$\exists y \forall x (x \in y \leftrightarrow \phi) \qquad \text{(Naive Comprehension)}$$

In this schema, ϕ may be replaced with any formula with no free occurrences of variables other than x.[4] Replacing ϕ with the formula $\sim x \in x$ yields the instance $\exists y \forall x \big(x \in y \leftrightarrow \sim x \in x\big)$; existentially instantiating to an arbitrary name r yields the contradictory sentence (R).

The dominant approach to the paradox has been to reject the second assumption (and thus Naïve Comprehension). The set r doesn't exist. We cannot simply assume that every formula corresponds to a set. Rather, we must carefully develop a theory of when sets exist and when they do not, a theory which implies the existence of all the sets we need in mathematics but does not imply contradictions like (R). Zermelo Frankel (ZF) set theory is an elegant theory of this sort, and is the dominant theory of sets today.

But there is another possible approach: reject the first assumption, according to which a term for a set and a term for one of its members have the same grammar. Russell and Whitehead adopted such an approach, known as the ramified theory of types, in *Principia Mathematica*. That theory was ungainly and soon became obsolete, but an improved version due to Church (1940) and others, known as the simple theory of types, continued to be studied by logicians (and its descendants by computer scientists); and it is this and related theories that have become so popular in recent metaphysics, philosophy of language, and philosophical logic.

Actually this type-theoretic approach has been used to develop a consistent theory of properties, relations, propositions, and the like, rather than sets. These entities are capable of playing a similar role to sets in the foundations of mathematics. And inconsistency threatens these entities just as it threatens sets:

[4] This is a special case. The more general version of the schema allows "parameters": instances may be prefixed with any number of universal quantifiers binding variables which may occur freely in ϕ.

uncritically assuming the existence of a property for each predicate yields the property of being a property that doesn't instantiate itself, which would then instantiate itself if and only if it does not instantiate itself. The type-theoretic resolution of this paradox is roughly that an expression for a property is not grammatically like an expression for one of its instances, so that a statement saying that a property instantiates itself will be ill-formed, and the paradox does not get off the ground.

In more detail: the simplest resolution of the paradox takes talk of properties to be formalized using the language of second-order logic, in which there are variables whose grammar is that of predicates, in addition to first-order logic's variables whose grammar is that of names. Thus the grammar of second-order (but not first-order) predicate logic allows formulas such as these:

$$\exists F F(a)$$
$$\forall R\big(R(a,b) \to R(b,a)\big)$$
$$\forall x \exists F\, F(x)$$

"Quantification over properties" is thus understood as quantification into predicate position: $\exists F$ and $\forall F$. And the "attribution of a property" to a thing is achieved, not by attaching a dyadic predicate of instantiation to a singular term naming the thing and a singular term naming the property, as in 'x instantiates y' (compare '$x \in y$') but rather by attaching a predicate to a singular term for the thing: '$F(x)$'. The attempt to formulate a claim that a property instantiates itself then becomes '$F(F)$', which is as ungrammatical in second-order logic as it is in first.

The language of second-order logic is a special case of the higher-order languages that are now popular, which in general go beyond second-order languages in two ways: they allow constants and variables of arbitrary grammatical categories, and they allow lambda abstraction.

The notion of "arbitrary grammatical categories" is made precise by the device of *types*. Types are conventional entities used to represent, or code up, grammatical categories; thus we speak of expressions in formal languages as *having* or being *of* types. The purpose of representing grammatical categories as entities is to allow us to quantify over them in the metalanguage, in order to make generalizations: "for any type, τ, if an expression has type τ, then …".

Here is one typical development of the idea. We begin with a type, e, which will represent the grammatical category of expressions that stand for *e*ntities (i.e., singular terms). What entity, exactly, is the type e? It doesn't matter; the association between types, construed as entities, and the grammatical categories they represent is purely conventional. We might as well take the type e to be the *letter* 'e'.

e is called a "primitive" type because it isn't constructed from any other types. All other types are constructed from simpler types, according to this rule, where n may be any natural number (including 0):

$$\text{If } \tau_1, \ldots, \tau_n \text{ are types, then } (\tau_1, \ldots, \tau_n) \text{ is also a type} \tag{T}$$

(Again, keep in mind the conventionality of types. $(\tau_1, \ldots, \tau_n)$ can be regarded as nothing more than the left paren, followed by τ_1, followed by a comma, followed by τ_2, then a comma, ..., followed by τ_n, followed by the right paren.) The type $(\tau_1, \ldots, \tau_n)$ represents the grammatical category of an expression that combines with n expressions, of types $\tau_1, \ldots, \tau_n$, respectively, to make a *formula* (that is, to make an expression that can be either true or false, if none of its variables are free). That is, an expression of type $(\tau_1, \ldots, \tau_n)$ is an *n-place predicate* whose arguments are of types $\tau_1, \ldots, \tau_n$.

An important special case is when $n = 0$; the resulting type () is the type of formulas. (An expression that doesn't need any arguments in order to make a formula is already a formula.) Further examples: (i) (e) is the type of an expression that combines with an expression of type e (a singular term) to make a formula. That is, (e) is the type of the familiar one-place predicates of first-order logic. (ii) $\big((),()\big)$ is the type of an expression that combines with two expressions of type () (i.e., with two formulas) to make a formula. That is, it's the type of two-place sentence operators, such as & or ∨.

The rule (T) can be applied iteratively, since $\tau_1, \ldots, \tau_n$ may be any types, including complex ones. Since (e) is a type, so is $\big((e)\big)$; but then $\Big(\big((e)\big)\Big)$ is also a type; and so on. There are infinitely many types.

In a typical higher-order language based on this simple type theory, constants and variables of each of the infinitely many types are allowed. Thus in addition to quantifying into singular-term-position (as in first-order logic), or predicate position (as in second-order logic), one can quantify into sentence position (variable of type ()), as in:

$$\forall p \exists q (q \leftrightarrow \sim p)$$

("for every proposition, there exists a proposition that is true iff the first is not true"), or into one-place sentence-operator position:

$$\exists O \forall p \big(O(p) \leftrightarrow \sim p\big)$$

("There exists a property of propositions that is had by a given proposition iff the proposition is not true"), or any other position represented by a type.[5]

In addition to quantification into positions of all types, the currently popular higher-order languages include a second innovation (also due to Church): lambda abstraction. The purpose of lambda abstraction is to allow for complex expressions of arbitrary type. For example, in addition to a simple predicate F for 'frolics' and a simple predicate G for 'gallops', we might want a complex predicate for their "conjunction", a predicate that means 'frolics-and-gallops'. (Why? For one thing, to be an allowable substitution for predicate variables in quantified sentences. 'For all properties, if Thunder has the property then Laredo has the property' should imply something like 'If Thunder frolics-and-gallops then Laredo frolics-and-gallops'.) We represent 'frolics-and-gallops' as $\lambda x.(Fx \mathbin{\&} Gx)$, which is a predicate, read as "is an x such that x is F and x is G". In general, where $v_1, \ldots, v_n$ are any variables, of types $\tau_1, \ldots, \tau_n$, respectively, and ϕ is any formula, then $\lambda v_1, \ldots, v_n.\phi$ is an expression of type $(\tau_1, \ldots, \tau_n)$, meaning "are $v_1, \ldots, v_n$ such that ϕ".

A nice perk of lambda abstraction is that it can take over the job of variable binding from quantifiers. For instance, the standard first-order quantifiers $\forall$ and $\exists$ can be treated as predicates of one-place first-order predicates—that is, expressions of type $\big((e)\big)$. Thus instead of:

$$\forall x F(x) \qquad \exists x\big(F(x) \mathbin{\&} G(x)\big) \qquad \forall x \exists y R(x,y)$$

we could write, in official contexts anyway:

$$\forall(F) \qquad \exists\Big(\lambda x.\big(F(x) \mathbin{\&} G(x)\big)\Big) \qquad \forall\Big(\lambda x.\exists\big(\lambda y.R(x,y)\big)\Big)$$

In general, a quantifier over "τ-entities" is a predicate of predicates of τ-entities, and thus has type $\big((\tau)\big)$.

2. "Innocent" higher-order quantification

Quine famously said that second-order logic is set theory in sheep's clothing, meaning that a second-order sentence like $\exists F F(x)$ really just means that x is a member of some set. (Or that x instantiates some property; but Quine

[5] Here I am using p and q as variables of type (), and O as a variable of type $\big((),()\big)$. Often the types of expressions are represented explicitly by superscripting: $p^{()}$, $q^{()}$, $O^{((),())}$.

took talk of properties to be strictly less clear than talk of sets.) That is, a more perspicuous way of saying what you were trying to say would use a first-order language: $\exists y\, x \in y$ (or $\exists y\, I(x, y)$, where I means "instantiates"). More perspicuous because it makes clear our ontological commitments: by saying it in the first-order way, we no longer hide our commitment to an ontology of sets (or properties).

The central presupposition of higher-order metaphysics is that Quine was mistaken about this. To be sure, one could choose to use a higher-order language to express the same claims that might otherwise be expressed in a first-order language quantifying over sets or properties. But it is also possible to take the higher-order languages as sui generis, expressing claims that cannot be expressed in first order languages. Thus understood, a sentence like $\exists F\, F(x)$ can be true even if there are no such things (first-order quantifier) as sets or properties; it does not mean that there exists some entity (first-order quantifier) that x instantiates or is a member of. So what does it mean, then? It means, well, that $\exists F\, F(x)$. Similarly for sentences containing variables of other types. $\exists p\, p$ does not mean that there exists (first-order quantifier) some proposition that is true; it can be true even if there are no such things as propositions. It means, well, that $\exists p\, p$.

If this view is correct, then many of the intuitive glosses of higher-order claims that I have been giving (and will continue to give) are misleading. Strictly speaking we should not gloss $\exists F\, F(x)$ as "x has some property", or $\forall p\big(p = (p \,\&\, p)\big)$ as "every proposition is identical to its self-conjunction", since each gloss suggests first-order quantification, over properties in the first case and propositions in the second. Indeed, it's debatable whether higher-order claims can be perspicuously stated in natural language at all.

This anti-Quinean view is sometimes put by saying that higher-order quantifiers are "ontologically innocent", that they are not "ontologically committing"; but this can seem to mean more than it does. What it does mean is that higher-order quantification does not commit us to there being anything *in a first-order sense*. $\exists F\, F(x)$ can be true without there being some entity (first-order quantifier) corresponding to the predicate variable F. Nevertheless, there is a perfectly good sense in which it is "ontologically committing". $\exists F\, F(x)$ is, after all, an existential sentence, and says that *there is* an F of a certain sort; it is false if there is no F that x has, in the second-order sense of 'there is no'.

George Boolos (1984) famously defended an anti-Quinean view in this vicinity, according to which *plural* quantifiers, such as 'some' in 'some pallbearers lifted the casket', are sui generis, and are not first-order quantifiers over sets or the like. Our discussion will encompass Boolos's view, but the current higher-orderists depart from Boolos in two main ways.[6] First, Boolos gives us a way of interpreting the language of monadic, second-order logic (in which the only non-first-order quantification is into the position of one-place first-level predicates, i.e., type (e)), whereas the current higher-order movement embraces quantified variables of arbitrary type, and thus of arbitrary 'adicy and level. Second, plural variables are "extensional": given the plural interpretation, the following sentence is true:

$$\forall F \forall G\big(\forall x(Fx \leftrightarrow Gx) \rightarrow F = G\big)$$

If it happens to be that all and only creatures with hearts are creatures with kidneys, then the creatures with hearts (plural variable) *are* the creatures with kidneys. The higher-orderists, on the other hand, do not accept the sentence. Thus their "$\forall F$" is more akin to "all properties" than "all pluralities" (setting aside the misleading suggestion of first-order quantification over properties).

In the claim just displayed, an identity predicate was flanked by second-order variables. Such higher-order identity predicates play a central role in many higher-order inquiries. Suppose that, for any type, τ, we introduce an identity predicate for that type, $=^\tau$. (It could be taken as primitive; or $\alpha =^\tau \beta$ could be defined as meaning that α and β have the same properties: $\forall F\big(F(\alpha) \leftrightarrow F(\beta)\big)$, where F is a variable of type (τ).) We may then raise the question of "fineness of grain" for type τ: under what conditions are "entities of type τ" the same or different? For instance, are propositions (forgive the first-order sound) individuated by truth value (the coarsest imaginable grain)? The thesis that they are may be stated as follows:[7]

$$\forall p \forall q\big((p \leftrightarrow q) \rightarrow (p =^{()} q)\big) \qquad \text{(extensional propositional grain)}$$

Or are they instead individuated by necessary equivalence? Helping ourselves to an operator $\Box$ for necessity, that claim would be:

$$\forall p \forall q\big(\Box(p \leftrightarrow q) \rightarrow (p =^{()} q)\big) \qquad \text{(intensional propositional grain)}$$

[6] A third way is that they allow properties with no instances, whereas Boolos does not recognize an empty plurality.

[7] '$\leftrightarrow$' is the material biconditional; $\phi \leftrightarrow \psi$ is true if and only if ϕ and ψ have the same truth value.

Or perhaps they are even finer-grained? Since $\Box\big(p \leftrightarrow (p \,\&\, p)\big)$, intensional propositional grain implies:

$$\forall p\big(p =^{()} (p \,\&\, p)\big)$$

which is incompatible with a competing idea about propositional grain: that propositions are "structured" in some sense.

Similar questions of grain may be raised for any type, with the help of lambda abstraction. Are properties identical to their self-conjunctions: $\forall F\big(F =^{(e)} \lambda x.(Fx \,\&\, Fx)\big)$? Is negation the same as triple negation: $\sim\, = \lambda p.({\sim}{\sim}{\sim} p)$? And so on.

All such claims are ungrammatical in first-order logic. Adopting the higher-order language opens up Dorr's continent of possible views about grain.

So: is higher-order logic "innocent"? Better: is higher-order quantification in good standing?

It isn't fully clear what Quine himself meant by calling second-order logic set theory in sheep's clothing, or what his reasons were.[8] Sometimes he simply assumes, begging the question, that all quantified variables range over entities (Quine, 1970, pp. 66–7). Sometimes he is insisting that second-order logic is no more part of logic proper than first-order set theory; but the current higher-orderists don't seem to view logicality as an especially important classification.

In my view, the prima facie case against accepting irreducibly higher order quantification is simply parsimony. Posits that make the world more complex are, other things being equal, to be avoided. And the posit of higher-order quantification makes the world much, much more complex. Dorr's continent, exciting as it admittedly is, is exactly the problem. The posit commits one to an ocean of new facts, and the size of the continent suggests the size of the ocean. When I embraced the logical apparatus of first-order logic, using an embeddable negation sign, I didn't sign up for questions such as whether $\sim\, =^{((),())} \lambda p.({\sim}{\sim}{\sim} p)$. The jump in expressive power when we adopt the higher order framework does indeed lead to exciting opportunities for new research, but the downside is increase in the complexity of the world. This is a basic and common sort of recoil from a proposed metaphysical commitment.[9]

Higher-order-ism might be worth its cost in complexity, just as properties in physics such as charge and mass are presumably worth their cost. But the

[8] See Boolos (1975) and Turner (2015).

[9] An even more common sort is less defensible (in my view): recoil from apparently unknowable facts. The epistemic recoil has a quite different source, neo-verificationist rather than Occamist. See Sider (2020*b*, section 3.15) for a discussion of related issues.

cost is nevertheless real, not to be paid lightly. The remainder of this paper will examine arguments that the cost is indeed worth paying.[10] But it may be objected right at the start that there is no cost at all, precisely because higher-order quantification is innocent. I have already indicated in a preliminary way my objection to this thought: although higher-order quantifications are not ontologically committal in a first-order sense, they are ontologically committal in a higher-order sense. But more can be said.

First, there are many other cases in which a "metaphysical commitment" doesn't involve a first-order ontological commitment. For instance, the adoption of a modal outlook—using modal operators in our theorizing—presumably has no distinctive ontological commitments: the modal operators do not correspond to new entities, but rather to "new modes of truth", so to speak. For "modalists", reality has a modal aspect, an aspect unrecognized by anti-modalists like Quine. Modalists accept an ocean of facts, resulting in a continent of new questions, such as whether reality might have been exactly as it actually is physically but not mentally, whether I could have been born from different parents, and so on. The world is a more complex place according to modalists than it is according to Quine, despite the fact that modalists don't (or needn't) recognize any new entities. (My quantification over aspects, modes, and facts in this paragraph is inessential, present only because of natural language's preference for nouns—illustrated at least twice by this very sentence.) Similarly, the adoption of predicates of 'charge' and 'mass' in physics, which everyone acknowledges as involving an increase in worldly complexity, don't (or needn't) involve postulation of new entities.

In each case, one might argue that new entities should indeed be postulated: universals of charge and mass, properties of necessity and possibility possessed by propositions (or perhaps instead: possible worlds). Whether there are such

[10]Could the continent itself justify the cost? The ideological and ontological posits of medieval angelology are hardly justified by its continent of questions. But higher-order metaphysics is more formally disciplined than inquiry into angels dancing on pin-heads. We seem to be getting some formal traction: the investigation can be carried out with mathematical rigor, certain initially natural-seeming ideas can be demonstrated to be inconsistent, formally natural groupings of viewpoints have begun to emerge, and so on. It's an interesting question whether this is any sort of evidence that we are making contact with reality. It's important not to confuse the admitted value of mathematically disciplined cartography of logical space (note 2) with evidence of contact with reality. Still, some may claim that in pure mathematics itself, formal traction is evidence of contact with reality. The most extreme version of this view is one commonly associated with Hilbert (e.g., 1899), namely that consistency implies truth in mathematics; a variant might say that "consistent ideology" is ipso facto in good standing.

entities is a matter of controversy: "platonists" say there are, "nominalists" (like me) say there aren't. But it shouldn't be controversial that if nominalism is true (and if possible worlds don't exist), the adoption of modal operators or predicates of charge and mass still amounts to the recognition of added worldly complexity, of a sort that ought, other things being equal, to be minimized.

I myself think of the metaphysical commitment in a certain way, as in part including a commitment to the expressions in question "carving at the joints" (Lewis, 1983; Sider, 2011).[11] But the point is not tied to this metaphysical baggage. Even those who are skeptical of it should, unless they reject realist philosophy of science in general, agree that the adoption in physics of predicates of 'charge' and 'mass' is "costly" in the Occamist sense. And then, unless they claim some special exemption for metaphysics, they should agree that the adoption of modal operators is costly in the same sense. And then, unless they claim some special exemption for logic, they should agree that the adoption of higher-order languages is also costly.[12]

I have said that the higher-order viewpoint increases reality's complexity because it recognizes an ocean of new facts. But it might be objected that much of this ocean may not be new at all, in a higher-order sense of 'new'. For if propositions (in a higher-order sense) are sufficiently coarse-grained, many of the allegedly new propositions expressed by higher-order sentences will in fact be identical to old propositions recognized all along, propositions expressed by first-order sentences. In the most extreme case, if propositions are individuated maximally coarsely, by their truth value, then there will be just two propositions, The True and The False, so that no proposition expressible in the higher-order language is new.

This argument makes two mistakes. First, it evaluates a theory's complexity on the basis of the *truth* about grain, rather than what the theory *says* about grain. The parsimony argument is epistemic: convinced that the world is probably simple, we give more credence to theories that say that the world is simpler. A higher-order theory that says that grain is coarse might well gain credibility for that reason, but it doesn't matter whether the theory's propositions are in fact identical to old propositions. Second, the argument evaluates complexity solely on the basis of *propositional* grain. Two theories, each of which says that

[11]See Fine (2001) for another approach in the same quadrant of logical space, though Fine does not particularly emphasize parsimony.

[12]Timothy Williamson forcefully opposes both exemptions, for instance in the preface to *Modal Logic as Metaphysics*. Other works in the anti-exceptionalist tradition include Almog (1989); McSweeney (2019); Paul (2012); Quine (1948).

propositional grain is maximally coarse, might still differ in their complexity in an epistemically relevant way. One might posit more physical properties, or posit more complex structure at higher levels.

The complexity judgments on which the prima-facie case against higher-orderism depends are admittedly fraught. (Not that we have an alternative to relying on them. We must simply temper the certainty of our judgments.) Compare, for instance, a higher-order view positing a certain level of grain with a first-order ontology of sui generis propositions, properties, and relations (including properties and relations of those propositions, properties, and relations; and so on) of the same grain. Although I am tentatively inclined to judge the higher-order view more complex, because of its larger ideology (sui generis variables and quantifiers at each level), it is hard to place much weight on this. A more secure kind of judgment is that a given higher-order theory is more complex than a first-order theory of sets—ZF, say. (The latter sort of theory is what I myself would like to defend.) However, it matters what the higher-order theory says about grain. The judgment becomes more tenuous if the posited level of grain is coarse—for instance, if the higher-order view includes "Booleanism" (Dorr, 2016, section 7), according to which, roughly, boolean equivalents are identical (for instance, negation = triple negation: $\sim =^{((),())} \lambda p.(\sim\sim\sim p)$). Booleanism does seem pretty simple; and although I prefer the beautiful austerity of first-order ZF, it's hard to imagine betting my life on this.

We are in murky epistemological territory. However, in some dialectical contexts and in much of the remainder of this paper, we can rely on a quite secure judgment: that first-order ZF is simpler than higher-order theories that include second-order ZF. For then the higher-order theory seems to include all the complexity of the first-order theory, plus more in addition.

Let us turn, now, to arguments for higher-order-ism.

3. Argument from natural language

Some defend higher-order quantification by arguing that natural language already contains it. Boolos (1984) famously argued that natural language contains plural quantification; and Agustín Rayo and Stephen Yablo (2001) (following Arthur Prior (1971) and Dorothy Grover (1992)) argue that natural language contains devices tantamount to both monadic and polyadic second-order quantification. Just as Boolos claims that the natural language sentence 'Some critics admire only one another' doesn't carry a commitment to sets of critics, so Rayo and Yablo argue that 'Somehow things relate such that everything is so related to something' (the putative natural-language analog of $\exists R \forall x \exists y R(x, y)$) doesn't carry a commitment to relations (as entities).

But it isn't clear why any of this matters. If natural language doesn't contain higher-order quantification, couldn't we just introduce it, provided reality can support such talk? New concepts in physics (charge, spatiotemporal separation) are rarely introduced by defining them in pre-existing terms; rather, we lay out an inferential role for the concepts and posit that the role is filled. Conversely, if natural language does contain higher-order quantification, but reality *can't* support it, wouldn't natural language higher-order claims then be subject to some sort of discourse failure, or ultimately be made true by purely "first-order" facts? The real issue is whether reality can support higher-order talk, not whether natural language already has it. Moral and modal skepticism of various sorts persist (including error theories, expressivist theories, and aggressively reductive theories) despite the presence of modal and moral natural language; why should matters be different with higher-order quantification?

The primordial issue, as I say, is the metaphysical one of whether reality can "support" higher-order languages. That is, does reality contain higher-order facts? Do irreducibly higher-order claims make adequate contact with reality? Now, it is difficult to express this issue in some canonical way that is both theoretically satisfying and neutral on certain questions of metaphysics. As in the previous section, one might understand the issue in terms of a metaphysics of carving at the joints. But it's again important to see that such metaphysical baggage is not required. It is plain that there are analogous, gripping questions about both modal and moral language, and that those questions do not grip us solely because of an antecedent acceptance of any particular inflationary metaphysics.[13] So let us continue with an atheoretical, baggage-free statement of the question: does higher-order talk have the underpinnings in reality needed to be free from either reduction or discourse failure of various sorts?

4. The ZF argument

Boolos also gave a second argument, which has nothing to do with natural language, and which bears on the metaphysical standing of higher-order logic.

[13] It is compatible with this methodological point that the best way of understanding the questions in fact makes use of inflationary metaphysics.

The argument begins by assuming the correctness of the Zermelo-Frankel approach to set theory. (Thus the argument supports higher-order logic—second order logic, in the first instance—as a supplement to, not a replacement for, set theory; and the parsimony judgment it must overcome is the comparatively secure one that first-order ZF is simpler than higher-order theories that include second-order ZF.) In its standard, first-order axiomatization, ZF set theory contains the following axiom schema:[14]

$$\forall z \exists y \forall x \big(x \in y \leftrightarrow (x \in z \;\&\; \phi)\big) \qquad \text{(Separation)}$$

Despite its similarity to Naive Comprehension, Separation does not imply Russell's contradiction. Whereas Naive Comprehension says that any condition picks out a set, Separation says merely that any condition picks out a *subset* of any *given* set z. Substituting '$x \notin x$' for ϕ yields an instance saying that for any set z, there exists a set containing all and only the members of z that are not members of themselves. But this isn't contradictory. (Since another axiom of ZF guarantees that no set is a member of itself, this subset will simply be all of z.)

The idea behind ZF is to replace Naive Comprehension's method for "generating" sets with a two-step process of generation. Very roughly, certain axioms (the null set, pairing, unions, and powerset axioms, and the axiom of infinity) generate sets that are the "raw materials", which we then "cut to size" using the Separation schema, which lets us "carve out" arbitrary subsets of the sets generated by the other axioms.[15]

But there is an apparent problem with this procedure: the Separation schema doesn't really imply the existence of "arbitrary subsets". Rather, it generates only those subsets that can be picked out by a formula, ϕ, in the language of set theory. Thus the infinitely many instances of the Separation schema, taken together, fall short of expressing what one might have thought is the intuitive idea, namely that *any collection of members of a given set*, z, *forms a set*. Collections that are inexpressible by formulas are left out. (Since the language of set theory has an enumerable vocabulary and its formulas are finitely long, its set of formulas is enumerable, whereas the set of subsets of any infinite set is not enumerable.)

But in the intuitive idea that "any collection of members of a given set forms a set", what does "any collection" mean? It can't mean "any set", since

[14] As with Naïve Comprehension, y cannot be free in ϕ, but parameters are allowed (note 4).

[15] This metaphor is strained by the fact that the procedure iterates, and is further strained by the fact that Separation derives from the Replacement schema.

the comprehension principle would then become the trivial statement that any subset of a given set, z, forms a set. The higher-orderist has an answer: "any collection" can be understood in second-order terms: for any *property*, in a second-order sense of "any property", there exists a subset of any given set z consisting of z's members that have the property. Or, as Boolos himself argued, "any collection" can be understood in plural terms: for any things (universally quantified plural variable), there exists a subset of any given set z consisting of z's members that are among those things. Taking the former approach, we can replace the separation schema with a second-order separation axiom (not a schema, a single sentence):

$$\forall F \forall z \exists y \forall x \Big(x \in y \leftrightarrow \big(x \in z \,\&\, F(x) \big) \Big) \qquad \text{(Second-order Separation)}$$

The argument in favor of higher-order quantification, then, is that with it we can better formulate the axioms of ZF set theory:

> It is, I think, clear that our decision to rest content with a set theory formulated in the first-order predicate calculus with identity… must be regarded as a compromise, as falling short of saying all that we might hope to say. Whatever our reasons for adopting Zermelo-Fraenkel set theory in its usual formulation may be, we accept this theory because we accept a stronger theory consisting of a *finite* number of principles, among them some for whose complete expression second-order formulas are required. We ought to be able to formulate a theory that reflects our beliefs. (Boolos, 1984, p. 441)

Now, the alleged problem with first-order ZF cannot be that there are truths about sets that we can't state using its language. Assuming that *any* language we can speak has an enumerable vocabulary with finitely long sentences, there will always be too many truths about sets to state them all individually, no matter what language we speak.

The quotation suggests that the problem involves our *beliefs*: their statement requires expressive resources beyond first-order ZF. This is still weak: we might believe things that don't correspond to anything in the world, or are even incoherent. If the beliefs of some sect of medieval angelologists require distinctive language to state, this is no argument that the language is in good standing.

An initially more promising argument is that second-order ZF is a *better theory*, and that this gives us a reason to accept any conceptual resources needed to state it.

Compare the posits we make in the physical sciences for the sake of theoretical gain. The best reason to believe that there are fundamental physical properties and relations is that only by positing them can we state the laws of dynamics.[16] For instance, take the introduction of charge in classical electrodynamics. We observe regular patterns in the motions of things. When certain things get close to each other, they move apart—some quite sharply, others less so. When certain other things get close to each other, they move even closer together—some quite sharply, others less so. It is natural to posit that the particles have different properties, which we name charges; that there are laws relating force to charge; and that there are dynamical laws saying how things move as a function of the forces acting on them. Thus we posit properties like charge in order to be able to state strong laws of nature. If we didn't posit charge, we couldn't state laws of dynamics; we could only have list-like statements: "these particles moved in these ways, those moved in those ways, …".[17] Similarly, perhaps, we should treat higher-order quantification as a theoretical posit, like the posit of charge. Without it, we have only the list-like separation axioms of first-order ZF. With it, we can state the *law* of second-order separation.

Now, the analogy is flawed since the Separation axioms aren't really list-like: all instances of the schema share a certain syntactic form. The problem with stating laws of motion without properties like charge is that in the list of descriptions of particles' motions, no predicates beyond spatiotemporal ones occur, so there is no pattern; but in the list of instances of Separation, each instance has a rich syntactic structure, in virtue of which it makes sense to say that each shares a single syntactic form. So the set of instances of the Separation schema is *not* a complex, miscellaneous set in the way that the list-like set of descriptions of particle motions is. It exhibits a simple pattern; it is unified by a single syntactic property instantiated by all its members.

In fact, once we remember that our total theory includes laws of logic, not just laws of nature and laws of mathematics, it becomes clear that the sort of unification-by-pattern that we observed with the instances of the Separation schema is ubiquitous. The laws of propositional logic, for instance, are usually stated using schemas, such as:

$$\big(\phi \to (\psi \to \phi)\big)$$

[16] Nominalist restatement: those laws can be stated only by introducing fundamental physical predicates.

[17] See Sider (2020*b*, sections 4.4 and 4.12) for a discussion of some related issues.

That can be avoided by replacing the schema with a single axiom, involving particular sentences, such as 'Snow is white → (grass is green → snow is white)', and then adding a rule of substitution allowing us to infer a uniform substitution instance from any theorem (thus from the new axiom we may infer by substitution an arbitrary instance of the original schema, by replacing 'Snow is white' with ϕ and 'Grass is green' with ψ). But this just shifts the bulge in the carpet. For rules of inference themselves, even simpler ones like modus ponens, are not statements in the language of first-order logic, but rather are relations between formulas. The set of theorems of propositional logic form a simple, unified whole because of the patterns they collectively instantiate; and these patterns are not exhausted by a statement of the axioms that generate the whole, but in addition involve the fact that the set of theorems is closed under the rules. (Compare Lewis Carroll's (1895) point.) Their collective simplicity involves the holding of the relations that are the rules.

So the mere fact that the axioms of first-order ZF must be stated schematically does not detract from its theoretical simplicity. Its theorems, as a whole, exhibit a pattern, in part captured by the fact that each instance of the Separation schema shares a single syntactic form; and this sort of pattern is akin to patterns that unify any theory closed under rules of inference.

The ZF argument should not be understood as targeting the *simplicity* of first-order ZF, but rather its *strength*. (This is clearly how Boolos understood it, though perhaps not with my spin, analogizing to dynamical laws.) "Schematic laws"—that is, sets of statements unified by a syntactic property—can be simple, as we've seen; but the instances of the Separation schema are collectively weaker than the second-order axiom of Separation.[18] So the argument is this. A good theory should have laws that are both simple and strong; and although first-order ZF has laws that are simple (in some cases in the schematic sense that we have been discussing) and somewhat strong, it does not match second-order ZF's combination of simplicity and strength.

The argument can't be left there, as will emerge if we consider how Separation is applied. Let A and B be any two sets. A simple theorem of ZF is the statement that there exists such a set as the intersection of A and B—the set of things that are elements of both A and B. In first-order ZF, this is proven by

[18]Weaker in a model-theoretic sense, for instance (assuming the full interpretation of the second-order quantifiers): the Löwenheim-Skolem theorem implies the existence of countable models of first-order ZF, but there are no such models of second-order ZF.

beginning with this instance of the Separation schema:

$$\forall z \exists y \forall x \big(x \in y \leftrightarrow (x \in z \,\&\, x \in B) \big)$$

(which we obtain by changing the schematic letter ϕ to "$x \in B$".[19]) Then we let z be A, and infer:

$$\exists y \forall x \big(x \in y \leftrightarrow (x \in A \,\&\, x \in B) \big) \tag{1}$$

y is our desired intersection of A and B.

The proof in second-order ZF is a little different. We can begin by inferring the following from the second-order Separation principle (letting z be A):

$$\forall F \exists y \forall x \Big(x \in y \leftrightarrow \big(x \in A \,\&\, F(x) \big) \Big) \tag{2}$$

Next we need to instantiate the variable F to the property of "being a member of B". But how do we do that? The answer depends on what *logical* principles we take to govern the second-order quantifiers.

Typical axiomatizations of second-order logic include the following schema (where ϕ may be replaced with any formula with no free occurrences of variables other than x):[20]

$$\exists F \forall x \big(F(x) \leftrightarrow \phi \big) \tag{Comprehension}$$

(Despite its similarity to Naive Comprehension, this schema does not lead to Russell's paradox given the grammar of second-order logic.) An instance of Comprehension (letting ϕ be '$x \in B$') is:

$$\exists F \forall x (F(x) \leftrightarrow x \in B)$$

which, together with (2), implies (1).

[19] Really we need the version of the schema with parameters—see note 4.

[20] And perhaps parameters; see note 4. An axiom schema of Comprehension is not needed in a language with λ abstraction, a rule of universal instantiation in which the universally quantified variable can be instantiated to λ abstracts, and the schema of "β conversion": $\lambda v.(A)\alpha \leftrightarrow A_v(\alpha)$, where $A_v(\alpha)$ is the result of changing vs to αs in A (in accordance with the usual restrictions). (To derive an arbitrary instance of Comprehension, begin with $\forall G \exists F \forall x (Fx \leftrightarrow Gx)$; infer $\exists F \forall x (Fx \leftrightarrow \lambda x.(A)x)$ by universal instantiation; then use β conversion to derive the instance.) But this wouldn't affect the argument, since the theory of properties would share the same expressive weakness as the theory based on Comprehension: it would only imply the existence of properties corresponding to formulas.

Thus in the second argument we still needed a schema, namely the logical axiom schema of Comprehension, to reach the desired conclusion. Now, given what we have learned, the need for a schema in the overall theory here—second-order ZF, including the underlying logical theory—does not compromise its simplicity. But there is a question of strength. Although second-order ZF (plus its logic) has simple and strong laws governing the existence of *sets*, it would seem to lack simple and strong laws governing the existence of *properties*. Better: its theory of the existence of properties is on a par with first-order ZF's theory of the existence of sets. It is schematically simple; and it is somewhat strong; but its strength is limited by the inherent limits of schemas in enumerable languages. The failure of strength occurs, to be sure, within the logical part of the theory (if we count second-order logic as logic); but it's hard to see how that matters.

The challenge facing the ZF argument, then, is to say why the elimination of the weakness in first-order ZF justifies the posit of the second-order quantifiers, when the resulting theory, second-order ZF, is weak in a parallel way.

The second-orderist might try to meet the challenge by saying that, while there is a way to eliminate the weakness in first-order ZF (by moving to the second order), there is no conceivable way to eliminate the weakness in second-order ZF. But that just isn't true. We could, for instance, posit a sort of "super-second-order quantification", and use it to state a principle of plenitude for the original second-order quantifiers: for any super-property there is a corresponding property. The second-orderist will presumably regard such additions as misguided, since they too will need a comprehension schema, just as the second-order quantifiers needed one, and thus will have a weakness structurally like the one they were trying to avoid in first-order ZF. (The alternative would be to formulate a series of theories, each allegedly superior to the last, based on second-order quantification, super-second-order quantification, super-duper-second-order quantification, and so on; but surely this is unattractive, not least because of its complexity.) But this invites the question of whether we should say the same thing about the shift from first- to second-order ZF and the addition of the second-order quantifiers—that *this* addition is also misguided, given the structural similarity between the weakness of the Comprehension schema in second-order logic and the weakness of the Separation schema in first-order ZF that we were trying to avoid. Why add the second-order quantifiers to deal with a problem that will only re-arise?

I am not denying that second-order ZF would be a stronger theory (in the sense of "strength" relevant to theory choice as described above), if its vocabulary were in good standing.[21] It would indeed be stronger, for it would contain a sentence saying that there is a set for every property. But notice that the description of the added strength uses the very vocabulary (namely, second-order quantification) whose legitimacy is at issue. It isn't as if first-orderists can be convicted by their own lights of leaving a lawlike generalization out of their theory, since the allegedly omitted generalization ("there is a set for every property") isn't even statable in their language.

The second-orderist claims that first-order ZF has a kind of "gap" (because certain of its existence claims about sets must be merely schematic) and that this gap should be filled by this law-like statement: 'for every property, F, there is a set, z, such that an object x is a member of z iff $F(x)$'. The gap-filler is not stateable in the vocabulary of first-order ZF, and thus its existence as a possible content is not common ground; the second-orderist is trying to simultaneously persuade us of the existence of the gap and the means to fill it. But second-orderists are unmoved by the attempt to simultaneously persuade them of the existence of a gap in their theory and the means to fill it: a missing law-like statement 'for any super-property $\mathscr{F}$, there is an F such that for all x, $F(x)$ iff $\mathscr{F}(x)$', and the associated additional vocabulary of super-second-order quantification. The crucial point is that the allegation of weakness involves allegedly omitted content that is only recognized by the accuser.

The second-orderist is dangling a carrot: here is some new vocabulary you could adopt (the second-order quantifiers), and in terms of it, new generalizations you could state. Moreover, an axiom schema you accept (Separation) can be subsumed under a single principle statable in the new vocabulary (each instance of Separation is entailed by the second-order Separation axiom; the argument of course uses the second-order comprehension schema). But the carrot still dangles, even after the move to the second order. When we appreciate this dialectical situation, we should realize that we should never have followed the carrot in the first place. For after all, it will be forever out of reach.

(Some will insist that the notion of a *property* is intuitively legitimate, or well-understood, or needed, or whatever, in a way that neither the notion of set nor the notion of super-second-order quantification are; and that this breaks the symmetry I've been emphasizing between the move from first- to second-order ZF, and the move from second- to super-second-order ZF. Fair enough, but that would be a different argument. We are here considering the ZF argument taken in isolation.)

[21] Skolemite metasemantic challenges are not at issue here.

It is fruitful to compare this dialectic with a similar one involving laws of nature. Consider the law that like-charged particles repel one another. "Deflationists" about laws of nature, such as defenders of the Humean or best-system theory, think that this amounts to nothing more than the regularity that all like-charged particles in fact repel, plus some bells and whistles.[22] "Inflationists" about laws, on the other hand, think that there is some kind of further fact, the law, which explains the regularity that like-charged particles repel. But even the inflationists reject an extreme inflationism, according to which a good explanation of the regularity requires a still further fact, a **Meta-Law** governing the law; for the **Meta-Law** seems to be explanatorily superfluous.[23] Further, although deflationists regard inflationists' robust law as explanatorily superfluous, many of them reject the extreme deflationism of someone like Michael Esfeld (2020), who thinks that the posit of charge is superfluous since we could just as well state the law of motion by saying (roughly) that particles move in certain ways, namely the ways they would move if there were a property of charge; according to Esfeld, adding that the differences in motion are due to differences in charge does not improve the explanation. It is not obvious who is right in this dialectic; the answer turns on difficult questions about explanation.[24] But I suspect that most will agree either with the standard deflationist or the standard inflationist, and will reject both extreme deflationism (on the grounds that eliminating charge from physics is a genuine explanatory loss) and certainly extreme inflationism (on the grounds that the **Meta-Law** is explanatorily superfluous).

To my mind, the comparison to this dialectic about laws of nature weakens the ZF argument. What seems particularly objectionable about extreme inflationism is the parallelism between its explanation of regularities and the inflationist's explanation: the two are exactly alike except for an added layer in the former case. But the same objection, as we saw, seems to apply to the defender of second-order ZF. Unlike the positing of charge, which *does* result in a structurally different explanation (this is partly why Esfeld does not convince), the positing of the second-order quantifiers does not seem to result in a structurally different explanation.

[22] See, for instance, Lewis (1994).

[23] Some inflationists, such as Lange (2007), accept laws governing other laws in certain cases, but this doesn't affect the point. Lange doesn't think that laws cannot govern without meta-laws, or that his meta-laws need further meta-meta-laws.

[24] See Dorr (2007, section 3; 2010, p. 160–3) and Sider (2020*b*, section 4.12).

5. The argument from semantics

The status of the ZF argument is somewhat parallel to that of another argument for higher-order languages, namely that such languages are needed to give a semantic account of unrestricted quantifiers.[25]

The standard approach to semantics in logic—to giving an account of how sentences come to be true and false, and of when sentences imply one another in virtue of meaning—is the model-theoretic approach. A model is normally defined as an ordered pair $\langle D, I\rangle$, where D, the domain, is a set, and where I, the interpretation function, is a function that assigns to each nonlogical expression in the language some appropriate set-theoretic construction based on D, such as members of D to names, and sets of n-tuples of D (extensions) to predicates. Using methods developed by Tarski, one can define what it means for an arbitrary sentence of the language to be "true in" such a model.

Thus the standard approach defines the domain of any model, and the extension of any predicate in any model, as sets. But that is limiting. In the intended interpretation of the language of first-order ZF set theory, for instance, the quantifiers range over all sets and '$x \in y$' means that x is a member of y. So the domain of a model corresponding to this intended interpretation should be a set containing all sets, and the extension of '$\in$' in this model should be a set containing all and only ordered pairs $\langle x, y\rangle$ where x is a member of y. But there are no such sets according to ZF set theory (which is assumed in the metalanguage), since either would lead to Russell's contradiction.

Given the standard approach, then, no model captures the intended interpretation of the language of first-order set theory. At best, models can provide a distorted representation of that intended interpretation, by treating the quantifiers as being restricted to a mere part of the set-theoretic hierarchy.[26]

Against this backdrop, the higher-order outlook becomes attractive. For one can, in a second-order language, define a sort of "model" in which the "domain" can contain all the sets, and in which the "extension" of '$\in$' contains all and only the ordered pairs $\langle x, y\rangle$ where x is a member of y. The trick is to abandon

[25] See, for instance, Williamson (2003).

[26] There are subtle arguments that this limitation does not affect which sentences the standard approach counts as valid or implying one another. But such arguments are less clearly correct when the languages move beyond the first order, and break down if the language contains certain sorts of expressions. And even if it always gives the right answers, the standard approach would seem still to be mis-modeling the semantic facts. See Boolos (1985); McGee (1992); Rayo and Uzquiano (1999) for discussion.

first-order quantification over models conceived as entities, and instead to treat quantification over models as being second-order.[27] We define 'R is a model', for R a second-order dyadic variable, in such a way that when R is a model, '$R(x,y)$' can be thought of as meaning that y is a semantic value of the linguistic expression x. So if 'S' is a two-place predicate, then "R('S', $\langle u,v\rangle$)" can be thought of as meaning that $\langle u,v\rangle$ is "in the extension of 'S' in R", although this is misleading since we are not accepting the existence of an entity, the extension of 'S' in R. In one of these second-order "models" the "extension" of the two-place predicate '$\in$' will consist of all and only the ordered pairs whose first coordinate is a member of the second coordinate. That is: for some model R, R('$\in$', z) if and only if z is an ordered pair $\langle u,v\rangle$ such that u is a member of v. No paradox results because we do not recognize an entity as the extension of '$\in$'. One can then define the notion of an arbitrary sentence in the language of first-order set theory being true in such a "model" R.

By adopting a second-order language, we can thereby state more adequate semantic theories for first-order languages. But we will then naturally aspire to state semantic theories for second-order languages, such as the metalanguage just used in the semantic theory for the first-order language. And as Øystein Linnebo and Agustín Rayo (2012) argue, if we want to acknowledge the full range of "models" for second-order languages, a third-order metalanguage will be needed. Let α be some one-place second-order predicate constant in the second-order language (i.e., a predicate that can be attached to one-place predicates—i.e., type $\big((e)\big)$). For any property, $\mathscr{F}$, of properties (variable of type $\big((e)\big)$), it would be possible to interpret α so that it applies to exactly the properties G (variable of type (e)) such that $\mathscr{F}(G)$. So for each such $\mathscr{F}$, there must be a new model R. But it can be shown that there are more such $\mathscr{F}$s than there are second-order relations R. (The argument is analogous to the usual Cantorian diagonal argument showing that a set has strictly lower cardinality than its power set, but here the "cardinality comparison" and argument are made in a higher-order language.) The models for the second-order language must be third-order relations.[28]

[27] See Boolos (1985); Rayo and Uzquiano (1999); Williamson (2003).

[28] The argument, notice, is that a higher-order metalanguage is needed to recognize all the possible interpretations for the original language. If we only wished to give the *intended* interpretation of the second-level language, a second-level metalanguage could be used (although given Tarski's theorem it would require a new primitive semantic predicate). But an explanatory theory of meaning should include a theory of all the semantic possibilities, of how variation in the meanings of the parts induces variation in the meanings of wholes, just as an explanatory

In fact, Linnebo and Rayo show how to generalize this argument into the transfinite. For each language in a certain transfinite hierarchy (and thus languages with syntactic types beyond those that I defined earlier, each of which was finite), stating its semantics requires a still higher order metalanguage.

This fact weakens the argument from semantics, if that argument is taken in the metaphysical spirit we have been considering. We began with an explanatory ambition: to give a certain sort of semantic theory for the language of first-order set theory. If we take that ambition to justify recognizing second-order quantification, with its attendant worldly complexity[29], then we are saddled with a new explanatory ambition, the satisfaction of which requires a new language, which results in a new explanatory ambition; and so on. The explanatory demand is insatiable, in that there is no language we could speak in Linnebo and Rayo's hieararchy in which we could state a semantics for all languages in that hierarchy.

Will we eventually resist the explanatory demand, and say that for some final language L_α, no explanatory semantic theory can be given (after having recognized all the worldly complexity of the preceding higher-order languages)? Embracing explanations up to L_α but no further seems akin to embracing, in addition to robust laws, also meta-laws, meta-meta-laws, and so on, up to a certain point and then stopping; or embracing, in addition to sets, second-order quantification, super-second-order quantification, super-duper-second-order quantification, and so on, and then stopping. If we must eventually resist the explanatory demand for L_α, wouldn't it have been better to resist it right at the start, with the language of first-order set theory? Doing so wouldn't mean saying that this language is meaningless, or doesn't quantify over all sets after all; it rather means that we won't give a certain sort of *theory* of the meaningfulness of that language. Nor does it mean that no theory whatsoever of its meaningfulness is available. The usual models of model theory can still be models in the philosophy-of-science sense, albeit imperfect ones, of meaning. And we could still give other sorts of semantic theories, such as a Davidsonian (1967) theory of truth for the language of set theory. It's a bit sad, but not the end of the world.

The alternative would be to continue to acquiesce to the explanatory de-

physical theory should encompass all the physical possibilities, not just the actual one. Good explanations must be general.

[29]Here again we rely only on the most secure sort of complexity judgment from section 2: that a higher-order theory that includes an ontology of sets is more complex than first-order set theory.

mand, formulating more and more semantic theories in increasingly higher-order languages, limited only by patience and lifespan. There is no vicious regress here, since the theory we formulate at a given stage does not convey meaning on the preceding language; rather, the meaningfulness of the preceding language is taken as a pre-existing fact, which the subsequent language is used to explain. The series of semantic explanations might be compared to a series of causal explanations of what occurs at some time in terms of what occurs at some preceding time—a series that we might legitimately continue indefinitely.[30] I don't wish to deny that explanatory gains could indeed be made at each stage in this series of semantic explanations. But each explanatory gain comes with a cost in complexity; and when we look ahead at the looming indefinite series, the appeal of its first step is diminished, in comparison with cutting the whole thing off at the start and sticking with first-order logic and imperfect semantics.

6. Arguments from categoricity

Second-order logic has distinctive model-theoretic properties.[31] For instance, using a second-order language containing expressions for 'zero', 'successor', 'plus', and 'times', one can state a theory of the arithmetic of natural numbers that is "categorical": any two of its models are isomorphic. (Each model has the familiar "shape" of the natural number line: all elements are arranged in a line in which each element is reachable from an initial element by some finite number of discrete jumps.) But in a first-order language using these nonlogical expressions, no theory, not even one containing infinitely many axioms, has this feature. If such a theory has a model of the familiar sort, it will also have models of arbitrarily large cardinality, and even "nonstandard" models which fail to be isomorphic to the familiar sort despite having the same number of elements.

Facts like these weigh heavily with some fans of second-order logic. But they don't add up to a convincing argument that second- (or higher-) order logic is in good standing.

One argument that might be based on the facts is metasemantic: without second-order resources, arithmetic language could not have the determinate

[30] Thanks to Timothy Williamson here.

[31] For the formal results alluded to in this section, see Shapiro (1991, chapter 4) and Väänänen (2019).

interpretation that it in fact has. If our theory were merely first-order, according to this argument, nothing would rule out intended or correct interpretations based on nonstandard models.

This argument appears to presuppose an "interpretationist" (Williams, 2007) approach to metasemantics, according to which the only constraint on the correct interpretation of arithmetic language is that our theory of arithmetic come out true under the interpretation. But interpretationism isn't true for *all* language.[32] If it were, then any consistent theory of physics (for instance) would have a correct interpretation under which it comes out true, provided there are enough entities in the world. So if the argument is to even get off the ground, it must provide a reason to think that interpretationism is true of mathematical language despite being false generally. For instance, it might be thought that in physics, the additional constraint on correct interpretation is causal in nature—that our usage of physical predicates must bear certain causal relations to their putative semantic values—whereas causal constraints are inapplicable to mathematical language.

Thus understood the argument relies on the contentious assumption that the only additional constraint on correct interpretation is causal. That premise would be rejected, for instance, by David Lewis (1983, 1984), whose proposed additional constraint *is* applicable to mathematical language: correct interpretations must, other things equal, assign "natural" properties and relations as semantic values.

And that isn't the argument's only contentious assumption. For although the models of second-order arithmetic are all isomorphic, they are not identical. If nothing constrains the correct interpretation of arithmetic language beyond that our arithmetic theory must come out true, then there are no constraints on *which* particular entities count as natural numbers, or on which particular entities count as which natural numbers. There will be correct interpretations, for instance, in which Julius Caesar is the denotation of the symbol 'zero'. Thus the argument cannot be: "arithmetic language has an *absolutely* determinate interpretation, and only with second-order resources can this be secured", since even with second-order resources, arithmetic would not be absolutely determinate (given interpretationism for arithmetic language).[33] It must rather be: "arithmetic language, although not absolutely determinate, is determinate

[32]This is the (unintentional) lesson of Putnam's "model-theoretic argument (1978, part IV; 1980; 1981, chapter 2).

[33]Compare Weston (1976, section V).

up to ismorphism, and only with second-order resources can this be secured". Thus the argument must rely on a sort of "structuralism".

When we turn from arithmetic to set theory, even more structuralism will be required, so to speak. For in that case, the second-order mathematical theory—namely, second-order ZF—does not quite constrain its models up to isomorphism. Rather, all that is guaranteed is that in a certain sense, any two models have isomorphic initial segments. The argument in this case would thus need to rely on the premise that set-theoretic language is that determinate, but no more.

But all this skirmishing is beside the main point, which is the argument's uncritical stance towards the second-order quantifiers.[34]

The claim that the models of second-order arithmetic are all isomorphic relies on certain definitions from (standard, set-theoretic) model theory. In the "standard" (or "full") definition of model for second-order logic, the monadic second-order quantifiers are treated as ranging over all subsets of the domain; it is this definition that was presupposed above. But if the second-order quantifiers are instead treated as ranging only over a restricted range of subsets of the domain, with the range varying from model to model (as in "general", or "Henkin" second-order models), nonisomorphic models of second-order arithmetic reappear. The standard definition of model is appropriate given the interpretation of the second-order quantifiers that the higher-orderist advocates, namely as ranging over absolutely all properties (to put it intuitively); but what is the metasemantics of that interpretation? Its correctness cannot be secured solely by our putting forward the theory of second-order logic (including all instances of the comprehension scheme), since that theory has non-full Henkin models.

The second-orderist seems to be exempting the second-order quantifiers from interpretationist metasemantics. But if they are exempt, then why not arithmetic or set-theoretic language as well? "The second-order quantifiers are part of logic" is no answer without a specification of the scope of "logic" under which the second-order quantifiers but not arithmetic or set-theoretic vocabulary count as logical, and without an argument that logical, and only logical, vocabulary thus understood is exempt from interpretationism. The metasemantic argument is a dialectical failure, since the higher-orderist is in no position to object to a first-orderist who denies interpretationism, and claims that the languages of first-order arithmetic or set-theory are determinate

[34] Again compare Weston (1976).

(whether absolutely or in some restricted sense) despite having nonstandard models.

Instead of the metasemantic argument, the higher-orderist might simply argue that the virtue of second-order logic is that it allows us to single out, by purely logical means, the intended class of structures in arithmetic, and to nearly single out the intended class of structures in set theory. But what is so special about singling out these structures by "logical" means? Assuming that the determinacy of set-theoretic vocabulary is no longer in question (we have left the metasemantic argument behind), even without second-order logic we can single out the arithmetic structures by *set-theoretic* means, and we can of course single out the unique set-theoretic "structure" by set theoretic means.

7. The collapse argument

I have been questioning whether higher-order quantifiers are in good standing. But suppose we construe "not in good standing" as indeterminacy, in the sense of there being multiple candidate meanings, or precisifications, for a given higher-order quantifier, no one of which is determinately meant.[35] One might then try to answer my challenge with a "collapse argument".

Collapse arguments aim to show that for certain logical connectives, any two meanings that individually obey certain inference rules for the connective must be equivalent. Following Cian Dorr, we will understand talk of these meanings in higher-order terms.[36] For instance, to quantify over candidate meanings for '&', we use variables c of the type of that connective, namely $\big((),()\big)$. Say that any c "obeys conjunction introduction" if for any propositions p and q, p and q together entail $c(p,q)$, and that c "obeys conjunction elimination" if $c(p,q)$ entails p and also entails q, for any propositions p and q. Then consider any c_1 and c_2 each of which obeys conjunction introduction and elimination, and let p and q be any propositions. Since c_1 obeys conjunction elimination, $c_1(p,q)$ entails p and also entails q; and since c_2 obeys conjunction introduction, p and q together entail $c_2(p,q)$; thus (given standard assumptions about entailment), $c_1(p,q)$ entails $c_2(p,q)$. A similar argument (using conjunction elimination for

[35] This needn't be tied to supervaluationism; all that is assumed is that indeterminacy requires precisifications.

[36] See Dorr (2014). Dorr considers the collapse argument for the first-order quantifiers; and that argument isn't subject to the difficulty I raise here, because it's less plausible that the higher-order quantifiers Dorr uses in the argument are penumbrally connected to the first-order quantifiers.

c_2 and conjunction introduction for c_1) shows that $c_2(p,q)$ entails $c_1(p,q)$. Thus c_1 and c_2 always generate mutually entailing propositions.

A collapse argument for a higher-order existential quantifier 'Q' could be given as follows. Consider any two candidate meanings for 'Q', Q_1 and Q_2, which we may take to be properties of properties (recall the end of section 1). Thus if 'Q' is a quantifier over "τ-entities" (for some type τ), it is of type $\big((\tau)\big)$; the variables 'Q_1' and 'Q_2' are therefore of this type.[37] Let 'F' be a predicate variable of τ-entities, i.e. a variable of type (τ). Continuing to follow Dorr (2014, section 5), the argument assumes that any candidate existential quantifier meaning Q over τ-entities must obey existential introduction and existential elimination—"intro" and "elim", for short—understood as follows:

Q obeys intro $=_{\text{df}}$ Any property F entails the property being such that $Q(F)$

Q obeys elim $=_{\text{df}}$ If a property F entails the property of being such that a certain proposition, p, is true, then $Q(F)$ entails p

where 'F' has type (τ) and 'p' has type $()$. (In addition to an entailment relation over propositions we are employing an entailment relation over properties. As Dorr explains, there are different ways this can be understood.) Then for any F, since Q_2 obeys intro, F entails the property of being such that $Q_2(F)$; but since Q_1 obeys elim, the proposition $Q_1(F)$ entails the proposition $Q_2(F)$. A similar argument (using intro for Q_1 and elim for Q_2) shows that $Q_2(F)$ entails $Q_1(F)$. Thus Q_1 and Q_2 generate mutually entailing propositions from any property F.

The problem with this attempt to show that the expression 'Q' is not vague is that, as we will see, it uses terms that are, plausibly, "penumbrally connected" (Fine, 1975) to that very expression.[38] The following obviously fallacious argument illustrates how arguments of this sort can go wrong:

> Any precisification of 'bald' must apply to all and only the things that instantiate the property of baldness. Thus any two precisifications, b_1 and b_2, must be extensionally equivalent: $b_1(x)$ iff x instantiates baldness iff $b_2(x)$.

[37] The quantifiers binding 'Q_1' and 'Q_2', or rather, attaching to λ abstracts for properties of entities of this type, therefore have the type $\Big(\big(((\tau))\big)\Big)$.

[38] This is not to deny that collapse arguments can play a useful role in other contexts, for instance in justifying the introduction of new expressions using old expressions whose interpretation is not in question.

or even:

> Any precisification of 'bald' must apply to all and only the bald things. Thus any two precisifications, b_1 and b_2, must be extensionally equivalent: $b_1(x)$ iff x is bald iff $b_2(x)$.

The fallacy derives from the penumbral connections between the expression whose vagueness is at issue (namely 'bald') and the expressions used in the arguments (the singular term 'baldness' in the case of the first argument, and the predicate 'bald' in the case of the second argument, which of course is penumbrally connected to itself). Insofar as it is compelling at all, the thought that any precisification of 'bald' must apply to all and only the things that instantiate baldness, or to all and only the bald things, is based on the idea that the following claims are "analytic" (to use a placeholder):

> Something is bald iff it instantiates baldness

> Something is bald iff it is bald

But what the analyticity of these two sentences requires is that each sentence be true under every global precisification of the language. (All that matters about our placeholder notion of "analyticity" is that it implies truth on all precisifications. It could even just mean "truth on all precisifications".) That is, where i is any global precisification, which includes a precisification t_i of each vague term t, the analyticities require:

> Something is bald$_i$ iff it instantiates baldness$_i$

> Something is bald$_i$ iff it is bald$_i$

But neither of these claims implies, when $i \neq j$, that something is bald$_i$ iff it is bald$_j$.

In broad strokes, the problem for the collapse argument is similar. The thought that any precisification of 'Q' must obey intro and elim is based on the idea that 'Q obeys intro and elim' is "analytic", and must therefore be true on all global precisifications. But it is plausible that if 'Q' is vague, then certain expressions used in the collapse argument—namely, certain other higher-order quantifiers, and 'entails'—are also vague, and indeed, penumbrally connected to one another and to 'Q'. Thus what the analyticity of 'Q obeys intro and elim' requires is that *for each global precisification* i, Q_i *obeys intro*$_i$ *and elim*$_i$, where

'intro_i' and 'elim_i' are the results of substituting terms for the i-precisifications of the vague terms in the definitions of 'intro' and 'elim'. That is, Q_i obeys intro_i and elim_i, Q_j obeys intro_j and elim_j, and so on. But then the collapse argument that $Q_i(O)$ and $Q_j(O)$ (for arbitrary O) are mutually entailing will break down. For that argument involves "diagonally" applying intro and elim: first intro for Q_i and elim for Q_j are combined, and then intro for Q_j and elim for Q_i are combined. But what we derive from the obedience of intro_i by Q_i will contain expressions subscripted with i, whereas what we derive from the obedience of elim_j by Q_j will contain expressions subscripted with j, which blocks their combination; and similarly for the second half of the argument.

To see in detail how this plays out in the collapse argument, we'll need to make some concrete choices. (I believe my conclusions to be robust, but I won't consider alternate choices.) The first choice has to do with the notion of entailment used in the collapse argument. For higher-order variables 'x' and 'y' of any single type, I will understand 'x entails y' as meaning that x is identical with the conjunction of x and y (compare Dorr (2014)). Thus a proposition p entails a proposition q (variables of type ()) iff $p = (p \,\&\, q)$, and a property of propositions O entails a property of propositions P (variables of type $\big(()\big)$) iff $O = \lambda p.\big(O(p) \,\&\, P(p)\big)$. The second choice is to define higher-order identity as the sharing of all properties. Thus, for instance, $p = q$ ("proposition p is identical to proposition q") will be understood as meaning $\forall O\big(O(p) \leftrightarrow O(q)\big)$ ("p and q have the same properties").

Consider, now, the case of the collapse argument where 'Q' is the propositional quantifier. Applying the choices of the previous paragraph, the definitions of obeying existential introduction and existential elimination become:[39]

$$Q \text{ obeys intro} =_{\text{df}} \forall O\, \forall X\Big(X(O) \leftrightarrow X\Big(\lambda q.\big(O(q) \,\&\, Q(O)\big)\Big)\Big)$$

"for any property O of propositions, O has the same properties as its conjunction with being-such-that-$Q(O)$"

$$Q \text{ obeys elim} =_{\text{df}} \forall O \forall p\Big(\forall X\Big(X(O) \leftrightarrow X\Big(\lambda q.\big(O(q) \,\&\, p\big)\Big)\Big) \\ \rightarrow \forall P\Big(P\big(Q(O)\big) \leftrightarrow P\big(Q(O) \,\&\, p\big)\Big)\Big)$$

[39] I have simplified by β-converting.

"for any property O of propositions and any proposition p, if O has the same properties as its conjunction with being-such-that-p, then $Q(O)$ has the same properties as its conjunction with p"

where 'Q' and 'X' are type $\big((())\big)$, 'p' and 'q' are type (), and 'O' and 'P' are type $(())$.

The thought that any precisification of 'Q' must obey intro and elim is based on the idea that 'Q obeys intro and elim' is "analytic". Thus that sentence must be true on all global precisifications. But it is plausible that if 'Q' is vague, then all the higher-order quantifiers are vague, and, moreover, are penumbrally connected to one another, so that in any global precisification, each higher-order quantifier must receive a corresponding precisification. (For instance, perhaps one global precisification corresponds to a conception of higher-order entities individuated extensionally, another to a modally-individuated conception, and so on.) Thus what the analyticity requires is that for each global precisification i we have:

$$\forall_i O\, \forall_i X\Big(X(O) \leftrightarrow X\big(\lambda q.(O(q) \,\&\, Q_i(O))\big)\Big) \qquad (\text{intro}_i)$$

$$\forall_i O \forall_i p\Big(\forall_i X\big(X(O) \leftrightarrow X(\lambda q.(O(q) \,\&\, p))\big) \rightarrow \forall_i P\big(P(Q_i(O)) \leftrightarrow P(Q_i(O) \,\&\, p)\big)\Big) \qquad (\text{elim}_i)$$

Let's now examine the obstacles facing the collapse argument that for any precisifications Q_1 and Q_2 of 'Q' and any property, O, of propositions, $Q_1(O)$ and $Q_2(O)$ are mutually entailing. The first obstacle is this: what meanings of the higher-order quantifiers will we use to state the argument? I just used higher-order quantifiers to quantify over precisifications of 'Q' and over properties of propositions. Are these the original, potentially vague quantifiers; or are they to be understood under one of the global precisifications?

It might be thought that we can simply choose a "large enough" global precisification, 3, so that Q_1 and Q_2, for instance, are both in the range of the quantifier $\forall_3$ of the appropriate type. Perhaps this would be workable, but the picture of quantifier precisifications "containing" one another is suspect. Compare a putatively "small" propositional quantifier, on which propositions are individuated by truth value, with a putatively "large" propositional quantifier on which propositions are individuated modally, by necessary equivalence. This is as good a case as any for the containment picture. But which of the true propositions recognized under the large quantifier should be identified with the sole true proposition ("The True") of the small quantifier?

(A related further obstacle is how to understand the quantification over global precisifications, since it appears to involve a sort of type-general quantification that isn't allowed in the higher-order languages considered in this paper. But this isn't a very deep obstacle, since the proponent of the collapse argument wouldn't really need to take the quantification over global precisifications seriously, and could instead regard the indices i as schematic.)

Continuing with the collapse argument: to show that $Q_1(O)$ entails $Q_2(O)$, the argument applies elim to Q_1 and intro to Q_2. And as we saw, what we in fact have is that Q_1 obeys elim$_1$ and Q_2 obeys intro$_2$. The former is this:

$$\forall_1 O \forall_1 p\Big(\forall_1 X\Big(X(O) \leftrightarrow X\big(\lambda q.(O(q) \,\&\, p)\big)\Big) \rightarrow \forall_1 P\Big(P\big(Q_1(O)\big) \leftrightarrow P\big(Q_1(O) \,\&\, p\big)\Big)\Big) \tag{elim$_1$}$$

The idea would be to instantiate $\forall_1 O$ to O and $\forall_1 p$ to $Q_2(O)$, yielding:

$$\forall_1 X\Big(X(O) \leftrightarrow X\big(\lambda q.(O(q) \,\&\, Q_2(O))\big)\Big) \rightarrow \forall_1 P\Big(P\big(Q_1(O)\big) \leftrightarrow P\big(Q_1(O) \,\&\, Q_2(O)\big)\Big) \tag{*}$$

Now, even this is potentially suspect, because of the first obstacle. 'O' and '$Q_2(O)$' are tied to the higher-order meanings we are using to state the argument—some global precisification 3, perhaps. What guarantee is there that these meanings are in the range of the appropriate 1 quantifiers, so to speak? Better: when we are speaking under the 3 precisification, what reason is there to assume that '$\forall_1 O$' and '$\forall_1 p$' obey universal instantiation with respect to O and $Q_2(O)$, respectively?

The next move would then be to derive the antecedent of (*) from the fact that Q_2 obeys intro$_2$:

$$\forall_2 O\, \forall_2 X\Big(X(O) \leftrightarrow X\big(\lambda q.(O(q) \,\&\, Q_2(O))\big)\Big) \tag{intro$_2$}$$

by instantiating $\forall_2 O$ to O. Here, too, the first obstacle rears its head: why think that O is "in the range of" '$\forall_2 O$'? But there is a further obstacle, which to my mind is the most serious. Even if this use of universal instantiation is allowed, the result wouldn't be the antecedent of (*), but rather:

$$\forall_2 X\Big(X(O) \leftrightarrow X\Big(\lambda q.\big(O(q) \,\&\, Q_2(O)\big)\Big)\Big) \tag{+}$$

This says that O and its conjunction with being-such-that-$Q_2(O)$ have the same 2-properties, whereas the antecedent of (*) says that they have the same 1-properties.

It might be replied that (+) entails the antecedent of (*) if the 1-properties are "contained in" the 2-properties. But even setting aside our concerns about this notion of containment, the respite would be temporary. For in the other half of the collapse argument we will use the facts that Q_2 obeys elim$_2$ and Q_1 obeys intro$_1$, and then we will have the reverse situation: the counterpart of (+) will now have the 1 quantifier and the counterpart of (*) will have the 2 quantifiers, and there will be no case (from "containment", anyway) that the former entails the antecedent of the latter.

References

Almog, Joseph (1989). "Logic and the World." In Joseph Almog, John Perry and Howard Wettstein (eds.), *Themes from Kaplan*, 43–65. New York: Oxford University Press.

Bacon, Andrew (2022). *A Philosophical Introduction to Higher-order Logics*. Routledge. Forthcoming.

Boolos, George (1975). "On Second-Order Logic." *Journal of Philosophy* 72: 509–27. Reprinted in Boolos 1998: 37–53.

— (1984). "To Be Is to Be the Value of a Variable (or to Be Some Values of Some Variables)." *Journal of Philosophy* 81: 430–49. Reprinted in Boolos 1998: 54–72.

— (1985). "Nominalist Platonism." *Philosophical Review* 94: 327–44. Reprinted in Boolos 1998: 73–87.

— (1998). *Logic, Logic, and Logic*. Cambridge, MA: Harvard University Press.

Carroll, Lewis (1895). "What The Tortoise Said To Achilles." *Mind* 4: 278–80.

Church, Alonzo (1940). "A Formulation of the Simple Theory of Types." *Journal of Symbolic Logic* 5: 56–68.

Davidson, Donald (1967). "Truth and Meaning." *Synthese* 17: 304–23. Reprinted in Davidson 1984: 17–36.

— (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Dorr, Cian (2007). "There Are No Abstract Objects." In Theodore Sider, John Hawthorne and Dean W. Zimmerman (eds.), *Contemporary Debates in Metaphysics*, 32–63. Oxford: Blackwell.

— (2010). "Of Numbers and Electrons." *Proceedings of the Aristotelian Society* 110: 133–81.

— (2014). "Quantifier Variance and the Collapse Theorems." *The Monist* 97: 503–70.

— (2016). "To Be F is to Be G." *Philosophical Perspectives* 30(1): 39–134.

Esfeld, Michael (2020). "Super-Humeanism: The Canberra Plan for Physics." In David Glick, George Darby and Anna Marmodoro (eds.), *The Foundation of Reality: Fundamentality, Space and Time*, 125–38. Oxford: Oxford University Press.

Fine, Kit (1975). "Vagueness, Truth and Logic." *Synthese* 30: 265–300. Reprinted in Keefe and Smith 1996: 119–50.

— (2001). "The Question of Realism." *Philosophers' Imprint* 1: 1–30.

Gabriel, Gottfried, Hans Hermes, Friedrich Kambartel, Christian Thiel, Albert Veraart, Brian McGuinness and Hans Kaal (1980). *Gottlob Frege: Philosophical and Mathematical Correspondence*. Oxford: Blackwell.

Grover, Dorothy (1992). *A Prosentential Theory of Truth*. Princeton: Princeton University Press.

Hilbert, David (1899). "Letter to Frege." In Gabriel et al. (1980).

Hirsch, Eli (2011). *Quantifier Variance and Realism: Essays in Metaontology*. New York: Oxford University Press.

Keefe, Rosanna and Peter Smith (eds.) (1996). *Vagueness: A Reader*. Cambridge, MA: MIT Press.

Lange, Marc (2007). "Laws and Meta-Laws of Nature: Conservation Laws and Symmetries." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38(3): 457–481.

Lewis, David (1983). "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–77. Reprinted in Lewis 1999: 8–55.

— (1984). "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221–36. Reprinted in Lewis 1999: 56–77.

— (1994). "Humean Supervenience Debugged." *Mind* 103: 473–90. Reprinted in Lewis 1999: 224–47.

— (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

Linnebo, Øystein and Agustín Rayo (2012). "Hierarchies Ontological and Ideological." *Mind* 121(482): 269–308.

McGee, Vann (1992). "Two Problems with Tarski's Theory of Consequence." *Proceedings of the Aristotelian Society* 92(n/a): 273–92.

McSweeney, Michaela Markham (2019). "Logical Realism and the Metaphysics of Logic." *Philosophy Compass* 14(1): e12563.

Paul, L. A. (2012). "Metaphysics as Modeling: The Handmaiden's Tale." *Philosophical Studies* 160: 1–29.

Prior, A. N. (1971). "Platonism and Quantification." In *Objects of Thought*, 31–47. Oxford: Oxford University Press.

Putnam, Hilary (1978). *Meaning and the Moral Sciences*. Boston: Routledge and Kegan Paul.

— (1980). "Models and Reality." *Journal of Symbolic Logic* 45: 464–82. Reprinted in Putnam 1983: 1–25.

— (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.

— (1983). *Realism and Reason: Philosophical Papers, volume 3*. Cambridge: Cambridge University Press.

Quine, W. V. (1948). "On What There Is." *Review of Metaphysics* 2: 21–38. Reprinted in Quine 1953: 1–19.

— (1953). *From a Logical Point of View*. Cambridge, MA: Harvard University Press.

— (1970). *Philosophy of Logic*. Cambridge, MA: Harvard University Press. Second edition, 1986.

Rayo, Agustín and Gabriel Uzquiano (1999). "Toward a Theory of Second-Order Consequence." *Notre Dame Journal of Formal Logic* 40: 315–25.

Rayo, Agustín and Stephen Yablo (2001). "Nominalism through De-Nominalization." *Noûs* 35: 74–92.

Shapiro, Stewart (1991). *Foundations without Foundationalism: A Case for Second-Order Logic*. Oxford: Clarendon Press.

Sider, Theodore (2011). *Writing the Book of the World*. Oxford: Clarendon Press.

— (2020*a*). "Crash Course on Higher-Order Logic." `http://tedsider.org/teaching/higher_order_20/higher_order_crash_course.pdf`.

— (2020*b*). *The Tools of Metaphysics and the Metaphysics of Science*. Oxford: Oxford University Press.

Thomasson, Amie L. (2007). *Ordinary Objects*. New York: Oxford University Press.

— (2015). *Ontology Made Easy*. New York: Oxford University Press.

Turner, Jason (2015). "What's So Bad About Second-Order Logic?" In *Quantifiers, Quantifiers, and Quantifiers. Themes in Logic, Metaphysics, and Language*, 463–87. Springer.

Väänänen, Jouko (2019). "Second-order and Higher-order Logic." Stanford Encyclopedia of Philosophy. Available at `http://plato.stanford.edu/entries/logic-higher-order/`.

Weston, Thomas (1976). "Kreisel, the Continuum Hypothesis and Second Order Set Theory." *Journal of Philosophical Logic* 5(2): 281–98.

Whitehead, Alfred North and Bertrand Russell (1910). *Principia Mathematica*. Second edition. Cambridge: Cambridge University Press, 1925.

Williams, J. Robert G. (2007). "Eligibility and Inscrutability." *Philosophical Review* 116: 361–99.

Williamson, Timothy (2003). "Everything." *Philosophical Perspectives* 17: 415–65.

— (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.