What's So Bad about Second-Order Logic?*

Jason Turner

April 16, 2018

According to its detractors, second-order logic is 'not logic'. Philosophical orthodoxy seems to side with the detractors, even if it's not quite clear what their complaint amounts to. By and large, contemporary philosophers tend to regard second-order logic with suspicion, or worse.

Among *bona fide* philosophers of mathematics and logic the debate is considerably refined, and often hinges on points of especial interest to those disciplines. But second-order logic's bad name among the general philosophical populace isn't thanks to these high-level debates. No, most philosophers think second-order logic is bad thanks to some stuff Quine said about ontological commitment and its being set theory in sheep's clothing.

I am come neither to praise second-order logic nor to bury it. Rather, I want to get to the bottom of what these complaints might be, and whether they are any good. The goal isn't Quine exegesis — I'm not particularly concerned with whether Quine's overall philosophy of logic gave him good reason to reject second-order logic. The question instead is whether either of these Quine-inspired themes gives us conclusive reason to reject second-order quantification. I claim neither neutrality nor completeness; my biases on both fronts will be in plain view. Still, my hope is not to persuade you of second-order logic's virtues, but just to give you a good feel for the relevant issues and considerations.

1 Logic

Second-order logic's detractors claim that it is 'not logic'. But what does it mean to call something 'logic'? It will be helpful to consider both a philosophical and a technical answer.

1.1 Formal Systems

Logicians study *logical systems*. At their most basic, each such system consists of three things: (i) a language; (ii) a *syntactic calculus* for the language; and (iii) a *model theory* for the language.

A syntactic calculus is, at its heart, a procedure for granting some sort of 'good' status to arguments based solely on the syntactic shape of their premises and conclusions. For a given system S, when S's syntactic calculus regards an argument

*Thanks to Aaron Cotnoir, Daniel Elstein, Robbie Williams, and an audience at the University of Leeds Centre for Metaphysics and Mind for helpful comments and discussion.

from a set of premises Δ to a conclusion ϕ as good we write '$\Delta \vdash_S \phi$'. The $\vdash_S$ relation is S's *proof-theoretic consequence* relation.

A model theory specifies a class of objects — usually, set-theoretic ones — called *models*, and defines a 'true in' relation that holds between sentences of the language and these models. If a sentence ϕ is true in a model M, we say that M is a *model of* ϕ, and if all the sentences in a set Δ are true on M, then it is a model of Δ. When every model of Δ in a given system S is also a model of ϕ in that system, we write $\Delta \vDash_S \phi$. The $\vDash_S$ relation is S's *model-theoretic consequence* relation.[1]

If a given system is *sound*, then proof-theoretic consequence in it guarantees model-theoretic consequence in it. If it is *complete*, then model-theoretic consequence guarantees the proof-theoretic kind. If we're lucky, we will have a sound and complete system. Occasionally we're not so lucky. Some systems are unsound or, more commonly, incomplete.

1.2 Genuine Logic

Formal systems, and the consequence relations they give rise to, are cheap. I can cook up a formal system where ϕ is both a proof-theoretic and a model-theoretic consequence of Δ if and only if ϕ has at least one instance of the name 'Jason' in it. But such a system is manifestly uninteresting. We want to study more interesting systems than that.

We might care about a formal system for one of two reasons. First, it might be *mathematically* interesting — it might have mathematical properties worth studying in its own right. But second, it might be *philosophically* interesting, because it may give us a mathematically tractable way to study something we care about: *genuine logical consequence*.

Here's the idea. Between premises and a conclusion there can hold a particular relation (or, perhaps, one of a handful of relation*s* — see [2]) that we care about. We aim to give arguments where this relation (or one of these relations) holds between the premises and conclusions of arguments. Arguments where this happens are good, and command our attention, in ways that arguments without it aren't and don't.[2] Call this relation *logical consequence*.

Perhaps (for a given language) this relation is co-extensional with the proof-

[1] This can be generalized in various ways. For instance, in many-valued logics, models assign sentences one of several truth-values, one or more of those values is 'designated', and model-theoretic consequence is understood as preservation-of-desigation-in-all-models. Other generalizations are possible, of course. For our purposes, though, we can stick with our first-pass understanding.

[2] Gillian Russell [20] argues that there is a crucial ambiguity in what we take 'arguments' to consist in which leads to a logical pluralism of a different sort than that advocated in [2]. I have slid over this ambiguity, and am going to proceed by assuming (hoping!) it won't affect any of what's to come.

or model-theoretic consequence relation of some formal system. Perhaps (for a given language) this relation is even *identical* to some proof- or model-theoretic consequence relation. Still, in the first instance we care about the relation not under its proof- or model-theoretic guise, but under a guise that we already implicitly grasp, however darkly, in our native ability to tell what 'follows from' what.

Let's use '$\Rightarrow$' for this genuine consequence relation. Its importance stems in large part from the theoretical role it is supposed to play. That role is complex, but is often thought to include:

Modality: If $\Delta \Rightarrow \phi$, then it is in some sense impossible for all of Δ to be true while ϕ is false.[3]

Normativity: If $\Delta \Rightarrow \phi$, then it is an error in reasoning to accept all of Δ while rejecting ϕ.

Topic Neutrality: Whether $\Delta \Rightarrow \phi$ should not depend on the truths of any particular subject-matter.

Ontological Innocence: Whether $\Delta \Rightarrow \phi$ should not depend on the existence of things of any particular kind K.

(Cf. also [2: §§2.3–2.5]. Note that Ontological Innocence is plausibly taken as a corollary of Topic Neutrality, since whether there are things of a particular kind K looks like a truth of a particular subject-matter.)

To say that second-order logic *is* logic is to say that there is at least one second-order system *SO* where, for any Δ and ϕ in that system, $\Delta \Rightarrow \phi$ iff $\Delta \vDash_{SO} \phi$. Conversely, to say that second-order logic is *not* logic is to deny that any second-order system coincides with genuine consequence this way.

The arguments to be considered each aim to show that every second-order system violates one of the four principles just outlined. If they're right, and if those principles are indeed constraints on genuine consequence, then no second-order consequence relation coincides with genuine consequence, and so second-order logic is not logic.

You may doubt that there *is* any relation of 'genuine consequence'. Perhaps our native grasp of 'following from' is too dark and muddled to have settled on any single, unambiguous relation between premises and conclusion. If it hasn't, then the question 'Is second-order logic logic?' may have no well-defined answer. But this won't rob the arguments to be considered of their force, for (if successful) they

[3]These principles make free use of truth and falsity, and thus may be subject to worries of the sort Field [6] has levelled against what he calls 'The Validity Argument'. I think the situation here can be finessed here with use of conditionals and infinite conjunctions, but for our purposes won't bother.

can be taken to show that second-order logic fails to have certain nice features we commonly associate with 'logic'.

2 Second Order Systems

As I've interpreted it, the claim that second-order logic is 'not logic' depends on certain logical systems counting as 'second-order'. I should say something about what this means.

2.1 Second-Order Languages

Second-orderness is, in the first instance, a property of *languages* — the property a language has when it allows variables to occur in predicate position and quantifiers to bind those variables.[4] We get a standard second-order language by taking a first-order language, adding a new stock of predicate-variables, and allowing those variables to be bound by the quantifiers. We'll call a system second-order just in case its language is second-order.

Generally, the second-order predicate variables have a fixed adicy: there will be one-placed predicate variables 'X^1', 'Y^1', ..., two-placed predicate variables 'X^2', 'Y^2',..., and so on. If a first-order formula is well-formed, then the result of replacing any of its n-adic predicates with an n-adic predicate variable will be well-formed, as will the result of binding those variables with a universal or existential quantifier.

2.2 Second-Order Model Theories

Second-order systems come in two varieties, depending on what their model-theory is like. Very roughly, the two model theories correspond to different conceptions of what the second-order quantifiers are doing.

Suppose we think of second-order quantifiers as ranging over properties and relations. On an *abundant* conception of properties, every set of things corresponds to some property, and every set of n-tuples corresponds to some n-adic relation.[5] On a *sparse* conception, properties and relations are scarce: some sets of individuals (or n-tuples of individuals) may correspond to no genuine property (or relation).

[4]Systems that allow the variables but not the binding are also possible — see e.g. [21: p. 62] — but we won't consider them here.

[5]Since properties are generally taken to be intensional entities, a better characterization would use sets of possibilia. We will stick to extensional contexts here, though, so we can safely ignore the difference.

Full second-order systems hardwire the abundance of properties and relations into the model theory. *Henkin* second-order systems, by contrast, allow the properties and relations to be sparse.

In standard first-order model theory, models consist of a domain — a non-empty set — and an interpretation function, which assigns individuals from that domain to names of the associated language and extensions over that domain to its predicates. (An extension for a monadic predicate is simply the set of its satisfiers; the extension for an n-adic one is the set of n-tuples that satisfy it.) A variable assignment assigns individuals from the domain to variables; truth on a model is first defined for open formula relative to a variable assignment. Truth on the model *simpliciter* is simply truth on all variable assignments.

To get a model theory for *full* second-order logic, we use the same sorts of models, but modify the truth-on relation. First, we extend variable assignments so that they also assign n-adic extensions to n-adic predicate variables. We then add a clause to our definition of truth-on-a-model so that $\ulcorner \forall X \phi(X) \urcorner$ is true exactly when $\ulcorner \phi(X) \urcorner$ is true on all variable assignments. This essentially gives the second-order n-adic variables the effect of quantifying over all n-adic extensions drawn from the domain.

Henkin model theory differs from full model theory by including a 'second-order domain'. We can give it one in several ways, but however we do it, the effect will be to specify a range of extensions (drawn from the first-order domain) for the second-order variables to range over. We then modify the definition of 'variable assignment' so that the second-order variables can only be assigned extensions from the second-order domain, and keep the definition of truth-on-a-model the same as it is for full second-order logic. This essentially gives the second-order n-adic variables the effect of quantifying over a privileged class of n-adic extensions drawn from the domain.

This small difference has profound knock-on effects. In particular, it gives full second-order logic incredible expressive power. One example, which we will return to later, is that it allows us to form sentences (CH) and (NCH) such that, if the continuum hypothesis is true, (CH) is true on all full second-order models, and if the continuum hypothesis is false, (NCH) is true on all full second-order models.[6] As a result, full second-order systems are essentially incomplete. There are no syntactic calculi that capture all of their model-theoretic consequences. More precisely, for every full second-order system F, there is some sentence ϕ where $\vDash_F \phi$ but $\nvdash_F \phi$.[7]

Conversely, there are many complete Henkin second-order systems. In fact, this

[6] [21: p. 105] Note that (NCH) is not simply the negation of (CH). On models with countable domains, both (CH) and (NCH) are true.

[7] This is failure of what is called 'weak completeness' — not all model-theoretic truths are theorems — and thus is stronger than the failure of what is called 'strong completeness,' which happens when not all model-theoretic consequences are proof-theoretic ones.

kind of system got its name when Leon Henkin [8] proved that these systems are expressively equivalent to sorted first-order systems.

The arguments to be given in §3 target second-order logic *generally*, and thus aim to show that *no* second-order system is logic. The arguments in §4 are focused primarily on showing that full second-order logic is not logic, and leave the Henkin version untouched. It's worth keeping the complaints, and their targets, distinct.

3 Ontological Guilt

The first group of arguments complain that second-order logic isn't ontologically innocent, and therefore that it isn't logic. (These arguments really target second-order *quantification*, rather than second-order *logic*: the idea is that, whether logic or not, there's something untoward about quantifying in predicate position.)

3.1 Textbook Quineanism

If asked what was wrong with second-order logic, the average philosopher-on-the-street is likely to offer an argument like the following, and attribute it to Quine:

Textbook Quineanism

(i) You are ontologically committed to something of a particular kind if and only if the bound variables in your system have to range over things of that kind to be true.

(ii) The bound variables of second-order logic have to range over something predicate-like.

(iii) Therefore, theorems of second-order logic ontologically commit you to something predicate-like.

(iv) So second-order logic is not logic.

('Predicate-like' things here include sets, properties, Fregean concepts, and anything else that could in principle be the semantic value of a predicate.)

Note first that, whatever its merits, this argument cannot be Quine's. The first premise supposes we can take an arbitrary language, check its quantifiers, and figure out what its ontological commitments are. But Quine consistently held that 'ontological commitments' were only well-defined for *first-order* theories; to get the ontological commitments of some other sort of theory, we first had to translate it into a first-order one. Quine also famously held that, in general, there was no uniquely 'right' translation for any given theory. As a result, non-first-order theories don't give rise to unique ontological commitments.

The argument can be made more Quinean by replacing the first premise with one that says we have to translate into a first-order system and replacing the second premise with the claim that we'll thereby translate second-order quantifiers as first-order ones ranging over something predicate-like. But it's not clear this argument would be any better, because it's not clear Quine's insistence on first-order translation is well-motivated. That insistence is related to his belief that serious theorizing shouldn't use anything but extensional first-order resources, and would require us to translate not just second-order quantifiers, but modal and tense operators, generalized quantifiers, and a host of other resources before we can get down to serious theory. Most philosophers nowadays reject the Quinean demand when it comes to these other resources. If they want to insist on it for second-order quantification, they owe us an explanation for the differential treatment.[8]

But even if the textbook argument isn't Quine's, it *is* an argument, and fuels contemporary suspicion of second-order quantification. Can the friend of second-order logic say anything about it?[9]

We can distinguish two very different pictures of second-order quantification. Which forms of resistance are available will depend on which picture we are attracted to. According to one picture going back to Frege, second-order quantifiers range over a special, distinctive type of object. Let's call this the *Fregean* picture. More recently, a number of philosophers have started to endorse a *qualitative* conception, according to which second-order quantifiers shouldn't be thought of 'ranging over' anything at all. To think of the quantifiers this way is to simply misunderstand the kind of business they're in.

Friends of the Fregean picture will answer the Quinean argument by denying the ontological innocence principle driving the move from (iii) to (iv). The Fregean thinks that there is a distinctive kind of *logical* entity, and that the second-order quantifiers range over those. But according to the Fregean it's no mark against a system's logicality — or even its topic neutrality — if it commits us to things of a distinctively logical kind. So the arguments of this section won't trouble the Fregean.[10]

Friends of the qualitative conception will resist premise (ii). Considering *how* they might do this will give us a better picture of what the qualitative conception

[8]The argument of §3.2 may provide such an explanation; but then we can consider it in its own right, rather than as an adjunct to Quine's.

[9]Another question: If this argument isn't Quine's, what *is* Quine's, and is it any good? Unfortunately, the closest thing I can find to an argument in Quine is at p. 66 of his *Philosophy of Logic*, and it's a howler. I've left it out in interests of space, but had I included it I would not have said anything more (or better) against it than was said by Boolos [3: pp. 510–511].

[10]The Fregean has other troubles, though: expressibility problems that relate to the so-called 'concept horse problem'. I cannot hope to pursue that huge literature here. I'll simply focus on the qualitative conception instead.

amounts to.

3.1.1 The Nature of Quantification

Let's start by considering an argument *for* premise (ii). It runs like this: '*To use a quantifier* is to ontologically commit. The very idea of $\ulcorner\exists X^n\phi(X^n)\urcorner$ being *not* ontologically committing simply misunderstands the nature of ontological commitment. So $\ulcorner\exists X^n\phi(X^n)\urcorner$ must ontologically commit us to something that corresponds to the bound 'X^n'. Furthermore, since the bound 'X^n' is the kind of thing that can be used predicatively, it will have to be of a distinctive 'predicate-like' kind — a set, or a Fregean concept, or what-have-you. And that's all it means to say that 'X^n', ranges over something of that kind.'

According to the qualitative conception of second-order quantification, this line of thought is mistaken. Consider the sentence

(1) Fido is brown.

What does it ontologically commit us to? Intuitively, it commits us to Fido, and to brown things, although not to any predicate-like thing of 'brownness'. Following Prior [12: p. 35], we can ask why it commits us to Fido. The answer seems to be because it uses the name 'Fido,' and it uses that name in such a way that the sentence can only be true if 'Fido' refers to something. So consideration of (1) suggests that terms in name position have the power to commit, and terms in the predicate position do not.

If that is right, then the reason

(2) $\exists x(x \text{ is brown})$

commits us to something brown is not because a variable is bound, but because that variable is in name position. On the other hand, since 'is brown' doesn't commit us to anything predicate-like in (1), the bound second-order variable in

(3) $\exists X^1(X^1(\text{Fido}))$

shouldn't, either. To say otherwise is to say that it is somehow the distinction between the *particular* and the *general* that engenders ontological commitment. But why should this be? Why is it instead not the distinction between the *subject* and the *qualitative way* the subject is said to be that engenders the commitment? It is the presence of expressions (names or variables) in name-like position that gives rise to ontological commitment — bound variables have nothing to do with it. When we assert (3), we're not saying that there is some predicate-like entity or other that Fido participates in, any more than when we assert (1) we say that there is a particular predicate-like entity that Fido participates in. In (1) we say how Fido is particularly; in (3) we say how he is in more generality.

3.1.2 Model Theory

Another argument for premise (ii) appeals to second-order model theory, as described in §2.2. 'Just look at that model theory. It explicitly has second-order quantifiers ranging over extensions drawn from the domain. They thus range over something predicate-like — extensions — and premise (ii) follows.'

The friend of the qualitative conception has several potential responses. First, he might sharply distinguish between a *model theory* and a *semantic theory*. The former is a set-theoretic device for mapping consequence relations. The latter gives *truth-conditions* for claims of a language. The two are, at first glance, simply different projects. Since model theory uses a notion of 'truth in a model,' it is tempting to run the two projects together. But we needn't. In the first instance, model theory provides us with a class of models and a relation between models and sentences. We could have called this relation 'zapping,' and then said that ϕ is a model-theoretic consequence of Δ whenever any model that zaps all of Δ also zaps ϕ. We may think that the thus-defined relation helps us investigate genuine logical consequence without thinking that zapping is some species of *truth*, or that zapping-on-a-model conditions gives us any particular insight into meaning or truth-conditions. If truth-on-a-model conditions have little to do with truth-conditions, then the fact that the former treat quantifiers as ranging over extensions gives us little reason to think that the latter do.

Not everyone can endorse this response. In particular, standard Tarskian accounts of logical consequence identify it with the preservation of truth on all *interpretations* — all ways of assigning meanings to the non-logical terms. They then identify interpretations with models. Among the interpretations — which are the models — is the *intended* interpretation, the one that gives genuine truth-conditions for sentences of the language. On such a picture, models *are* in the business of giving genuine truth-conditions, and — if the second-order model theory is as described in §2.2 — those truth-conditions say that $\ulcorner \exists X^n \phi(X^n) \urcorner$ is true if and only if some set of n-tuples from the domain satisfies $\ulcorner \phi(X^n) \urcorner$.

But Tarskians friendly to second-order quantification may offer one of two other responses. First, they might say we gave the wrong model theory in §2.2. That model theory was given in a first-order metalanguage. But if second-order languages are kosher — as the second-order Tarskian believes they are — then their model theories should be second-order as well. Down this route lies the technical work of providing such a model theory and verifying that it works the way we want it to, a task undertaken in [17]. If successful, though, it defuses the argument, for the second-order model theory does not interpret the second-order quantifiers as first-order ones ranging over subsets of the domain.

A second response accepts the model theory of §2.2 but denies that it makes for unwanted ontological results. This response starts by noting that, on standard

Tarskian accounts, '$F(a)$' is true if and only if the denotation of 'a' is in the extension of 'F'. Still, we don't think that simply asserting '$F(a)$' commits us to the existence of extensions. (If we did, we could argue that first-order logic isn't ontologically innocent and therefore not logic.) Although extensions show up in the official truth-conditions for '$F(a)$', they are mere theoretical apparatus of the semantics, not part of what the sentence demands *of the world*. (Cf. [16] and [26].) The sentence only demands that a be F, even if sets show up in the machinery we use for semantic theorizing. And it is what a sentence demands of the world, rather than semantic machinery, that we must take seriously when figuring out a sentence's ontological commitments.

The quantification over sets given by the model theory is, according to this response, simply more semantic machinery. If that's right, the model-theoretic treatment of the second-order quantifiers, by itself, gives us no reason to think those commit us ontologically to predicate-like entities.

3.2 The Intelligibility Argument

A strengthened version of the Textbook argument — one which may underlie Quine's thinking, as the seeds for it can be found in his work — stems from the thought that the ontologically innocent picture of second-order quantification just sketched is somehow unintelligible.

Here, in rough form, is the idea. Formal systems are nothing but squiggles on paper until we do something to give them meaning. But the only way we can give these expressions meaning is by explicitly defining them using expressions we already understand.[11] And there's no way to do this for second-order quantifiers except by identifying them with first order quantifiers over predicate-like entities. So either the second-order quantifiers are meaningless, or else they are ontologically committal in a way that bars them from being logic.

Consider how this works for first-order quantifiers. Quine ([13: pp. 65–71] and [14: pp. 161–163]) and van Inwagen [24: pp. 18–22] tell us that sentences of the form

(4) $\exists x(\ldots x \ldots x \ldots)$

should be understood as

(5) There is something x such that $\ldots x \ldots x \ldots$.

This then tells us how to understand '$\exists$'. Furthermore, the variables themselves should be understood as pronouns, with indices to disambiguate their anaphors. That is, (5) is understood as

[11] Or, perhaps, by ostending their meanings; but I take this option to be unavailable for second-order quantifiers.

(6) There is something such that ... it ... it ....

If there were extra quantifiers and pronouns, we would get confused which pronouns went with which quantifiers. We can index them with little 'x' and 'y' subscripts to avoid disambiguation; and we can write the thus-indexed pronouns simply as 'x' and 'y' rather than 'it_x' and 'it_y'. That is how to understand variables.

According to the argument, if second-order quantifiers are to be intelligible, they will need a similar story about

(7) $\exists X^n(\ldots X^n \ldots X^n \ldots)$.

The most natural way to read (7) is as

(8) There is some way X^n that ... X^n ... X^n ...,

which is simply first-order quantification over 'ways' — predicate-like things that might be identified with properties, relations, extensions, or what-have you. Thus, if second-order quantification is intelligible, it's just first-order quantification over 'ways', making it not ontologically innocent and thus not logic.[12]

The above argument essentially relies on three premises ((i)–(iii)) plus the Ontological Innocence principle, and runs:

The Intelligibility Argument

(i) A formal system is meaningless unless it is provided an interpretation — a specification of what its expressions mean.

(ii) We provide an interpretation by specifying the meaning of each expression using terms already understood.

(iii) The only plausible meaning for '$\exists X^n \ldots$' is 'There is a predicate-like thing X^n such that ...'.

(iv) So, if second-order logic is meaningful, the second-order quantifiers say that there are predicate-like things.

(v) If the second-order quantifiers say that there are predicate-like things, then second-order logic is not logic.

(vi) Therefore, either second-order logic is meaningless, or it is not logic.

I doubt the first premise can be rejected — uninterpreted squiggles on paper are just that. There is room to resist the other two, though. We'll consider them in reverse order.

[12] I don't know anywhere this argument is explicitly presented in this form; van Inwagen presents a similar argument against substitutional quantification in [23], and comes close to giving this one [25: p. 124]. In the latter he also ascribes something like the present argument to Quine in *Philosophy of Logic*, but I cannot quite find that argument there.

3.2.1 Premise (iii)

Premise (iii) may be resisted by finding a better interpretation for the second-order quantifiers than that suggested by the argument. Susan Haack [7: 52–55] has suggested that we can get away from these sorts of ontological-commitment worries by interpreting the second-order quantifiers *substitutionally*.

Truth-on-a-model conditions for quantified sentences are generally given 'objectually,' saying (roughly) that $\ulcorner \exists x \phi(x) \urcorner$ is true if and only if something in the model satisfies ϕ. But substitutional truth-on-a-model conditions say instead that it is true if and only if, for some name α, $\ulcorner \phi(\alpha) \urcorner$ is true. Notice this gets away from talking about things satisfying predicates, trading only in truth. If the truth-on-the-model conditions correspond to genuine truth-conditions, we can have a sort of 'ontologically innocent' reading of the quantifier. For instance, if

(9) Zeus is a greek god

can be true without there being any Zeus, then if '$\exists$' is read substitutionally,

(10) $\exists x(x \text{ is a greek god})$

can, too [1].

It is controversial whether (9) commits us to Zeus or not. It's much less controversial that (1) does *not* commit us to any predicate-like entity of brownness. In this case, if we interpret the second-order quantifiers substitutionally — that is, if we insist that $\ulcorner \exists X^n \phi(X^n) \urcorner$ is true if and only if, for some n-placed predicate Π^n, $\ulcorner \phi(\Pi^n) \urcorner$ is true — then we seem to have an interpretation of second-order quantification that shouldn't be understood as saying there are predicate-like things.

Well, if there really is such a thing as a 'substitutional interpretation' of a quantifier, at least. Van Inwagen [23] argues there is no such thing. My so-called 'substitutional interpretation' of the quantifier was in fact a substitutional *truth-condition*. An interpretation was supposed to be specification of *meaning*, and truth-conditions aren't obviously the same.

We could try to identify the meaning with the truth-conditions, in which case '$\exists X^n \ldots$' would mean something like 'there is an n-placed predicate which . . . ,' and thus would ontologically commit us to predicate-like things — namely, the predicates themselves. As van Inwagen notes, friends of substitutional quantification deny they mean *this* when they quantify substitutionally. But unless they tell us what they do mean, they haven't given us an interpretation at all.

Agustín Rayo and Steven Yablo [18] have suggested a different, non-substitutional strategy. They start by observing that most ordinary English renderings of second-order quantification look committal because they implicitly put second-order resources in nominal position. According to them, '$\exists X^n \ldots$' shouldn't be interpreted as 'there is some way . . . ,' because 'some way' is a noun phrase, and foes of second-order logic will see ontological commitment wherever noun phrases are used. If an interpretation wants to be non-committal, it should de-nominalize second-order quantifiers and variables.

Consider, for instance, the second-order sentence

(11) $\exists X^2(X^2(\text{Scooby, Shaggy}))$.

If we interpret this as 'There is a way Scooby and Shaggy are related,' we treat the second-order variables as ranging over 'ways'. But Rayo and Yablo suggest we should interpret (11) as

(12) Scooby and Shaggy are somehow related.

Here, the quantifier shows up adverbially, and there's no temptation to read it as quantifying over ways — or anything else.

Rayo and Yablo use this thought to interpret second-order quantifiers in a pattern that mirror's Quine and van Inwagen's interpretation of first-order quantifiers. Second-order variables are treated as the pro-adverb 'are so related', where we use different variables for anaphoric disambiguation. A sentences of the form

(13) $\exists X^n(\ldots X^n \ldots X^n \ldots)$

is understood as

(14) Things somehow$_{X^n}$ relate such that ... are so$_{X^n}$ related ... are so$_{X^n}$ related ...

The subscripted 'X^n's just help us keep track of which 'so related's go with which 'related somehow's.[13]

Yablo and Rayo's reply seems to undercut the intelligibility argument. If we can answer van Inwagen's challenge about substitutional 'interpretations', Haack's reply will undercut it as well. But it's worth noticing that neither reply, even if it vindicates second-order quantification, clearly vindicates *full second-order logic*.

When it comes to Haack's proposal, the reason is technical. The model theory for full second-order logic, recall, has the second-order variables ranging objectually over *all* extensions that can be drawn from the domain.[14] If we trade that in for substitutional quantification, we get the effect of, at best, ranging over only what are known as the 'constructible' extensions from the domain — those extensions that contain all and only the satisfiers of some open formula in the language. And it's well-known that ranging over only the constructible extensions is far weaker

[13]'Things somehow$_{X^n}$ relate is to be interpreted (roughly) as 'things are-or-aren't somehow$_{X^n}$ related; see [18: pp. 84].

[14]I'm assuming here that the second-orderist is happy to reason instrumentally with models in this way, even if she insists that, in all seriousness, second-order variables aren't in the 'ranging over' business.

than ranging over all of them. For instance, neither of (CH) and (NCH) will be true in every model. [21: pp. 110–116]

When it comes to Yablo and Rayo's proposal, the situation is less clear. But consider

(15) $\exists X^3 \forall x \forall y \forall z (X^3(x, y, z) \leftrightarrow x = \text{Scooby} \wedge y = \text{Shaggy} \wedge z = \text{Velma})$

This will be true on all full second-order models (which assign referents to 'Scooby', 'Shaggy', and 'Velma'), because for any three things in the domain one of the potential values for 'X^3' will be the set containing just one triple with those three things. But Rayo and Yablo interpret this as

(16) Things somehow$_{X^2}$ relate such that any three things that are so$_{X^3}$ related if and only if the first is Scooby, the second Shaggy, and the third Velma.

In other words — reifying 'hows' for clarification — there is a relation had by and only by Scooby, Shaggy, and Velma, in that order. This is far from clear, though; given our grasp on adverbial quantification, it may be that however Scooby, Shaggy, and Velma are related, there will be three other things that are also so-related.[15] If so, then while Yablo and Rayo's interpretation gets us second-order quantification, it won't get us full second-order logic.[16]

3.2.2 Premise (ii)

The intelligibility argument offers a challenge: provide an ontologically innocent natural-language interpretation of second-order quantification. Objections to premise (iii) take up this challenge. Objections to premise (ii), in contrast, reject the challenge itself.

We should agree that, if uninterpreted, second-order resources don't mean anything. But why think that meaning can *only* be assigned by explicit definition in already-understood terms? That's a pretty restrictive demand, and would seem to rule out theoretical terms like 'superposition' from quantum physics or '$\in$' from set theory, as neither has any explicit definition in more familiar terms. We'd better not commit to a constraint on interpretation that makes these meaningless.

Lewis [11] suggests that terms such as these are defined by *theoretical role*. We specify our theory (quantum mechanics, say, or set theory), and in doing so we use some new expressions, such as 'superposition' or '$\in$'. The theory links these new expressions with ones we're already familiar with. For instance, set theory has the theorem

[15] Likewise, it may be that however Scooby, Shaggy, and Velma aren't related, there will be three other things also so unrelated; this takes care of the 'don't' part of the 'do-or-don't' clause mentioned in note 13.

[16] See [19] for related worries.

(17) $\forall x \exists y(x \in y \land \forall z(z \in y \leftrightarrow z = x))$,

which uses not just the new '$\in$' but the old '$\exists$', '$=$', and so on.

Call a theory *realizable* if and only if its new expression can be interpreted in a way that — leaving the interpretations of the old expressions alone — makes the theory, or at least most of the theory, true. We can then give the new expressions meaning indirectly, by saying 'Let the new expressions be interpreted however they need to be in order for the theory to be true'. If the theory is not realizable, our new terms fail to get a meaning. If there is a uniquely best interpretation of the new terms that makes the theory true, the new terms will be thus interpreted. If several equally good interpretations each would make the theory true, the new terms will be indeterminate in interpretation among those several meanings. In any case, if the theory is realizable, we can interpret its new terms without explicit definition.

If we can do this for predicates we can presumably do it for expressions of other syntactic categories. In particular, we ought to be able to do it for quantifiers. We can write down the 'theory' of second-order logic, say 'Let the second-order resources be interpreted however they must be in order to make the theory true,' and let the interpretative chips fall where they may.

What would count as the 'theory' in this case? If we just want to quantify second-order, and are happy for Henkin consequence to govern our quantifying, the 'theory' can be the axioms of some Henkin system. But since full second-order logic is incomplete, no set of axioms will fix anything even close to a unique interpretation for second order quantifiers governed by the full consequence relation.

Can the friend of full second-order logic do better? Perhaps. We can specify the full consequence relation model-theoretically. Call an inference from Δ to ϕ *approved* if and only if ϕ is a full model-theoretic consequence of Δ. Then we can say 'Let the second-oder resources be interpreted however they must be in order to make all and only improved inferences valid.'

This response defuses the Intelligibility Argument. But it opens the door to skepticism. We *might* worry that the 'theory' (whether Henkin or second-order) isn't realizable at all. But it's plausible that any language which can be given a truth-conditional semantics can also be made meaningful, so we perhaps shouldn't worry to much about realizability. More worrisome is that, even if this gets us *an* interpretation for second-order quantification, it might get us an ontologically guilty one. For all that's been said, the only interpretation of '$\exists X^n \ldots$' that realizes the theory is 'there is a set-like entity X^n such that ...', and we're back in the arms of premise (iv).

Of course, skeptical worries aren't arguments, and the friend of second-order quantification may simply take the ontological innocence of the interpretation on faith. At best we end in a dialectical stalemate: the friend of second-order logic certain she has the innocent interpretation she needs, and the foe certain she does not.

4 Set Theory in Sheep's Clothing

The arguments from ontological guilt aim to undercut the very idea of second-order quantification. The arguments in this section, by contrast, have no special beef with putting variables in predicate position. They aim instead to show that there is something objectionable about *full* second-order logic.

I mentioned in §2.2 that, if F is a system using the 'full' model theory, then there are sentences (CH) and (NCH) (using only logical vocabulary) such that, if the continuum hypothesis is true, $\vDash_F$ (CH), and if it is false, $\vDash_F$ (NCH). In other words, it treats either one or the other as a logical truth.

At its most basic, the 'Sheep's Clothing' worry is that no genuine consequence relation should do this. It comes in two forms. In one form, consequence relations shouldn't do this because it keeps them from Topic Neutrality. In the other, they shouldn't do this because it gets them into trouble with Normativity.

4.1 Topic Neutrality

Logical consequence is a relation between premises and conclusions; logical truth — or what I'll call *validity* — is a property of individual sentences. It's the property that a sentence has if and only if it's a logical consequence of any set of premises (including the empty set).

Every consequence relation has a corresponding validity property: ϕ is an *S-validity* if and only if it is an S-consequence of any premises whatsoever. A sentence ϕ is a *full second-order validity* iff $\Delta \vDash_F \phi$ for every set Δ. It is a *genuine validity* iff $\Delta \Rightarrow \phi$ for every Δ.

Topic Neutrality gave us a constraint on consequence relations, but it gives rise to a further constraint on validity: whether or not a given sentence is genuinely valid should not depend on the truths of any particular subject-matter. If full second-order logic is logic, then the full second-order validities are genuine validities. So if second-order logic is logic, whether a sentence is full-second-order valid shouldn't depend on the truths of any particular subject matter. But (goes the objection), it *does* — it depends on the truths of set theory, as the example of (CH) shows — so second-order logic isn't logic.

Put more precisely, this argument runs:

Topical Sheep's Clothing

(i) (CH) is a full second-order validity if and only if the continuum hypothesis is true.

(ii) If (i), then if second-order logic is logic, the logical validities depend on the truths of set theory.

(iii) The logical validities do not depend on the truths of set theory.

(iv) So full second-order logic is not logic.

Premise (i) is supported by a simple model-theoretic fact, and premise (iii) follows from the Topic Neutrality constraint. But what of premise (ii)?

Note first that, for premise (ii) to have any plausibility at all, the 'if and only if' in (i) needs to be stronger than a mere truth-functional one. It's not the mere observation that '(CH) is a full second-order validity' and 'the continuum hypothesis is true' have the same truth-value that underwrites (ii); rather, (i) is supposed to express some sort of fairly deep connection between the two.

The deep connection seems to be there: we can *prove* that, if (CH) is a second-order validity, the continuum hypothesis is true, and vice versa.[17] But that might not be enough to make second-order logic 'depend' on the truths of set theory in any objectionable way.

Why not? Well, there are similar connections between first-order logic and set theory, but this doesn't seem to make first-order logic unacceptably topic-sensitive. Consider, for instance, the first-order inference:

(18)
$$\begin{array}{l} \forall x(Fx \to Gx) \\ \underline{\forall x(Gx \to Hx)} \\ \therefore \forall x(Fx \to Hx) \end{array}$$

We can show that this inference is first-order valid if and only if the subset relation is transitive. This is because (on the standard model-theoretic semantics) $\ulcorner \forall x(\Pi x \to \Xi x) \urcorner$ is true if and only if the set of Π-satisfiers is a subset of the set of Ξ-satisfiers.[18] This connection does not seem to keep first-order logic from being logic, though, so it's not clear why the similar connection between (CH) and the continuum hypothesis is a problem for second-order logic.

Fans of the Topic Neutrality argument may object that the relation between (18) and transitivity is crucially different from that of (CH) and the continuum hypothesis. Although the biconditionals connecting consequence (or validity) to set theory are true in both cases, in the first case we rely on the validity of (18) to show that the subset relation is transitive, whereas in the second we we rely on the (supposed) truth of the continuum hypothesis to show that (CH) is a full second-order validity. This suggests (goes the objection) that the dependencies are different

[17] Slightly more carefully, we can prove that, if $\vDash_F$ (CH), the continuum hypothesis is true, and vice versa; someone who follows the line of thought outlined in §4.2.3 will have room to resist concluding that if $\vDash_F (CH)$, (CH) is genuinely valid.

[18] Thanks to Aaron Cotnoir for suggesting this example to me.

in the two cases: the transitivity of subsethood depends on the validity of (18), whereas the validity of (CH) depends on the truth of the continuum hypothesis.

The observation about how we establish each biconditional seems right, as a mater of sociological fact. That's because we tend to use broadly first-order reasoning when reasoning model-theoretically about second-order logic.[19] It's not clear this sociological fact carries any weight, though. If we were happy to reason second-order, we might first formulate second-order set theory and then use the (assumed) validity of (CH) to show that the continuum hypothesis is true, and so on. Whether we did this or not, we might think that the continuum hypothesis (or its negation) is the set-theoretic *result* of the (logically prior) truth of (CH) (or (NCH)). No one who takes this perspective should grant that (i) makes validities depend in any worrisome sense on set theory.

4.2 *Normativity*

An alternative argument sees troubles stemming from the interaction between full second-order logic's expressive power and the normativity of logic. According to the normativity constraint of §1.2, it's an error of reasoning to accept the premises of a genuinely valid argument while rejecting its conclusion. As a corollary, it's an error of reasoning to reject any genuine validities. If second-order logic is indeed logic, then either (CH) or (NCH) will be a genuine validity, and we will thus have logically-based epistemic obligations towards it. But the continuum hypothesis — and so, by extension, (CH) and (NCH) — seem radically epistemically unsettled in a way that doesn't sit well with these obligations.[20]

In argument form, these observations run:

Normative Sheep's Clothing

(i) If full second-order logic is logic, then either (CH) is a genuine validity or (NCH) is.

(ii) If (CH) is a genuine validity, then it's an error of reasoning to reject (CH).

(iii) If (NCH) is a genuine validity, then it's an error of reasoning to reject (NCH).

(iv) So if full second-order logic is logic, then it's either an error of reasoning to reject (CH) or an error in reasoning to reject (NCH).

[19] At least, if the 'we' are philosophers; [21] argues at length that actual mathematical practice, which presumably includes that of model-theory, is rife with second-order reasoning. I cannot evaluate that claim here.

[20] I assume that if it is an error in reasoning to A, then we have an epistemic obligation to not A. I will sometimes slide between error-talk and obligation-talk in the text.

(v) It is neither an error in reasoning to reject (CH) nor an error in reasoning to reject (NCH).

(vi) Therefore, full-second-order logic is not logic.

There are a number of routes for resistance; some look more promising than others. Let's go through a few in turn.

4.2.1 Agnosticism

We might be skeptical about premise (v). We might think that, given the deep epistemic openness about the continuum hypothesis, being opinionated on *either* count would be epistemically unwarranted. Agnosticism about the continuum hypothesis is plausibly the only epistemically responsible route. But agnosticism plausibly entails that we reject *neither* (CH) nor (NCH), so it is not compatible with premise (v).

Since premise (v) tempts us, those who would reject it should explain its appeal. The story, presumably, is that we get tempted because we conflate non-rejection with acceptance. If Normativity told us that we had to either accept (CH) or accept (NCH), that *would* be bad. But telling us to not reject a claim is far weaker than telling us to accept it, because the former but not the latter leaves agnosticism as an option.

I can imagine two further responses to this reply. The first response tries to beef up the original argument by appealing to a stronger form of normativity,

Strong Normativity: If $\Delta \Rightarrow \phi$, then anyone who accepts all of Δ (and meets some condition C) makes an error in reasoning if they do not accept ϕ, too,

and using this stronger form to argue that we must accept either (CH) or (NCH).

Why the rider 'and meets some condition C'? Because otherwise counterexamples will be too easy to come by. If the argument from Δ to ϕ is too long and complex, or if the logical structure of ϕ itself is too difficult for us to get into our heads, then plausibly we do nothing epistemically wrong if we don't accept it.

The need for the rider saddles the proponent of this response with the unenviable task of navigating between the cliffs of counterexample and the shoals of satisfiability. For if C is too weak, Strong Normativity will tell us we have obligations we clearly don't have. If C is too strong, though, then even if the beefed-up argument shows that some possible agents have to accept either (CH) or (NCH), if we mere mortals don't satisfy the condition, it doesn't show anything untoward about *our* epistemic obligations. Whether the responder can chart just the right course is something I'll not consider further here.

The second line of response is more subtle, and runs like this. 'Okay, so (v) is false. We should reject neither (CH) nor (NCH) because we lack the relevant evidence to come to an opinion. But this kind of impermissibility-of-rejection is *weaker* than the sort had by genuine validities. For instance, if I suppose that p and then reason myself into rejecting a genuine validity, that gives me epistemic warrant to reject the supposition. That's how *reductio* reasoning works. But supposing that p and then reasoning myself into rejecting something which, as a matter of fact, I don't have evidence to reject gives me no warrant for rejecting p. (CH) and (NCH) may be both impermissible-to-reject in the weaker, evidential sense, but not the stronger sense — but one of them would be if it were a genuine validity.'

Call a claim ϕ *reductio fodder* iff, if we reason to the rejection of ϕ under a supposition, that gives us warrant to reject the supposition. Then the second line of response suggests a modification of the original argument:

Reductio Sheep's Clothing

(i*) If full second-order logic is logic, then either (CH) is a genuine validity or (NCH) is.

(ii*) If (CH) is a genuine validity, then it is reductio fodder.

(iii*) If (NCH) is a genuine validity, then it is reductio fodder.

(iv*) So if full second-order logic is logic, then either (CH) is reductio fodder or (NCH) is.

(v*) Neither (CH) nor (NCH) is reductio fodder.

(vi*) Therefore, full-second-order logic is not logic.

It is far more difficult to deny (v*) than (v). Denying it would seem to undercut the considerably entrenched mathematical practice of showing various results to hold in the presence or absence of the continuum hypothesis. Presumably, if I suppose the truth of the continuum hypothesis, I can reason my way to the rejection of (NCH); but even if the continuum hypothesis is false (and so (NCH) reductio fodder), we would not think we could use reductio reasoning to come to know it this way.

I suspect many readers will find the original argument less compelling than this new one, so it so it might be prudent to focus on the latter from here on in. But I won't follow this wise counsel. The next few sections consider further objections to the original argument. This is of small moment, though, as the objections to be considered will apply, with only small modifications, to the Reductio argument too.

4.2.2 Indeterminacy

A number of philosophers think there's just no fact of the matter about whether the continuum hypothesis is true. Suppose that's right, and suppose second-order logic is logic. Then there will just be no fact of the matter as to whether (CH) or (NCH) is genuinely valid, either — even if, definitely, one of them is.

Normative Sheep's Clothing uses a form of dilemma reasoning, and many treatments of indeterminacy rule certain dillema-like inferences. Foes of the argument may be tempted to challenge the argument's validity on those grounds.

But that would be a mistake. Distinguish *direct* dilemma reasoning from the *indirect* sort, or 'proof by cases'. The direct form runs:

> *A* or *B*.
> If *A*, then *C*.
> If *B*, then *C*.
> Therefore, *C*.

The indirect form removes the second and third premises and replaces them with a procedure: suppose the antecedents and then prove the consequents under those suppositions. Proof by cases takes (direct) dilemma reasoning and replaces each conditional premise with an indirect proof.

Some treatments of indeterminacy invalidate indirect proof for conditionals. For instance, in supervaluational treatments we can derive $\ulcorner$Determinately, $\phi\urcorner$ from ϕ, but we cannot then go on to prove $\ulcorner\phi \rightarrow$ Determinately, $\phi\urcorner$. So indirect proof fails. In these systems, proof by cases fails for essentially the same reason. But direct disjunctive dilemma is valid in this system. Since Normative Sheep's Clothing relies on direct disjunctive dilemma, it can't be charged with straightforward invalidity, even on these treatments of indeterminacy.

If the continuum hypothesis is indeterminate, though, there may be a better criticism of Normative Sheep's Clothing in the neighborhood. Suppose that we are looking at a rose with a color in the penumbra between pink and red. Now consider this argument:

The Rose Argument

(i′) The rose is either red or pink.

(ii′) If the rose is red, we ought to believe that it is red.

(iii′) If the rose is pink, we ought to believe that it is pink.

(iv′) So we either ought to believe that the rose is red, or we ought to believe that it is pink.

Supposing vagueness is a species of indeterminacy, the Rose Argument parallels Normative Sheep's Clothing in interesting ways. The first premise is (on many theories of vagueness) true, and the argument valid. But the conclusion seems to commit us to the unacceptable claim that we ought to hold a definite opinion about whether a rose is red, even when it is a borderline case of being red.

One plausible diagnosis of the Rose Argument has it that premises (ii′) and (iii′) are false, but tempting. They are tempting because we tend to evaluate them by first imagining that we are in a position to assert 'the rose is red' or 'the rose is pink,' and then asking what we ought to believe in those cases. On many treatments of indeterminacy, we will only be in a position to assert either of these if the rose is *determinately* red or pink. We thus evaluate the conditionals by an imaginative analogue of indirect proof; but just as indirect proof gives us the wrong result when indeterminacy is involved, this imaginative process does, too.

In other words, we're tempted to accept the false premises (ii′) and (iii′) because our evaluative process confuses them with the true

(ii″) If the rose is determinately red, we ought to believe that it is red.

(iiii″) If the rose is determinately pink, we ought to believe that it is pink.

These are true; but since it is *not* true that the rose is either determinately red or determinately pink, we can't use (ii″) and (iii″) to get (iv′).

If this is the right diagnosis of the Rose Argument, and if the continuum hypothesis is in fact indeterminate, then a similar diagnosis may fit Normative Sheep's Clothing.[21] For if the continuum hypothesis is indeterminate and second-order logic is logic, then although it may be true (and determinately true) that either (CH) or (NCH) is genuinely valid, it should be indeterminate which one is. But plausibly we only need to be governed by the determinate facts of logic: if it's unsettled which of several claims is genuinely valid, then our epistemic obligations ought not favor one over the other.

If *that's* right, though, then the Normativity Constraint from §1.2 needs to be tweaked: it's only an error in reasoning to accept all of Δ but reject ϕ when the latter is *determinately* a genuine consequence of the former. Since we can't get premises (ii) and (iii) of Normative Sheep's Clothing from this modified constraint, the argument fails.

4.2.3 Weakening the Logic

Premise (i) of the Normative Sheep's Clothing argument is underwritten by two claims. The first is that, if the continuum hypothesis is true, $\vDash_F$ (CH), and if it is false, $\vDash_F$ (NCH). The second is that, if full second-order logic is logic, then if $\vDash_F \phi$, ϕ is a genuine validity.

Motivation for the second claim stems from our definition of (full) second-order logic being *logic*: There is some second-order system SO where $\Delta \Rightarrow \phi$ iff $\Delta \vDash_{SO} \phi$.

[21] It may not be the right diagnosis, of course. Another plausible diagnosis [4] has it that the Rose Argument is sound and we should satisfy the obligations of (iv′) by getting ourselves into a position where it's indeterminate whether we believe that the rose is red or believe instead that the rose is pink.

Motivation for the first claim comes from the following mathematical fact: First-order model theory, as standardly defined, is such that every model is a model of (CH) if the continuum hypothesis is true, and a model of (NCH) if the continuum hypothesis is false.

Here is a different conception of 'being logic', one endorsed by Field [5], who claims Kreisel [10] as inspiration. Recall that formal systems, as described in §1.1, include both a proof-theoretic and a model-theoretic component. On Field's picture, each of those components has an important job to do. When devising a formal system S for genuine logical consequence, we want to pick one where the proof system is 'genuinely sound': if $\Delta \vdash_S \phi$, then $\Delta \Rightarrow \phi$. And we want to choose a system with a model theory that is 'genuinely complete': if there is a model of Δ without ϕ — that is, if $\Delta \not\models_S \phi$ — then $\Delta \not\Rightarrow \phi$.

If the system is complete — if every model-theoretic consequence of the system is also one of its proof-theoretic consequences — then we get a 'squeezing argument' to show us that genuine consequence coincides with both the model-theoretic and the proof-theoretic relation. If the system is incomplete, though — as every full second-order system will be — the best we can say is that genuine consequence lies somewhere in between the proof- and model-theoretic relations. All proof-theoretic consequences are genuine consequences, and all genuine consequences are model-theoretic consequences. But that's all we can say.

From this perspective, we might think that full second-order logic is 'logic' iff there is a second-order system F which is genuinely sound and genuinely complete. In that case, saying that second-order logic is logic will only commit us to genuine consequence lying somewhere 'in between' $\models_F$ and $\vdash_F$. So, so long as F won't let us *derive* either (CH) or (NCH) (which, presumably, it won't), we can grant that, if the continuum hypothesis is true, $\models_F$ (CH), while insisting that it's still not a genuine validity — for this may be a case where the model-theoretic relation outstrips genuine consequence.

This undercuts the argument. But it does so at a cost. Friends of full second-order logic generally like it precisely because of its greater expressive power. Some may *want* either (CH) or (NCH) to be genuinely valid; they'll have no truck with this style of response. But even those who are happy to let (CH)'s or (NCH)'s validity go will still want full second-order logic to be genuinely stronger than the Henkin variety. If *all* we can say is that consequence lies between proof- and model-theory, then — since the proof theory can give us no logic stronger than a Henkin logic — we'll have no guarantee that genuine logic is any stronger than Henkin.

Of course, friends of this response may simply dig in their heels and insist that it *is* stronger. 'Formal systems are for helping us *investigate* consequence, not for *telling us what it is*. I don't guarantee anything about logic by specifying a formal system. Rather, the consequences are what they are; formal systems just help us reason about them. I happen to think that genuine consequence is almost-but-not-quite as strong as $\vDash_F$, and whether I can specify some model theory that exactly tracks it or not is beside the point.'

Fair enough; but we may legitimately wonder whether, after making this speech, the speaker can go on to learn very much at all from a formal system. Suppose he discovers that ϕ is a model-theoretic consequence of Δ. Can he conclude anything about their logical relationship? Unfortunately not — this may be one of those cases where model-theoretic consequence outstrips genuine consequence. In fact, his formal system will only let him be sure that ϕ is a consequence of Δ when ϕ is provable from (and then a Henkin consequence of) Δ. Of course, he might just insist he knows the genuine consequences when he sees them — but then he can forget the formal system entirely and just rely on his consequence-seeing faculties.

We may instead stick with our original gloss on 'being logic' and resist (i) by resisting the first claim that underwrote it. We cannot resist that claim by denying that, as standardly defined, every full second-order model is a model of (CH) if the continuum hypothesis is true and a model of (NCH) if the continuum hypothesis is false. That's a mathematical claim and not legitimately up for grabs. But we might try to deny it by giving a *non*-standard definition of a full model-theoretic consequence relation. We won't define $\vDash_F$ as 'truth in all second-order models', but as something more sophisticated cobbled together from these resources.

One way it might go stems from our thoughts about indeterminacy from §4.2.2. In that section, we considered the idea that the Normativity constraint ought to govern only determinate consequence. In other words, we granted that, determinately, if $\vDash_F \phi$, then ϕ is genuinely valid, but suggested that if it is indeterminate whether ϕ is valid, we are free to reject ϕ. But we could break the link earlier in the chain, insisting that if it is indeterminate whether $\vDash_F \phi$, then ϕ is *not* genuinely valid. We then identify the genuine validities as those which, determinately, are true on every model.

More precisely, given the standard model-theoretic relation $\vDash_F$, we can define another one, $\vDash_{F^D}$, where $\Delta \vDash_{F^D}$ iff, determinately, $\Delta \vDash_F$. The formal system F^D is the one just like F except that we swap $\vDash_{F^D}$ for $\vDash_F$. Clearly, F^D should count as second-order system if F does.[22] The friend of second-order logic can think that it 'is logic' because $\Delta \vDash_{F^D} \phi$ iff $\Delta \Rightarrow \phi$ without being committed to either (CH) or (NCH) being genuinely valid.

Unlike the Kreisel-based rejection of (i), this move gives us a clearer picture of exactly which model-theoertic consequences correspond to genuine consequences: they are the *determinate* ones. Even if only the determinate holding of $\vDash_F$ gives us genuine consequences, that's enough to show that genuine second-order con-

[22] Should it count as a *full* second-order system? It's genuinely weaker than the (usual) full system, but stronger than Henkin systems (see the next note). I doubt usage is fixed enough to settle this question.

sequence outstrips Henkin consequence.[23] And, insofar as we have a fairly good grip on which model-theoretic results depend on claims of dubious determinacy and which do not, we can use the second-order model theory as a tool to investigate genuine consequence.

4.2.4 Rejecting Normativity

A final option is nuclear: simply reject the normativity constraint.

Let's be a bit more precise. Aside from a brief parenthetical remark in §1.2, I have been talking as though there is a unique genuine consequence relation that we care about. But that's not at all obvious. There may be a number of different relations, each of which deserves to be called a consequence relation.

If that's right, it may be that not all such relations obey, or obey equally well, all of the constraints of §4.2.2. Perhaps, for instance, the kind of normativity we care about in logic ones apart from the kind of necessity we care about. This idea isn't new: David Kaplan ([9]; see also [20]) suggests that 'I am here now' is a kind of logical validity, and the reason is at least partly because no one should reject it. Yet it's clearly not necessary: I may be here now, but I could have been somewhere else instead.

If normativity and necessity come apart, then there may be several candidate 'logical consequence' relations: one that tracks the distinctive normativity of logic, one that tracks the distinctive necessity of logic, and so on. Shapiro [22: pp. 772–773] has suggested something more-or-less along these lines. And he has suggested that full second-order logic is logic precisely because it tracks the distinctive necessity — normativity be hanged.[24]

5 Scorecard

Rather than trying to exonerate second-order logic, I have merely provided a brief for the defense. That defense will find some challenges more worrying than others. In particular, the Intelligibility and Normative Sheep's Clothing arguments present deeper difficulties than Textbook Quineanism or Topical Sheep's Clothing. Still,

[23] The quick-and-dirty way to show this is to note that full second-order logic has sentences that characterize infinite models. So long as our notion of infinity isn't itself indeterminate, this means that (determinately) we have a sentence that is entailed by an infinite set ('there is at least one thing,' 'there is at least two things,' ...) but not any of its finite subsets. So, unlike Henkin systems, the system F^D is not compact.

[24] This may not be an entirely fair characterization of Shapiro's view, in large part because he is considering a normative constraint somewhat stronger than the one outlined in §1.2. Still, this captures the basic idea, and is a move available here to friends of second-order logic worried about our weaker Normativity constraint.

the friend of second-order logic has several defensive avenues available. I leave it to her to decide which, if any, she wishes to take.

References

[1] Barcan Marcus, R. (1972). Quantification and ontology. Nous 6(3), 240–250.

[2] Beall, J. C. and G. Restall (2006). Logical Pluralism. Oxford: Oxford University Press.

[3] Boolos, G. (1975). On second-order logic. The Journal of Philosophy 72(16), 509– 527.

[4] Dorr, C. (2003). Vagueness without ignorance. Philosophical Perspectives 17(1), 83–113.

[5] Field, H. (1991). Metalogic and modality. Philosophical Studies 62(1), 1–22.

[6] Field, H. (2008). Saving Truth from Paradox. Oxford: Oxford University Press.

[7] Haack, S. (1978). Philosophy of Logics. Cambridge: Cambridge University Press.

[8] Henkin, L. (1950). Completeness in the theory of types. The Journal of Symbolic Logic 15(2), 81–91.

[9] Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (Eds.), Themes from Kaplan, pp. 481–563. New York: Oxford University Press.

[10] Kreisel, G. (1967). Informal rigor and completeness proofs. In I. Lakatos (Ed.), Problems in the Philosophy of Mathematics, pp. 138–171. Amsterdam: North-Holland.

[11] Lewis, D. (1970). How to define theoretical terms. The Journal of Philosophy 67, 427–446. Reprinted in Lewis, D. (1983), Philosophical Papers, Volume 1. Oxford: Oxford University Press, pp. 78–95.

[12] Prior, A. N. (1971). Objects of Thought. Oxford: Clarendon Press.

[13] Quine, W. (1940). Mathematical Logic. Cambridge, Mass.: Harvard University Press.

[14] Quine, W. (1960). Word and Object. MIT Press.

[15] Quine, W. (1970). Philosophy of Logic. Englewood Cliffs, N.J.: Prentice-Hall.

[16] Rayo, A. (2008). On specifying truth-conditions. The Philosophical Review 117, 385–443.

[17] Rayo, A. and T. Williamson (2003). A completeness theorem for unrestricted first-order languages. In J. C. Beall (Ed.), Liars and Heaps: New Essays on Paradox, Chapter 15, pp. 331–356. Oxford: Oxford University Press.

[18] Rayo, A. and S. Yablo (2002). Nominalism through de-nominalization. Noûs 35(1), 74–92.

[19] Rossberg, M. (forthcoming) Somehow things do not relate: On the interpretation of polyadic second-order logic. *The Journal of Philosophical Logic*.

[20] Russell, G. (2008). One true logic? The Journal of Philosophical Logic 37(8), 593– 611.

[21] Shapiro, S. (1991). Foundations Without Foundationalism. Oxford University Press.

[22] Shapiro, S. (2005). Higher-order logic. In S. Shapiro (Ed.), The Oxford Handbook of Philosophy of Mathematics and Logic, pp. 751–780. Oxford University Press.

[23] van Inwagen, P. (1981). Why I don't understand substitutional quantification. Philosophical Studies 39, 281–285. Reprinted in van Inwagen, P. (2001). Ontololgy, Identity, and Modality. Cambridge: Cambridge University Press, pp. 32–36.

[24] van Inwagen, P. (1998). Meta-ontology. Erkenntnis 38, 223–250. Reprinted in van Inwagen, P. (2001). Ontololgy, Identity, and Modality. Cambridge: Cambridge University Press, pp. 13–31.

[25] van Inwagen, P. (2004). A theory of properties. In D. W. Zimmerman (Ed.), Ox- ford Studies in Metaphysics, Volume 1, pp. 107–138. Oxford: Oxford University Press.

[26] Williams, J. (210). Fundamental and derivative truths. Mind 119(473), 103–141.