QUANTIFICATIONALISM

ANTONIO MARIA CLEANI

Introduction

Philosophers disagree about whether tense is real.[1] What exactly that means is a point of contention. One familiar piece of technology used to articulate this disagreement is the distinction between *temporary* and *eternal* propositions.[2] A *temporary* proposition is one that is neither always true nor always false: the proposition that I am sitting is true now, but will be false later. An *eternal* proposition, on the other hand, is one that has a fixed truth value across all times: the proposition that I am sitting at 11:00 am on February 12, 2026 is true now, has always been true, and will always be true. At a first approximation, to be a realist about tense is to hold that there are temporary propositions, while to be an anti-realist about tense is to hold that all propositions are eternal.[3]

There is a parallel disagreement about whether modality is real.[4] We can elucidate it in a similar way. A *contingent* proposition is one that is neither necessarily true nor necessarily false. The proposition that Timothy Williamson is a contingentist is false, but could have been true. A *non-contingent* proposition, on the other hand, is one that has a fixed truth value across all possible worlds: the proposition that Timothy Williamson is actually a contingentist is false and could not have been true. To be a realist about modality is to hold that there are contingent propositions, while to be an anti-realist about modality is to hold that all propositions are non-contingent.[5]

In this paper, I introduce another debate of this sort, which has been all but ignored in the extant literature, and explore the realist side of it.[6] The debate concerns quantification. I argue that there is sense to be made of a distinction

Draft of November 6, 2026.

[1]The debate is often framed as the question whether we should "take tense seriously." Contemporary discussion is usually traced back to McTaggart (1908). Realism was championed by Prior (1957, 1967); Fine (2006a) is an influential contemporary realist. Prominent anti-realists include Russell (1915); Goodman (1951); Quine (1960); Mellor (1981); Lewis (1986); Sider (2001). The debate on the reality of tense is sometimes distinguished from that over whether the present is "metaphysically privileged" (Zimmerman, 2005; Fine, 2005), though not always (Sider, 2001, Ch. 1).

[2]See especially Richard (1981); Brogaard (2012); Bacon (2018); Dorr and Goodman (2020).

[3]Not all the anti-realists mentioned earlier count as anti-realists in my sense. Sider (2001), for example, thinks tense should be simply paraphrased away, so it is not clear whether he would accept the claim that every proposition is eternal.

[4]The realist side features Prior (1957); Prior and Fine (1977) again; see also Yagisawa (2010)—especially his notion of "modal tense"—and Williamson (2013). For the anti-realist side see Spinoza (1985); Quine (1953); Schaffer (2012).

[5]As before, it is not clear whether Quine (1953) is an anti-realist in this sense, though Spinoza (1985) and Schaffer (2012) arguably are.

[6]The lone exception appears to be Stojanovic (2012), though her framing of the debate differs significantly from mine.

between *domain-variable* and *domain-invariable* propositions. This distinction is to quantification what the distinction between temporary and eternal propositions is to temporal discourse. A domain-variable proposition is one whose truth value varies across *domains of quantification*: the proposition that there are no even prime numbers is false, but true *with the exception of* 2. A domain-invariable proposition, on the other hand, is one that has a fixed truth value across all domains of quantification: the proposition that except for 2, there are no even prime numbers is true, and indeed true across all domains of quantification.

A realist about domain-variability (for lack of a better term) is someone who holds that there are domain-variable propositions, while an anti-realist about domain-variability is someone who holds that all propositions are domain-invariable. I precisely formulate a particular version of realism about domain-variability, which I call *quantificationalism*, and explore the metaphysical picture that comes with it.

I argue that quantificationalism should be taken seriously because it is philosophically promising. It has a wide range of applications in topics as diverse as metaontology, the metaphysics of modality, predicativity, and (more tentatively) potentialism in the philosophy of mathematics. At the same time, it is not without costs. While precisely formulating versions of realism about tense and modality is relatively straightforward, doing the same for quantificationalism requires considerably more care. I show that this can be done by understanding the notion of domain-occasional truth using the apparatus of *metaphysical substitutions* (Bacon, 2019, 2020). The resulting framework is revisionary, but ultimately illuminating.

Here is the plan for the paper. Section 1 articulates a minimal version of realism about tense, which serves as a model for the formulation of quantificationalism in Section 2. Section 3 sketches a logical framework for quantificationalism and clarifies key pieces of its ideology. Section 4 develops one application of quantificationalism in some detail, and Section 5 outlines a few more. I conclude in Section 6.

The paper also contains two appendices. Appendix A discusses the motivations for a non-standard design choice in the syntax of my logic for quantificationalism, and explores the revisionary metaphysics of identity that arises from it. Appendix B presents the technical details of the framework in compressed form; an extended treatment can be found in the companion paper Cleani (Manuscript).

1. Warm-Up: Temporalism

I begin by giving a more careful articulation of quantificationalism. It will be helpful to proceed by analogy: I will first articulate a precise version of realism about tense I call *temporalism*, then introduce quantificationalism as a version of realism about domain-variability modeled after temporalism.

I have formulated realism about tense as the claim that there are temporary propositions—propositions that are neither always true nor always false. In this statement, I use quantification over propositions as a mere shorthand for higher-order quantification in sentence position. A realist about tense, in my sense, is someone who accepts

Temporalism: For some p, it is neither always the case that p nor always the case that not p.[7]

[7]Temporalism in the sense just defined is discussed in Bacon (2018); Dorr and Goodman (2020), and it is implicit in Arthur Prior's work.

We should take care not to conflate temporalism with different, and generally stronger, ways of articulating realism about tense. First, we should distinguish temporalism from what we may call *first-order* tense realism. This view articulates tense realism in terms of first-order quantification over propositions, a special sort of abstract object satisfying a certain theoretical role.[8] I find first-order tense realism an unsatisfactory formulation of tense realism, for two main reasons.

First, it threatens to make the debate over the reality of tense a merely verbal one, since it is resolved differently depending on what we take propositions to be. Propositions have been taken to serve both as the objects of attitudes like belief and desire, and as the semantic content of sentences. But it is not obvious that there is a single kind of thing that can play both these roles. Notably, Lewis (1980) argues the semantic contents of sentences cannot be the same as the objects of belief, on the basis that the former but not the latter vary in truth value across times. If Lewis is right, we cannot take propositions to play both roles. We must make a choice, and the truth of realism about tense depends on which choice we make.

Second, one might worry that this formulation saddles realism about tense with irrelevant ontological commitments. No matter whether propositions are taken to be semantic values of sentences or objects of attitudes, most philosophers agree that propositions are abstract objects. Nominalists, who do not believe in abstract objects, do not believe in propositions, and a fortiori do not believe in temporary propositions. Thus, no nominalist can be a realist about tense.

Temporalism does better than first-order tense realism on both these fronts. In classical higher-order logic, the meaning of quantifiers in sentence position is uniquely determined by the introduction and elimination rules governing them, up to logical equivalence (Harris, 1982). Moreover, as argued at length by proponents of higher-order logic (e.g., Bacon, 2024a), quantification in sentence position does not carry any ontological commitment to propositions, understood as first-order abstract objects. I thus set aside first-order tense realism, and throughout the paper use expressions like 'some propositions' as shorthands for higher-order quantification in sentence position. Similar conventions apply to expressions like 'some properties,' 'all relations,' and so on.

Tense realism is also sometimes understood in terms of fundamentality-theoretic notions. For example, one might understand tense realism as the claim that, *fundamentally*, there are temporary propositions.[9] On this construal of the debate, someone might accept the existence of temporary propositions while still counting as an anti-realist about tense. For example, they might think that all temporary propositions are *grounded* in eternal propositions, and therefore not fundamental.

Temporalism does not make any claim about the fundamentality of temporary propositions. It simply says that there are some. I chose this weaker formulation because I find fundamentality-theoretic strengthenings of temporalism unnecessary distractions at the present dialectical stage. My goal is to introduce a philosophical position, realism about domain-variability, by analogy with tense realism. Since realism about domain-variability has not been articulated before, it makes sense to

[8]This is the version of tense realism found in Brogaard (2012) and denied, e.g., by Richard (1981).

[9]Examples include Zimmerman (2005); Fine (2005). Of course, these views can be further distinguished by what sort of fundamentality-theoretic ideology they employ.

analogize it with the simplest version of tense realism, which I take to be temporalism.

In a similar vein, sometimes the central questions in the debate over the reality of tense are taken to be about how tense operators like 'always' and 'it was the case that' relate to quantification over times. A realist insists that tense operators are primitive, and quantification over times is legitimate only if analyzed in terms of tense operators (Prior and Fine, 1977). On the other hand, an anti-realist insists that quantification over times is legitimate in its own right, and tense operators must be reduced to quantification over times.

There is at least one sense in which temporalism is perfectly compatible with the claim that quantification over times is legitimate in its own right. Consider:

Tense Leibnizianism: For any p, for it to be the case that, sometimes, p is for there to be a time t such that, at t, p.

If we deny that quantification over times must be reduced to the tense operators 'sometimes' and 'always,' then this appears to be a definition of 'sometimes' in terms of quantification over times. But as long as there is some p and some times t and t' such that, at t, p but at t', not p, tense Leibnizianism is consistent with temporalism. Indeed, given tense Leibnizianism, a temporary proposition can be *defined* as a proposition p such that, at t, p but at t', not p, for some times t and t'.

This point bears emphasis. Temporalism, essentially, is the thesis that the tense operators 'sometimes' and 'always' are non-trivial: they can turn a truth into a falsehood or vice versa. This claim is orthogonal to the question whether we regard quantification over times as legitimate in its own right or not. A temporalist is free to accept times as legitimate entities and use them to introduce additional, more specific tense operators that regiment the intuitive notion of truth at a time. Namely, for each time t, one introduces a further tense operator 'at t,' which will also in general be non-trivial.[10] These operators can be used to illuminate the logical behavior of the more general tense operators 'sometimes' and 'always,' as done in tense Leibnizianism, but they need not be.

Last, temporalism can be combined with a natural analysis of tense in natural language, but should not be conflated with the latter. To illustrate, consider the sentences:

(1) No Western country has all five of Earth's climate zones.
(2) In 1783, no Western country had all five of Earth's climate zones.

Now, (1) is false, because the United States does have all five of Earth's climate zones. In fact, it is the only Western country to have enjoyed this much climate diversity in at least the last half century. But (2) is true: the western border of the United States in 1783 was the Mississippi River, so dry climates were not present in the United States at that time.

[10] Prior (1967); Prior and Fine (1977) construct instants as propositions that are sometimes the case and strictly imply, in the sense of eternal material implication, any proposition they are sometimes true together with. In such a setting 'at t' can be defined as the operator of being eternally implied by t. To Prior and Fine, this kind of quantification over times is legitimate because times themselves are construed in terms of tensed ideology. A temporalist in my sense, on the other hand, is free to embrace tense Leibnizianism while rejecting any analysis along these lines, so long as they maintain that 'at t' is a sentential operator.

What explains the difference in truth value between (1) and (2)? A temporalist has a simple explanation: (2) is constructed by applying a non-trivial tense operator, 'in 1783,' to (1). The latter expresses the same proposition when it occurs within the scope of 'in 1783.' Because 'in 1783, p' entails 'sometimes, p,' (1) is a temporary falsehood: false now, but true in the past. Call a temporalist who goes for an explanation along these lines a *linguistic temporalist*.[11]

Anti-realists who deny temporalism offer an alternative explanation for the difference in truth value between (1) and (2).[12] They insist that both (1) and (2) express eternal propositions, and deny that (1)—if it expresses a proposition at all—expresses the same proposition as its embedded occurrence in (2). Here is one way this thought can be fleshed out further. Somewhere in the syntax of (1) there is a free occurrence of a covert variable that takes times as values. When one utters (1) at a time t, the context of utterance assigns the value t to this variable, and the utterance expresses the proposition that, at t, no Western country has all five of Earth's climate zones. When t is, say, 11:00 am on February 12, 2026, the resulting proposition is an eternal falsehood. The modifier 'in 1783' in (2) works by binding this time variable and specifying some moment in 1783 as its value. Thus, for any t, when one utters (2) at t, that utterance expresses the eternal truth that, in 1783, no Western country had all five of Earth's climate zones. So we have an explanation of the variation in truth value between (1) and (2) that does not appeal to temporary propositions.

Call someone who gives an explanation along these lines of tensed discourse a *linguistic anti-temporalist*. Obviously, temporalism is compatible with linguistic temporalism. But I claim it is also compatible with linguistic anti-temporalism.

This point might seem surprising, so allow me to elaborate. Suppose you like the linguistic anti-temporalist explanation of (1) and (2). You think that, as a matter of natural language syntax and semantics, (1) and its embedded occurrence in (2) express different propositions, and that the modifier 'in 1783' is not a sentential operator at all. Even so, the following question can be raised: is there a way to *introduce* a sentential operator that behaves like 'in 1783' in the sense that it can be used to express the same propositions we express in English by prefixing sentences with 'in 1783'?

More precisely, is there a proposition p such that, for any time t, there is an operator $\mathsf{At}(t)$ with the following properties:

- The proposition expressed by 'No Western country has all five of Earth's climate zones' when uttered at t is identical to $\mathsf{At}(t)(p)$, and
- The proposition expressed by 'in 1783, no Western country had all five of Earth's climate zones' when uttered at t is identical to $\mathsf{At}(1783)(p)$?

Since $\mathsf{At}(t)(p)$ must be false but $\mathsf{At}(1783)(p)$ true when t is the time of writing, someone who answers 'yes' to this question counts as a temporalist, even though

[11] Prior (1957, 1967) is a linguistic temporalist in my sense. Brogaard (2012) endorses a similar linguistic thesis, though formulated in first-order terms. See also Kaplan (1989) and the subsequent literature on "operator arguments" (Yli-Vakkuri et al., 2019).

[12] This is the standard story told by all previously cited tense anti-realists, though in first-order terms.

they believe tensed discourse in natural language is best accounted for without using temporary propositions.[13]

In view of this, we should reformulate temporalism slightly:

Temporalism: For some p, it is neither always* the case that p nor always* the case that not p.

Here, 'it is always* the case that' is a quasi-technical expression that stands to the English expression 'it is always the case that' in more or less the same way that '$\mathsf{At}(1783)$' stands to the English expression 'in 1783.' A linguistic anti-temporalist might not think that 'it is always the case that' is itself a sentential operator, but for independent reasons might posit a sentential operator 'it is always* the case that' that can be used to express the same propositions we express in English by prefixing sentences with 'it is always the case that.' On these grounds, they count as a temporalist in my sense.

This example illustrates that temporalism is, at its core, a thesis about *granularity*. Temporalism asserts that reality is fine-grained enough for the notion of temporary truth to make sense. Extremely coarse-grained pictures of reality falsify temporalism. For example, assume propositions are individuated by material equivalence: two propositions are the same iff they have the same truth value. It follows that there are only two propositions, truth and falsity. In this case, temporalism must be false. For surely logical truths are always true, and contradictions are always false. So, all true propositions are always true, and all false propositions are always false. But any proposition is identical, on this picture, with either a logical truth or a contradiction. There simply are not enough propositions around for us to introduce operator expressions that faithfully regiment tensed discourse.

2. From Temporalism to Quantificationalism

At a first approximation, *domain-variability* is the apparent dependence of truth and falsity on what is included in the domain of quantification. It is false that all prime numbers are odd, but *with the exception of 2*, all prime numbers are indeed odd. Some mammals lay eggs, yet it is false that, *in addition to the platypus and the echidna*, some (other) mammals lay eggs.

A realist about domain-variability is someone who holds that there are domain-variable propositions, which change their truth values across domains of quantification much the same way as temporary propositions change their truth values across times. My preferred version of realism about domain-variability is *quantificationalism*, which I formulate as follows:

Quantificationalism: For some p, it is neither domain-invariably the case that p, nor domain-invariably the case that not p.

As in temporalism, quantificationalism cashes out quantification over propositions as higher-order quantification in sentence position. The expression 'it is domain-invariably the case that' is a theoretical term stipulated to occupy the position of a sentential operator, and to express generality with respect to domains of quantification the way that 'it is always the case that' expresses generality with respect to time. I will introduce the term 'it is domain-occasionally the case that' as an

[13] You might think that sentences express propositions not relative to a time of utterance alone but relative to a *context*, for some richer notion of context. The same point applies: replace 'when uttered at t' with 'when uttered at a context that specifies t as the salient time.'

abbreviation for the operator 'it is not domain-invariably the case that not,' which is the analog of 'it is sometimes the case that.' While these operators are certainly less familiar than their temporal analogs, I hope their content will get clearer as we unpack quantificationalism further.

It is important not to conflate quantificationalism with a number of theses in its vicinity. Here, the lessons of the previous section carry over. Quantificationalism is a higher-order generalization, and thus must be distinguished from the claim that there are domain-variable propositions, understood as a special sort of abstract object satisfying a certain theoretical role.[14] Quantificationalism makes no claim about the fundamentality of domain-variable propositions, and is independent of the idea that quantification over domains (whatever that may be) is legitimate in its own right.

Natural language contains expressions that appear to operate on domains of quantification, and we need to be careful not to conflate quantificationalism with a thesis about the syntax and semantics of these expressions. It is certainly true that one *might* be drawn to quantificationalism as a framework for providing a syntactic and semantic analysis of these expressions. Consider:

(3) No Western country has all five of Earth's climate zones.
(4) With the exception of the United States, no Western country has all five of Earth's climate zones.

We noted earlier that the United States is the only Western country to have all five of Earth's climate zones. Thus, (3) is false, while (4) is true. A realist about domain-variability can explain this difference in truth value in a way that mirrors the linguistic temporalist explanation of the difference in truth value between (1) and (2). They might say that (4) is constructed by applying a non-trivial operator expression, 'with the exception of the United States,' to (3). The latter expresses the same proposition when it occurs within the scope of 'with the exception of the United States.' Because 'with the exception of the United States, p' entails 'domain-occasionally, p,' (3) is a domain-occasional falsehood: false, but true relative to the domain of things other than the United States. A quantificationalist who goes for an explanation along these lines is a *linguistic quantificationalist*.[15]

Yet, as before, this sort of account is by no means forced upon the quantificationalist. Quantificationalists, who might have independent reasons to believe in domain-invariable propositions, are free to deny linguistic quantificationalism, and opt to explain the difference in truth value between (3) and (4) purely in terms of domain-invariable propositions.

The alternative explanation proceeds as in the case of tense. For whatever reason, when one utters (3) in a context where the salient domain of quantification includes the United States, the utterance expresses a domain-invariable falsehood: the proposition that, *in the contextually salient domain* d that includes the United States, no Western country has all five of Earth's climate zones. If, in the same context, one utters (4), the utterance expresses a domain-invariable truth: the proposition that, in the domain that results from removing the United States from d, there is no Western country having all five of Earth's climate zones. Perhaps

[14] Stojanovic (2012) defends a position of the latter sort. Stojanovic calls *domain sensitivity* roughly what I call domain-variability, though hers is a first-order notion.

[15] Stojanovic (2012) endorses a first-order version of linguistic quantificationalism.

this happens because the syntax of (3) contains covert occurrences of free variables ranging over domains of quantification.[16]

To the extent they wish to tie genuine domain-variability to natural language, a quantificationalist who denies linguistic quantificationalism might insist that, even though 'with the exception of the United States' is not itself a sentential operator, it is possible for us to *introduce* a sentential operator that we can use to express the same propositions we express in English by prefixing sentences with 'with the exception of the United States.' More precisely, this sort of quantificationalist will insist that there is a proposition p such that, for any domain D, there is an operator $\mathsf{At}(D)$ with the following properties:

- The proposition expressed by 'No Western country has all five of Earth's climate zones' when uttered in a context where the salient domain of quantification is D is identical to $\mathsf{At}(D)(p)$, and
- The proposition expressed by 'with the exception of the United States, no Western country has all five of Earth's climate zones' when uttered in a context where the salient domain of quantification is D is identical to $\mathsf{At}(D \setminus \{\text{United States}\})(p)$, where $D \setminus \{\text{United States}\}$ is the result of removing the United States from D.

At the very least, the logic of this operator $\mathsf{At}(D)$ will include the claim that '$\mathsf{At}(D)(p)$' entails 'it is domain-occasionally the case that p.' So, on a view of this sort, quantificationalism can be inferred from (3) and (4), even though (3) itself expresses a domain-invariable falsehood.

Just like temporalism, then, quantificationalism is, at its core, a thesis about granularity. It asserts that reality is fine-grained enough for the notion of domain-variable truth to make sense. Extremely coarse-grained pictures of reality, such as the previously considered view that propositions are individuated by material equivalence, falsify quantificationalism. Indeed, presumably logical truths are domain-invariably true, and contradictions are domain-invariably false. So, all true propositions are domain-invariably true, and all false propositions are domain-invariably false. There simply are not enough propositions around for us to introduce operator expressions that faithfully regiment discourse involving domain-manipulation.

3. Getting More Precise

Quantificationalism is only as clear as the notion of domain-occasional truth. While, I hope, a rough idea of what this unfamiliar notion is has emerged from the discussion so far, the best way to properly clarify it is to describe its logic. To that end, let me continue the analogy between temporary truth and domain-occasional truth. One approach to developing the logic of temporary truth is to connect the notion of temporary truth to the notion of truth at a time, as tense Leibnizianism does, then read off the logic of temporary truth from the logic of truth at a time.

This section outlines a logical framework for domain-occasional truth based on a similar idea. I introduce a notion of truth at a *particular* domain, connected to domain-occasional truth roughly the way that truth at a time is connected to

[16] This is the dominant approach for dealing with cases like (3) and (4), found, e.g., in Westerståhl (1985); Stanley (2002). Stanley and Szabó (2000) contains relevant discussion of contextual quantifier domain restriction, though it does not deal with explicit domain restriction by a phrase like 'with the exception of the United States.'

temporary truth. Then, I derive key principles for its logic by sketching an intended interpretation based on the notion of *metaphysical substitution* (Bacon, 2019, 2020).

The framework described in this section has some non-standard features, chief among them a restriction on λ-abstraction over variables occurring in certain "forbidden positions." This is a philosophically principled design choice that nonetheless has revisionary consequences for the metaphysics of identity in quantificationalism. I want to flag at the outset that I will not discuss either matter in detail here. The arguments are somewhat technical, so I decided to leave them to Appendix A. The interested reader can skip ahead to Appendix A after reading through the present section.

3.1. **A General Higher-Order Language.** We will work in a simply typed higher-order language. In this language, terms come categorized into *types*. There is a basic type e of singular terms, and a basic type t of formulas. Whenever σ and τ are types, we can form a new type $\sigma \to \tau$ belonging to expressions that combine with expressions of type σ to output expressions of type τ. We will fix a set of constants containing at least the usual Boolean connectives plus two infinite families of constants: a constant $\equiv_\sigma$ of type $\sigma \to \sigma \to t$ for each type σ, and a constant $\mathsf{E!}_\sigma$ of type $(\sigma \to t)$ for each type σ. We call the constants $\mathsf{E!}_\sigma$ *existence predicates*. The constant $\equiv_\sigma$ expresses the relation of *applicative indiscernibility*, that is, the relation two entities bear when they share all their properties in the absolutely unrestricted sense of 'all.' In my framework, applicative indiscernibility does not entail identity; I explain what the difference between the two is in Section A.2. To keep things simple, I ignore this complication in the main text and speak as if $\equiv$ coincides with identity.

Out of this basic vocabulary, plus an infinite stock of variables at each type, we can form complex terms by applying any one of five term formation rules, three of which are familiar and two of which are distinctly quantificational. Let me start with the familiar ones. First, we can *apply* a term M of type $\sigma \to \tau$ to a term N of type σ to obtain a term MN of type τ. Second, we can *abstract* a variable x of type σ in a term M of type τ to form a term $\lambda x.M$ of type $\sigma \to \tau$. Third, we can *quantify* a variable x of type σ in a formula P to form a formula $\forall xP$. As usual, we treat $\exists xP$ as an abbreviation for $\neg\forall x\neg P$.

Next, we have the two quantificationalist term formation rules. First, if P is a formula, we can form a new formula, $\blacksquare P$. This is how I regiment the claim 'it is domain-invariably the case that P.' So, in the present framework, quantificationalism can be stated as follows:

Quantificationalism: $\exists p(\blacklozenge p \wedge \blacklozenge\neg p)$.

(Here $\blacklozenge P$ abbreviates $\neg\blacksquare\neg P$.) Second, if F is a predicate of type $\sigma \to t$ and M is a term of type τ, we can form a new term, $\mathsf{At}(F)(M)$, also of type τ. An expression of the form $\mathsf{At}(F)$ is called a *domain specifier*.[17]

A domain specifier $\mathsf{At}(F)$ is an expression whose semantic function is to specify precisely the Fs as the domain of quantification within its scope. This is intended

[17]Officially, I allow At to take multiple arguments in its first argument place. More precisely, if $\bar{F} := F_{\sigma_1}, \ldots, F_{\sigma_n}$ is a sequence of predicates where each F_{σ_i} has type $\sigma_i \to t$ and all the σ_i's are distinct and M is a term of type τ, then $\mathsf{At}(\bar{F})(M)$ is a term of type τ. Intuitively, the domain specifier $\mathsf{At}(\bar{F})$ specifies the F_{σ_i}'s as the domain of quantification over things of type σ_i within its scope, simultaneously for all the listed σ_i's. I ignore this generalization in the main text.

to be analogous to the way that expressions like 'in 1783'—when thought of as sentential operators—force sentences within their scope to be evaluated with respect to the time 1783. In this framework, a sentence like 'with the exception of the United States, no Western country has all five of Earth's climate zones' is regimented as something like

$$\mathsf{At}(\lambda x.x \not\equiv \text{United States})(\neg\exists x(\text{WesternCountry}(x) \wedge \text{FiveClimateZones}(x))). \quad (5)$$

Henceforth, I will adopt the convention of pronouncing a domain specifier $\mathsf{At}(F)$ as 'among the Fs.' I want to stress, however, that this is really just a pronunciation convention, not a substantive hypothesis about what expressions of the form 'among the Fs' mean in English.

Domain specifiers regiment the informal notion of truth at a domain. Domains, in my theory, are identified with higher-order properties, and any such property corresponds to a domain.[18] For a proposition P to be true at the domain consisting of the Fs is simply for $\mathsf{At}(F)(P)$ to be true *simpliciter*. Domain specifiers relate to the domain-invariable truth operator $\blacksquare$ in roughly the same way tense nominals relate to the tense operator 'always.' In one direction, this means that if $\blacksquare P$ is true, then so is $\mathsf{At}(F)(P)$, for any domain specifier $\mathsf{At}(F)$. In the other direction, the intuitive idea is that if $\mathsf{At}(F)(P)$ is true for all domain specifiers, then $\blacksquare P$ is true, though some care is needed when formulating this claim precisely.[19]

I should draw attention to two non-standard features of my framework. The first concerns the distinction between *categorematic* and *syncategorematic* expressions. A categorematic expression is a typed term of the language, that is, an expression that can be iteratively built out of constants and variables through the term formation rules. A syncategorematic expression is a device for syntactically notating ways in which categorematic expressions themselves are built. In our setting, application (notated by term concatenation), λ, quantifiers, domain specifiers, and the domain-invariable truth operator are all syncategorematic. With five primitive syncategorematic expressions, my framework departs from orthodoxy in the higher-order metaphysics literature, which typically treats only application and λ as syncategorematic.

The other non-standard feature of my framework concerns λ-abstraction. In standard higher-order frameworks, a term of the form $\lambda x.M$ is well formed regardless of where the variable x occurs within M, if it occurs at all. The present framework imposes a non-standard restriction: $\lambda x.M$ is only well formed if x does not occur in a *forbidden position* in M. We say that x occurs in a forbidden position in M if it falls either inside the scope of an occurrence of the domain-invariable truth operator or within the second argument place of a domain specifier. Consequently, the following terms come out as ill-formed:

$$\lambda x.\mathsf{At}(F)(x) \qquad \lambda x.\mathsf{At}(F)(Gx) \qquad \lambda X.\mathsf{At}(F)(Xa) \qquad \lambda p.\blacksquare p.$$

Abstraction of variables within the scope of the first argument place of a domain specifier is permitted, so that, e.g., $\lambda X.\mathsf{At}(X)(M)$ is well formed. On the

[18] In treating domains of quantification as properties, rather than sets or pluralities, I am following Russo (Manuscript).

[19] The main problem is that At can take a predicate of any type in its first argument place, and our higher-order language lacks the ability to quantify over properties of any type at once. We can get around this complication by using inference rules: roughly, if we can prove $P \to \mathsf{At}(F)(Q)$ for an *arbitrary* domain specifier $\mathsf{At}(F)$, then we can also prove $P \to \blacksquare Q$.

other hand, in keeping with orthodoxy, I assume that $\beta\eta$-equivalent terms (such as $(\lambda x.M)N$ and $M[N/x]$) are synonymous and hence intersubstitutable *salva veritate* in all contexts when well formed. I do not take this to be a controversial claim, but rather a stipulation about the inferential role of the theoretical term λ.

An expression like $\lambda x.\mathsf{At}(F)(x)$, were it well formed, would be a *categorematic surrogate* of the domain specifier $\mathsf{At}(F)$: a categorematic expression that, due to the intersubstitutivity of $\beta\eta$-equivalents, could be used to *stand in* for $\mathsf{At}(F)$. Because of the restrictions I imposed on abstraction, my framework treats domain specifiers and $\blacksquare$ as *genuinely syncategorematic expressions*: syncategorematic expressions that lack categorematic surrogates. This design choice is philosophically motivated. On the interpretation I propose in Section 3.3, domain specifiers, and by extension $\blacksquare$, describe a way of putting higher-order entities together, *quantificational substitution*, that differs from and should not be reduced to application. The restrictions imposed on abstraction are what keeps these two modes of combination from collapsing.

As I remarked in the introduction to this section, showing why this is so requires somewhat technical arguments, so I decided to leave the matter for Section A.1. I want to remark again that my design choice, while a very natural consequence of my proposed interpretation of domain specifiers, is not philosophically innocent.[20] In particular, it leads to a revisionary, though ultimately illuminating, picture of the metaphysics of identity in quantificationalism, which I explore in Section A.2.

Last, I should note that the syncategorematic treatment of quantifiers has a different motivation: quantifiers are syncategorematic so that variables that cannot be abstracted can still be quantified. Without allowing this, we could not even formulate the formal statement of quantificationalism. Since quantifiers do not generate forbidden positions within their scope, we can always *construct* categorematic surrogates of syncategorematic quantifiers using λ-abstraction:

$$\hat{\forall}_\sigma := \lambda X.\forall y X y \qquad \hat{\exists}_\sigma := \lambda X.\exists y X y.$$

[20]The standard convention that λ-abstraction is unrestricted is not philosophically innocent. For example, assuming that $\lambda x.M$ is well formed even when x does not occur free in M implies that anything can be "metaphysically defined" from anything else, which contradicts certain fine-grained accounts of metaphysical structure (Dorr, 2016, §8). Similarly, if one can abstract variables in M in a different order than those variables appear in M, then every binary relation R has a converse, $\lambda yx.Rxy$. But the existence of converse relations has long been regarded as metaphysically suspicious (Williamson, 1985; Fine, 2000; Dorr, 2004).

The standard convention of having only two syncategorematic expressions, application and λ-abstraction, is not philosophically innocent either, and the choice of which expressions to treat as syncategorematic carries different weight depending on whether any restriction on λ-abstraction is imposed. There is a long tradition of treating logical vocabulary as syncategorematic, on the grounds that its semantic function is not to express particular entities but rather to describe *how* entities are combined to form propositions. Roughly, the idea is that the only constituents of, say, the proposition $P \wedge Q$ are the constituents of P and those of Q; conjunction does not contribute a constituent of its own. This idea stretches back as far as Buridan (1976, I.7.2) and is perhaps most famously championed by the early logical atomists, e.g., Wittgenstein (1921, 4.0312, 5.4611) and Russell (2010, pp. 39–40). But this approach cannot succeed when combined with unrestricted λ-abstraction, at least not if $\beta\eta$-equivalents are coextensive. If conjunction is syncategorematic but λ-abstraction is unrestricted, we can always form a categorematic surrogate of syncategorematic conjunction by abstracting within the scope of $\wedge$ to form the categorematic term $\lambda pq.p \wedge q$. Conjunction is *merely superficially* syncategorematic, whereas it could be made *genuinely syncategorematic* by restricting abstraction within its scope.

These are well-defined terms, and $\forall x^\sigma P$ is always intersubstitutable with $\hat{\forall}_\sigma \lambda x^\sigma.P$ whenever x does not occur in a forbidden position in P. These defined categorematic quantifiers are less expressive: for example, a claim of the form $\exists p\, \mathsf{At}(F)(p)$ cannot in general be rewritten as the result of applying the categorematic quantifier $\hat{\exists}_t$ to a predicate, since that would involve abstracting in a forbidden position. Still, categorematic quantifiers are useful to have around, as we shall see.

3.2. **Free Logic.** We are trying to study expressions capable of manipulating domains of quantification, so it is important to get a handle on what the logic of quantification itself must be like in order for this to be possible. The most important point to that effect is that quantifiers whose domain is capable of being so manipulated must obey a *free* rather than *classical* logic.

Classical quantificational logic validates the principle of unrestricted existential generalization:

$$Fx \to \exists x Fx. \tag{EG}$$

From the claim that Gavin Newsom is tall, I can infer that something is tall. The distinguishing feature of free logics is that they admit exceptions to this principle.

One reason to adopt a free logic is a desire to reason with empty terms. Using (EG), it is possible to derive every instance of the schema

$$\exists x(x \equiv t). \tag{6}$$

If we take t to be synonymous with the name 'Sherlock Holmes,' the corresponding instance of the above says that there is something identical to Sherlock Holmes, a claim you should resist if you think that 'Sherlock Holmes' does not refer to anything.

Another route to rejecting (EG), which is more pertinent here, comes from the idea that existence is metaphysically contingent: some things that do exist might have failed to exist. The idea of metaphysically contingent existence is ruled out by (EG) together with standard modal logic, from which one can derive:

Necessitism: Necessarily, everything necessarily exists:

$$\Box\forall x\Box\exists y(x \equiv y).$$

The modal logic required to show this is just the rule of necessitation for $\Box$: if P follows from no assumptions, then so does $\Box P$.[21]

Nothing in the underlying argument hinges on $\Box$ being metaphysical necessity. In exactly the same way, we can show that if (EG) is a logical truth and $\blacksquare$ obeys the rule of necessitation, then domain-occasional existence is impossible: anything that in fact exists domain-invariably does.

Domain-invariabilism: It is domain-invariably the case that everything is domain-invariably identical to something:

$$\blacksquare\forall x\blacksquare\exists y(x \equiv y).$$

A quantificationalist cannot accept domain-invariabilism. The very idea driving quantificationalism is that we can evaluate propositions for truth or falsity relative to non-current domains of quantification, some of which are smaller than the current

[21]The derivation is simple. From $x \equiv x$, which is a logical truth, infer $\exists y(x \equiv y)$ by (EG). Then apply necessitation, universal generalization, and necessitation again to get $\Box\forall x\Box\exists y(x \equiv y)$. For more discussion of this derivation and of necessitism, see Williamson (2013).

one. But domain-invariabilism says exactly that domains of quantification cannot shrink.

Quantificationalists, therefore, must either adopt a free logic or deny that $\blacksquare$ obeys the rule of necessitation. I will take the first option, but I should stress that I do not think the dispute over which of these two alternatives we should take is a deep one. It is a special case of the question of whether logical truths should remain true within the scope of necessity-like operators, which seems to boil down to a largely verbal matter about how to use the notion of logical truth.

To theorize in a free logic is to insist that our primitive quantifier $\forall$ does not have a classical logic, in the sense that it does not satisfy (EG) within the scope of $\blacksquare$. Doing so does not foreclose the possibility of introducing some other quantifier, distinct from $\forall$, which does have a classical logic. Views that theorize using a combination of a free quantifier and a classical quantifier are in fact quite common in the literature. *Contingentism*—the negation of necessitism—is a prominent example.[22] On pain of inconsistency, contingentists must either embrace a free logic or deny the rule of necessitation for metaphysical necessity; in practice, they tend to take the former option. Contingentists typically claim that existence is contingent not only in the sense that some things that do exist might have failed to exist, but also in the sense that there could have been things that do not in fact exist. For example, they might claim that while Wittgenstein could have had a first child, in fact there is nothing that could have been Wittgenstein's first child. When making generalizations, contingentists often help themselves to an "outer" quantifier whose logic is assumed to be classical, and whose range includes things such as Wittgenstein's merely possible first child (Fritz and Goodman, 2016).

Like the contingentist, I deny that $\forall$ itself has a classical logic, but I do make use of a distinct quantifier whose logic *is* classical. In fact, I will show that such a quantifier can be *defined* from $\forall$ using domain specifiers. What I take the discussion of domain-invariabilism above to show is merely that no quantifier whose domain can be manipulated by $\blacksquare$ and domain specifiers in the ways I shall describe can have a classical logic. Being "flexible" this way and having a classical logic are two interesting but incompatible properties of quantifiers that can be useful for different purposes. It can thus be desirable for a framework to be able to express both.

The system of free logic I shall use here is axiomatized using the *existence predicates* $\mathsf{E!}_\sigma$ I introduced earlier. (EG) says that if P, then there is something such that P. The present system replaces that with a weaker claim: if P is true *and* x *exists*, then there is something such that P:

$$Fx \to (\mathsf{E!}x \to \exists x Fx). \tag{E!EG}$$

The notions of existence expressed by these predicates are quantificational. The background logic proves that $\mathsf{E!}_\sigma$ and the predicate *is something*, $\lambda x^\sigma.\exists y(x \equiv y)$, are equivalent. Because of this, the addition of existence predicates does not correspond to an increase in ideological complexity.

A property is called *existence-entailing* when having it entails existing. I will be as liberal as possible regarding which properties are existence-entailing. Particular properties that can be logically built from existence or quantifiers are guaranteed to be existence-entailing: for example, *being an existing philosopher* ($\lambda x.\mathsf{E!}x \wedge Px$),

[22] See Williamson (2013); Fritz and Goodman (2016) for discussion. We find a similar approach among actualists, who believe that everything is actual. See especially Fine (1977, 2016).

being someone who loves themselves ($\lambda x.\exists y(Lxy \wedge x \equiv y)$) or existence itself. But the ambient logic is consistent with a wide range of hypotheses about whether properties that cannot be so built are existence-entailing. For example, it would be perfectly consistent to maintain that Sherlock Holmes is a fictional detective, or that he lives at 221B Baker Street, while denying that Sherlock Holmes exists (Bacon, 2013).

3.3. **The Interpretation of Domain Specifiers.** I said that domain specifiers function by specifying a particular property as the domain of quantification within their scope. I now elaborate on what that means.

Recall that I think of domains of quantification simply as higher-order properties. Each higher-order property F represents a domain consisting precisely of the Fs. At each type, there is a special domain, *the current domain*, represented by the property of *existing* (at that type). Part of what it means to say that a domain specifier, $\mathsf{At}(F)$, specifies the Fs as the domain of quantification within its scope can be spelled out in terms of how $\mathsf{At}(F)$ acts on the current domain (of the same type):

$$\mathsf{At}(F^{\sigma\to t})(\mathsf{E!}_\sigma) \equiv F^{\sigma\to t}. \tag{Gen}$$

In words: *to exist among the Fs just is to be F.*

It is important to draw a distinction between domain *specification* and domain *restriction*. Suppose that Sherlock Holmes is a fictional detective and yet Sherlock Holmes does not exist. Then (Gen) above predicts that Sherlock Holmes *exists among fictional detectives*, since this reduces to saying that Sherlock Holmes is a fictional detective. The domain specifier 'among fictional detectives' can thus shift the domain of quantification to one that includes "new" things that did not previously exist, so to speak.[23] Domain specification is more general than domain restriction, as the latter can be defined from the former: to restrict the domain to just the Fs is to specify the *existing* F*s* as the new domain.[24]

From (Gen), we can derive further principles describing how domain specifiers interact with quantifiers themselves. If to exist among the Fs is simply to be an F, then the result of combining $\mathsf{At}(F)$ with the categorematic quantifier $\hat{\forall}_\sigma$ (where $\sigma \to t$ is the type of F) should be a quantifier that ranges precisely over the Fs. Writing $\hat{\forall}^F$ to abbreviate $\mathsf{At}(F)(\hat{\forall})$, this means we should be able to derive:

$$\hat{\forall}^F x\, Fx \tag{7}$$

$$\hat{\forall}^F x\, Gx \to (Fy \to Gy) \tag{8}$$

Here and below, I write $\hat{\forall}^F x\, P$ for $\hat{\forall}^F(\lambda x.P)$, and likewise for $\hat{\exists}^F$. As with $\hat{\forall}$ itself, such terms are well formed only when x does not occur in a forbidden position in P. Both these principles are indeed derivable from (Gen) in my system.[25]

---

[23] More generally, this domain expansion effect can happen with any domain specifier $\mathsf{At}(F)$ where F is not existence-entailing.

[24] Formally: $\mathsf{At}^\cap(F) := \mathsf{At}(\lambda x.Fx \wedge \mathsf{E!}x)$.

[25] We can also derive a schema to the effect that the quantifier $\hat{\forall}^F$ is *normal*, i.e., closed under modus ponens:

$$\hat{\forall}^F x\,(Gx \to Hx) \to \hat{\forall}^F x\, Gx \to \hat{\forall}^F x\, Hx. \tag{9}$$

These three principles, together, capture the standard inferential role of a free quantifier that ranges over precisely the Fs. (Gen) also constrains the behavior of syncategorematic quantifiers binding in forbidden positions, though in a different way. When x occurs in a forbidden position

Again, since F need not be existence-entailing, the range of the quantifier $\hat{\forall}^F$ need not be a *restriction* of that of $\forall$. Quantifiers constructed via domain specifiers can have larger, smaller, or incomparable domains of quantification. In particular, as anticipated earlier, even though $\forall$ itself does not have a classical quantificational logic, we can *define* a quantifier that does by specifying a universal property as the new domain of quantification. By a universal property I do not merely mean a property that everything in the domain has, but rather a property that it is inconsistent to lack. The universal property *being such that* $\top$ (at a given type) can always be defined in purely logical terms, and I will use this property to introduce my classical quantifiers:

$$\hat{\forall}^{\top} := \mathsf{At}(\lambda x.\top)(\hat{\forall}), \qquad \hat{\exists}^{\top} := \mathsf{At}(\lambda x.\top)(\hat{\exists}).$$

I will henceforth adopt the convention of writing English quantifiers in boldface, as in **everything**, when I intend the classical reading.

Intuitively, our primitive free quantifiers behave like syntactic restrictions of the defined classical quantifiers just introduced. For example, it could be that although there are no fictional detectives, **there are** fictional detectives, because Sherlock Holmes is a fictional detective. To say that there are no fictional detectives just is to say that **there are** no fictional detectives who *exist*, in the sense of $\mathsf{E!}$. The picture is similar to that put forward by contingentists or actualists who theorize with a pair of quantifiers, one ranging over the "inner" domain of things that in fact exist and one over the "outer" domain of merely possible entities.[26] In fact, a quantifier $\hat{\forall}^F$ always behaves like the syntactic restriction of the defined classical quantifier by the predicate F. In other words, to be in the range of the quantifier $\hat{\forall}^F$ is simply to **exist** and to be F.

This settles how domain specifiers handle our primitive quantificational idioms. But (Gen) alone does not settle how domain specifiers should interact with arbitrarily complex terms. Consider the property of *being a lover*, in the sense of *loving someone*: $\mathsf{Lover} := \lambda x.\exists y\, Lxy$. What is it to be a lover *among the Fs*? Intuitively, it is to love **someone** *who is* F: once the domain is reset to the Fs, the quantifier 'someone' implicit in being a lover ranges over the Fs only. (Gen) alone does not deliver this result. It tells us what domain specifiers do to existence predicates, and derivatively to quantifiers, but not what they do to entities, like being a lover, that involve quantification only indirectly. To understand that, we can interpret domain specifiers in terms of *metaphysical substitutions* (Bacon, 2019, 2020).

This notion is easiest to introduce by pretending, for a moment, that reality is structured the way language is. On this pretense, the proposition that Alice loves someone is a complex entity built by applying the quantifier $\hat{\exists}$ to the property of being loved by Alice, La, which is in turn built by applying the relation L to Alice. If reality were structured this way, there would be a well-defined operation of *substituting* one constituent of a proposition (or property, or relation) for another. Substituting Bob for Alice in the proposition that Alice loves someone yields the

in P, we can derive $\mathsf{At}(F)(P[a/x])$ from $\mathsf{At}(F)(\forall x P)$ and $\mathsf{At}(F)(\mathsf{E!}a)$, but not in general from $\mathsf{At}(F)(\forall x P)$ and Fa.

[26] There are thorny philosophical questions regarding the interpretation of inner and outer quantifiers having to do with ontological commitment, the substantivity of metaphysical disputes, and the intelligibility of primitive higher-order quantification, both in the contingentist/actualist setting and in the quantificationalist setting. While I cannot satisfactorily address all of them here, I will briefly discuss some in Section 5.

proposition that Bob loves someone; substituting hate for love yields the proposition that Alice hates someone.

A domain specifier, on this picture, is understood as performing a special kind of substitution: $\mathsf{At}(F)$ replaces each quantificational constituent with its counterpart that ranges over the Fs in accordance with (Gen)—the quantifier $\hat{\exists}$ with $\hat{\exists}^F$, the existence predicate $\mathsf{E!}$ with F—while fixing every other constituent. The proposition that among the Fs, Alice loves someone, then, equals the proposition that Alice loves **some** F (Figure 1). By the same token, to be a lover among the Fs is to love **some** F, as desired.

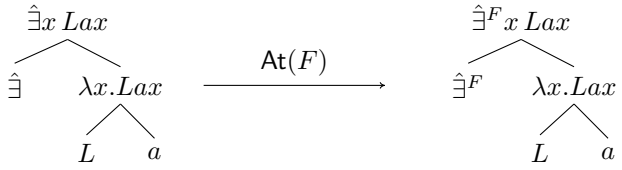


FIGURE 1. Domain specification as substitution: $\mathsf{At}(F)$ replaces the quantifier $\hat{\exists}$ with $\hat{\exists}^F$ and fixes every other constituent.

The structured picture is suggestive, but it should not be taken literally. If reality really were structured like language, we would expect that whenever FA is identical to GB, F is identical to G and A to B. But that conflicts with the identification of $\beta\eta$-equivalents: $\neg\neg p$ is $\beta\eta$-equivalent to $(\lambda q.\neg\neg q)p$, yet negation is not identical to double negation. Worse, the *Russell–Myhill paradox* shows that the principle is inconsistent in classical higher-order logic. The argument carries over to free logics where a classical quantifier is definable, such as the present one.[27]

The inconsistency of the quasi-syntactic picture does not imply that the notion of a metaphysical substitution is unintelligible. Instead of defining it in terms of notions like constituency supplied by the quasi-syntactic picture, we can take it as primitive. Its defining feature is that it respects the way that higher-order entities are combined to form new ones: a metaphysical substitution is a kind of transformation on higher-order entities of all types that *commutes with application*, in the sense that performing it on MN yields the same result as first performing it on M and N separately, then applying the results. This idea is intelligible independently of the assumption that reality has quasi-syntactic structure.

I propose, then, to interpret domain specifiers as follows. Each property F determines a metaphysical substitution that "moves" interpretations of existence predicates and quantifiers in accordance with (Gen) and the quantifier principles

[27]This argument is originally due to Russell (1903, Appendix B) and Myhill (1958) and recently revived, e.g., in Uzquiano (2015, 2023); Dorr (2016); Goodman (2017); Fritz et al. (2021); Bacon (2023). Here is a sketch. Given an arbitrary proposition p, consider the property of being *heteropredicative* (with respect to p): this is a property that a proposition q has when every property it predicates of p is a property that q itself lacks:

$$\mathsf{Het} := \lambda q.\forall X(q \equiv Xp \to \neg Xq).$$

We can prove in classical higher-order logic that the proposition $\mathsf{Het}p$ is not heteropredicative: there is some property of propositions Y that $\mathsf{Het}p$ predicates of p that $\mathsf{Het}p$ itself has. This is a counter-example to the structural principle mentioned at the beginning of the paragraph: $\mathsf{Het}p \equiv Yp$, yet $Y(\mathsf{Het}p)$ and $\neg\mathsf{Het}(\mathsf{Het}p)$, showing Het and Y must be different properties.

above, and that fixes everything that has nothing to do with quantification. I will call these substitutions *quantificational substitutions.* How quantificational substitutions act on arbitrary entities is entirely fixed by these facts and the requirement that they commute with application. A term of the form $\mathsf{At}(F)(M)$ expresses the result of executing the quantificational substitution determined by F on M.

This idea of an entity that has nothing to do with quantification is important and will come up repeatedly. I call such entities *stable.* We have seen many examples of properties that are not stable. Some involve quantification directly: existence properties and quantifiers themselves. Others do so less directly, like the property of loving someone. Others still indirectly involve quantification even though we do not express them using overtly quantificational language. For example, the term 'to eradicate' does not contain overt occurrences of quantifiers, but its meaning involves quantification: to *eradicate* a disease is to eliminate *all* (or perhaps *nearly all*) instances of it (cf. Lewis, 1986, p. 6). Stable entities are entities that do not involve quantification even in such indirect ways, and should therefore be fixed by all domain specifiers: whenever M is stable, the identity $\mathsf{At}(F)(M) \equiv M$ always holds.

I consider *combinators* and Boolean operations to be paradigmatic examples of stable entities. A combinator is the denotation of a closed term that contains no constants, for example $\lambda p.p$ (truth) or $\lambda X^{\sigma\to\tau} Y^{\tau\to\rho} z^{\sigma}.Y(Xz)$. We can describe the behavior of these entities without appealing to quantification at all. To compose X with Y is just to apply Y after X; the conjunction of P and Q is true precisely when P is true and Q is true.[28] I will also assume that the higher-order identity relations $\equiv_\sigma$ are stable.[29]

Whether any other entities are stable is a question that the framework leaves open. An attractive idea is that individuals (entities of type e) should always be stable; they simply are not the sort of thing that can involve quantification. Another interesting thought is that *fundamental* entities should be stable. This is especially fitting in the case of propositions if we think of fundamentality in terms of *grounding.* Grounding theorists typically regard universal and existential generalizations as grounded in their instances and therefore non-fundamental, so at least those kinds of non-stable propositions cannot be fundamental. I will not try to settle these matters here.

---

[28]You might think that Boolean operations actually do involve quantification indirectly. For example, you might think that negation is the property of materially implying *everything*, or that disjunction is the relation two propositions bear when *one of them* is true. If so, you will think Boolean operations are not stable, but you will also think that their logic is non-classical. For example, if neither P nor its negation exist, then $P \vee \neg P$ will be false. These glosses on Boolean operations are extensionally adequate if we replace the quantificational phrases above with their boldface classical counterparts, and intensionally adequate if we assume that our defined classical quantifiers are stable. As I note later, the latter claim is a theorem of my system, a consequence of the claim that Boolean operations, and therefore $\top$, are stable.

[29]This might seem more contentious. In higher-order metaphysics, it is common to define identity as the relation two entities bear when they share *all* their higher-order properties. In a free logical setting, we need to be careful with a definition of this sort: since universal generalizations fail to imply all their instances, it might be that two entities share all their higher-order properties and yet one is F while the other one is not, because F does not exist. This problem, of course, can be circumvented in our system by replacing the free quantifier with the defined classical one: identity is the relation two entities bear when they share **all** their higher-order properties. And the defined classical quantifiers are provably stable.

So much for stability. Once we embrace the substitutional picture, a number of further principles about domain specifiers naturally flow from it. Of course, domain specifiers should commute with application, since quantificational substitutions do:

$$\mathsf{At}(F)(MN) \equiv \mathsf{At}(F)(M)\mathsf{At}(F)(N). \tag{App}$$

In quantificationalist terms: to specify the domain within MN is the same as to specify it within M and within N, then apply the former to the latter. Application contributes no quantification of its own, so whatever quantification MN involves must come from M or N.

Next, we have two principles describing *closure conditions* on substitutions. Substitutions, it is natural to think, form a rich family. First, there is always an *identity* (or trivial) substitution, which moves nothing. This idea has a natural quantificationalist interpretation. If the current domain of quantification is given by *the existing things*, there should be a quantificational substitution that specifies *the existing things* as the domain of quantification within its argument. That should not change the domain of quantification at all: to be F among the existing things just is to be F. More generally, we should have the schema

$$\mathsf{At}(\mathsf{E!})(M) \equiv M, \tag{Id}$$

where M can have any type.

Second, given any two substitutions, there is always another substitution which *composes* the two, corresponding to applying one after the other. This idea, too, generalizes to quantificational substitutions. Every quantificational substitution is determined by a property. Thus, if we say that quantificational substitutions should be closed under composition, we must be able to find a property that determines the composition of the quantificational substitution determined by F with the one determined by G. The property *being G among the Fs* is the most natural choice. In other words, to specify the Fs and then the Gs as the domain one after the other is the same as to specify *the Gs among the Fs* as the domain in one step. Formally, when F and G have the same type,[30]

$$\mathsf{At}(F)(\mathsf{At}(G)(M)) \equiv \mathsf{At}(\mathsf{At}(F)(G))(M). \tag{Comp}$$

This principle helps us resolve a number of questions concerning nested domain specifiers. Here is one example. Naively, you might have thought that specifying the same domain "twice over" is redundant: something is true among the Fs just in case among the Fs, it is true among the Fs. (Comp) explains why this is not so: whenever F is not stable, this redundancy idea can fail, because to be F need not be the same as to be F among the Fs. To illustrate, consider again the property of being a lover (*loving someone*). Suppose Alice loves Bob and Bob loves only Carol, but Carol does not love anyone. Then Alice and Bob are both lovers, yet only Alice is a *lover among lovers*. So, it is true that Bob exists among lovers, but false that, among lovers, he exists among lovers.

Another consequence of (Comp) is that whatever you get by specifying a stable property as the domain of quantification must itself be stable. In particular, as I anticipated earlier, specifying the Fs as the domain within $\mathsf{At}(\lambda x.\top)(\hat{\forall})$ amounts to

[30] To state the correct composition principle for cases where F and G have different types, we need the more general conception of domain specifiers mentioned in Section 3.1, where a sequence of predicates of different types can occur in the first argument place of At. See Appendix B for details.

specifying $\mathsf{At}(F)(\lambda x.\top)$, which is just $\lambda x.\top$, as the domain within $\hat{\forall}$. So, our defined categorematic classical quantifiers must be stable. This is an important feature of the framework. Domain specifiers do not affect all quantifiers equally. There is one collection of quantifiers, our primitive free quantifiers, whose theoretical role is, in part, to be maximally *flexible*: their domain can be specified arbitrarily. When I informally say that an entity "involves quantification," what I mean is that it involves the kind of quantification that the primitive free quantifiers are doing. Other quantifiers, be they definable from the primitive free quantifiers or not, may have different theoretical roles. Classical quantifiers, in particular, are intended to be maximally *general*, rather than maximally flexible, and to remain so in every context. Their stability secures this: because the defined classical quantifiers are stable, they satisfy classical quantificational logic within the scope of arbitrary domain specifiers.

It is instructive to compare the principles (Gen), (Id) and (Comp) with some principles governing what Prior (1968) calls the *redundancy* interpretation of the present tense. Roughly, this interpretation of the present tense holds that *always, for P to be true just is for P to be true at the present time*. Here 'the present time' *non-rigidly* picks out whatever the present time is at the moment of evaluation rather than a fixed instant.

The claim

(10) For P to be true just is for P to be true at the present time

is a temporal analog of (Id). And just as the former always holds, (Id) holds relative to **every** domain. The redundancy interpretation also seems to predict that

(11) The present time at t is t,

which is an analog of (Gen). Finally, it also seems to predict that both of the following are true:

(12) For P to be true at the present time at t just is for P to be true at t;
(13) For P to be true 10 minutes after the present time at t just is for P to be true 10 minutes after t.

Both can be seen as tense-logical analogs of instances of (Comp).

With these principles in place, we can derive the original verdict that to love someone among the Fs just is to love **someone** who is F if we also assume that the *loving* relation is stable. Likewise, we can derive that the proposition that among the Fs, Alice loves someone is the same as the proposition that Alice loves **someone** who is F, given the additional assumption that Alice, too, is stable. Without these stability assumptions, matters are more complicated, though that is to be expected. What results we get depends on how, exactly, we take the *loving* relation and Alice herself to involve quantification.

3.4. **Granularity.** One last feature of the framework I want to emphasize before we move on to applications is its stance on how fine-grained reality is. I initially illustrated the concept of a quantificational substitution by leaning on the quasi-syntactic picture of metaphysical structure, only to later insist that the concept can be understood in its own right without defining it in terms of quasi-syntactic notions like constituency.

A more precise and forceful way of making this point is to show that the logic of domain specification can be developed in the context of a fairly *coarse-grained*

picture of reality, where logical equivalence suffices for identity.[31] Accordingly, the system I shall work with is closed under the rule

$$\text{if } \vdash P \leftrightarrow Q, \text{ then } \vdash \lambda\bar{x}.P \equiv \lambda\bar{x}.Q. \tag{Equiv$^+$}$$

An instance of (Equiv$^+$) can bind any sequence of variables $\bar{x}$ none of which occurs in a forbidden position in the body, including the empty sequence.

I call my system $\mathtt{QFC}_\dagger$, short for *quantificationalist free classicism*. A full axiomatization can be found in Appendix B. As the name suggests, $\mathtt{QFC}_\dagger$ is an expansion of the system of *free classicism* from Bacon (2024b), which can be axiomatized by adding (Equiv$^+$) to standard higher-order free logic.

The presence of (Equiv$^+$) should make clear that we are not smuggling any semblance of quasi-syntactic structure into the system. That said, $\mathtt{QFC}_\dagger$ does draw distinctions that free classicism does not. There are propositions that are neither provably equivalent nor provably inequivalent in standard free logic, but which become distinguishable in $\mathtt{QFC}_\dagger$ via domain specifiers. Two simple examples are $\top$ and the proposition that there is a truth, $\exists p\, p$. These propositions are neither provably equivalent nor provably inequivalent in standard higher-order free logic, since their material equivalence depends on whether the domain at type t contains a truth. But in $\mathtt{QFC}_\dagger$ they are provably distinct. We can use domain specifiers to shift freely from domains relative to which $\exists p\, p$ is true to domains where it is false, whereas $\top$, being stable, is domain-invariably true. Such additional refinements of metaphysical structure are nowhere near as problematic as quasi-syntactic principles: as I state in Theorem B.6 and prove in Cleani (Manuscript), $\mathtt{QFC}_\dagger$ is consistent.

4. Metaphysical Predicativity

I do not have a knock-down argument for quantificationalism, and I do not really know what such an argument would even look like. As I see it, the case for quantificationalism must be abductive: it should be accepted, if at all, because it provides a fruitful framework for theorizing about a range of philosophical topics, provided the benefits outweigh the costs. What I want to do in this section and the next is to try to convince the reader that there is a research program with legs here, which once carried out can serve as the basis for an abductive argument in favor of quantificationalism.

I begin by outlining an ambitious application of quantificationalism to the notion of *predicativity*. I propose a particular notion of predicativity, *metaphysical predicativity*, on behalf of the quantificationalist. It is a *worldly* notion: propositions, properties and relations, as opposed to linguistic entities like definitions or instances of axiom schemata, are said to be metaphysically predicative or impredicative, and these claims are made in the object language. Metaphysical predicativity earns its keep by helping articulate novel solutions to a number of paradoxes—or, more carefully, solutions that have been gestured at for a long time but arguably never precisely articulated. Here, I discuss solutions of this sort to Prior's paradox (Prior, 1961) and a paradox involving the relationship between grounding and quantification (Fine, 2010; Krämer, 2013), but I expect the approach to generalize to the

[31] The rule (Equiv$^+$) is not forced upon the quantificationalist. There are interesting finer-grained logics for quantificationalism that do not derive it. However, without (Equiv$^+$), some of the quantifier principles I rely on in this paper, such as (8) and (9), would have to be put in by hand or derived from seemingly more grain-neutral alternative inference rules that we could replace (Equiv$^+$) with. I lack the space to explore the range of options here.

Russell–Myhill paradox (briefly discussed in Section 3.3) and the class-theoretic paradoxes.

I want to emphasize up front that I am only offering a proof of concept. While I will go into some detail, a comprehensive treatment of metaphysical predicativity is beyond the scope of this paper.

4.1. **Impredicative Definitions.** A definition is called *impredicative* when it defines an entity by means of quantification over a domain that includes the very entity being defined. Otherwise, it is called *predicative*. For concreteness, consider:

(14) The set of *even numbers* is the set containing all and only those natural numbers n such that $n = m + m$ for some natural number $m \leq n$.

(15) To be a natural number is to possess every property had by 0 and closed under the successor function.

(16) The *Russell set* is the set containing all and only the sets that do not contain themselves.

(17) The *truth-teller* is the proposition that there is a true proposition.

Intuitively, (14) is a predicative definition. It defines the set of even numbers by quantifying over natural numbers, but the set of even numbers is not itself a natural number. On the other hand, (15) is impredicative. It defines the property of being a natural number by quantifying over all properties had by 0 and closed under the successor function. But being a natural number is itself one such property. (16) is also impredicative. It defines the Russell set by quantifying over all sets that do not contain themselves. But the Russell set itself is (and, paradoxically, is not) such a set. Likewise, (17) is impredicative, as it defines a proposition by quantifying over all true propositions, including itself.

The notion of predicativity originates in Russell (1906a), influenced by the work of Poincaré (1906).[32] Russell is led to the notion of predicativity by the idea that the class-theoretic paradoxes arise because of viciously circular quantification, and that the proper solution to the paradoxes is to prohibit such quantification. This prohibition, now known as the *vicious circle principle* (VCP), is articulated by Russell in different ways. Two explicitly mention definitions.

(VC1) If, provided a certain collection had a total, it would have members only definable in terms of that total, then the said collection has no total. (Whitehead and Russell, 1910, p. 37)

(VC2) No totality can contain members defined in terms of itself. (Russell, 1908, p. 237)

These principles, while far from transparent, have long been interpreted by logicians as asserting a ban on impredicative definitions. On this interpretation, for example, the argument underlying Russell's paradox is fallacious because the impredicative definition of the Russell set is illegitimate.

The distinction between predicative and impredicative definitions, as just now glossed, feels intuitive enough. But making it precise is not straightforward. Different authors have done so in different ways, and there is no single well-established notion of an impredicative definition in the literature.[33] I will begin by offering my own precisification of the distinction.

[32] Russell (1906b) had already introduced the term 'predicative,' but with a different meaning that did not involve vicious circularity.

[33] Feferman (2005) is the canonical survey.

My approach is to take the intuitive gloss of an impredicative definition, that of a definition that quantifies over a domain that includes the entity being defined, as the starting point, and to make precise sense of what it means for a linguistic expression to quantify over a domain that includes a given entity. This requires some care. Naively, we might propose to identify the domains over which a definition quantifies with the values that the variables bound by the quantifiers occurring within it can take. That is, a definition quantifies over a domain D (that includes a) iff it contains a quantifier binding variables that can take any value from D (and D contains a). Call this the *naive proposal*.

The naive proposal is too crude. The main problem is that in a language with unrestricted quantifiers whose logic is classical, any definition of an entity of type σ that contains quantifiers binding variables of type σ counts as impredicative. But the distinction between predicative and impredicative definitions makes intuitive sense in such languages. For concreteness, consider (14). It is intuitively a predicative definition of the set of even numbers. Yet when regimented in a first-order language for arithmetic and set theory, it will contain occurrences of the unrestricted first-order quantifiers 'all' and 'some.' In the usual interpretation, variables bound by these quantifiers can take any object as their value, including the set of even numbers. On the naive proposal, (14) is just as impredicative as (16).

Building on a suggestion of Chihara (1973, p. 6), I suggest we revise the naive proposal by considering not just what the quantifiers in fact range over, but also how what a definition picks out changes as we vary what the quantifiers range over.[34] Impredicative definitions quantify over a domain that includes the entity being defined in the sense that they are *quantificationally sensitive* to their definiendum: *what* they define can change depending on whether we add the definiendum to, or remove it from, the domain of quantification.

To illustrate, consider (15). Suppose we force the second-order quantifiers in the definiens to range only over the domain of properties whose extensions are finite initial segments of the natural numbers. This domain does not include any hereditary property, so (15) defines a property that vacuously applies to everything. Once we add the property of being a natural number back to the domain, on the other hand, the definiens picks out a property that does not apply to everything.

Or consider (16). Russell's paradox arises from the crucial assumption that the Russell set itself falls within the range of the quantifiers occurring in its definition.

---

[34] Chihara discusses the proof of Cantor's theorem, which shows that there is no surjection f from a set A to its powerset $\wp(A)$. The proof is by contradiction. We assume there is such a surjection f, which in set theory is represented by a set of ordered pairs. We then define a set B via the following membership conditions:

$$\forall x(x \in B \leftrightarrow x \in A \wedge \forall y((x, y) \in f \rightarrow x \notin y))$$

Chihara observes:

> Let us suppose that in our specification of B, 'x' and 'y' are restricted to relatively small totalities, viz. A and $\wp(A)$ respectively (clearly, we cannot restrict these variables to smaller ranges without specifying a different set from the one intended). Even then, the specification would be impredicative since B would be a member of the range of 'y'.

The crucial observation here is that there is a domain, namely the union of A with $\wp(A)$, such that if we let the quantifiers in our definition range over that domain, the definition picks out its intended definiendum. But if we remove the intended definiendum from that domain, the definition ends up picking out a different set.

Without this assumption, the argument cannot be run. And clearly, there are lots of choices of domains not including the Russell set, such that if we forced the quantifiers in the definiens to range over these domains, we would get something other than the Russell set as the result of the definition.

Analogously for (17). If we force the quantifiers in the definiens to range over, say, the actual truths of biology, then (17) defines a proposition that is false at a world where the truths of biology are wholly disjoint from the actual truths of biology. Once we add the proposition that there is a truth back into the domain, however, the definiens picks out a proposition that is true at such a world.

Paradigmatically predicative definitions, by contrast, do not behave this way. While (14) does define different sets depending on what we take the quantifiers to range over, adding or removing the set of even numbers itself from the domain of quantification cannot change what (14) defines, no matter what we take the domain to be. Sets are individuated extensionally, and since the set of even numbers is not a natural number, the membership conditions specified by (14) are equivalent to those we get by adding or removing it from the domain of quantification.

4.2. **Beyond Syntax.** The notion of predicativity I have been discussing is a property of linguistic entities, namely of definitions.[35] However, predicativity was not originally meant to be a purely linguistic notion. Following Russell, predicativists hold that some (or perhaps all) impredicative definitions are illegitimate because they exhibit a kind of vicious circularity. But this circularity is not purely a matter of linguistic form. Rather, it is the very entity being defined by an impredicative definition—or the class over which the impredicative definition quantifies—that, if it existed, would be involved in a vicious circle of some sort. There is, in other words, a *metaphysical* notion of predicativity underwriting the linguistic one. And we would be better served by theorizing directly in terms of the metaphysical notion instead of taking a detour through its linguistic approximation. This section elaborates on these points.

Why are impredicative definitions illegitimate? (VC1) and (VC2) appear to make ontological claims: there is nothing corresponding to an impredicative definition (VC1), or there is nothing that can only be defined impredicatively (VC2). An impredicative definition, then, is illegitimate because it either always fails to pick out anything, or it is at risk of doing so—because it might be attempting to define something that can only be defined impredicatively.

Why is that the case? If we think that at least *some* impredicative definitions are non-empty, the problem does not seem to lie in the grammar of impredicative definitions themselves. Rather, predicativists suggest that there would be something problematic in the *metaphysics* of the entities that only admit impredicative definitions, if there were any, or in the classes quantification over which can be used to state impredicative definitions. The definition of being a natural number as having every hereditary property is viciously circular, because it defines a property by quantifying over a domain that includes that very property. But if that is the *only* definition of being a natural number, that reveals a way in which

[35]Philosophers working with the Aristotelian concept of a *real definition* (e.g., Fine, 1994) might think there is a sense in which a definition is not merely a linguistic entity but rather some kind of worldly entity that specifies the nature of something. How predicativity interacts with the notion of a real definition is an intriguing question, but one I will not pursue here.

the property of being a natural number *itself*—or perhaps the class of hereditary properties—exhibits a kind of vicious circularity.

Something like this picture is suggested more strongly by other formulations of the VCP proposed by Russell, which do not mention definability at all. In particular,

(VC3) Whatever in any way *concerns all or any or some* [...] of a class must not be itself one of the members of a class (Russell, 1973, p. 209).

In (VC3), the relation *concerning all or any or some* is borne by an object to a class. No explicit reference to language is made. For brevity, let us call this relation *quantificationally concerning*. The VCP, as articulated by this statement, asserts that this relation cannot enter (presumably vicious) cycles of a particular form: no entity can quantificationally concern a class (set with a total) it is a member of.

In sum, impredicative definitions are illegitimate because, if I can only define an entity by means of quantification over a domain that includes that very entity, this is evidence that the entity itself quantificationally concerns a class that contains that very entity, and *that* is what is metaphysically problematic. The VCP is ultimately a metaphysical thesis, imposing the absence of such vicious circles of quantificational concern as a necessary condition on existence.[36]

I am suggesting that we can view predicative definability as an attempt to linguistically characterize the worldly property of not quantificationally concerning a class that contains oneself. I will now argue that a predicativist who thinks everything has this property should not regard predicative definability as a good approximation of it. For what counts as predicatively definable is language-dependent in a way that existence is not.

For example, whether the number 1 is predicatively definable depends on whether we work in a signature equipped with a constant symbol for 0 or not. If we do, we can define 1 predicatively as the successor of 0. If we do not, it is most natural to define 1 as the smallest natural number greater than 0, or as the smallest natural number with a predecessor, both of which are arguably impredicative definitions. Yet either 1 exists or it does not, regardless of the language we use to talk about it. More generally, predicative definability is best regarded as a relation rather than a unary property: an entity is predicatively definable *relative to* a language (or perhaps *given* some other entities the existence of which we assume to be unproblematic, see Feferman, 1964).

We can introduce unary versions of predicative definability following familiar blueprints. We could say that an entity is predicatively definable *simpliciter* if it is predicatively definable in *some* language. This would overgenerate: anything is predicatively definable in a language equipped with a constant symbol standing for it. Nor would it do to say that something is predicatively definable *simpliciter* if it is predicatively definable in *every* language: the existence of expressively impoverished languages would make this notion far too demanding. For the predicativist who wants to use predicative definability as a necessary condition on existence, perhaps the most promising approach is to say that something is predicatively

[36] Gödel (1944, p. 135), commenting on (VC3), distinguishes a description of an entity involving all of a class from the entity itself involving all of a class, supporting the present reading. Skolem (1967, p. 297) writes that what seems puzzling about impredicative definitions of sets is that they presuppose that there are "sets *whose existence* depends upon all sets" (emphasis added).

definable *simpliciter* if it is predicatively definable in some metaphysically distinguished language. That, of course, invites a host of thorny questions about what makes a language metaphysically distinguished in the target sense.

A cleaner approach, if it can be made to work, is to avoid the detour through language altogether. The sort of predicativist I am envisioning wants to assert that no entity quantificationally concerns a class that contains that very entity. The thesis that everything must be predicatively definable is a natural linguistic proxy for this claim, but ultimately falls short of capturing it. I propose, instead, to bypass the notion of predicative definability entirely, and make the thought precise by giving a purely non-linguistic account of what it means for an entity to quantificationally concern a class that contains that very entity. This account consists in a non-linguistic generalization of the notion of quantificational sensitivity introduced in Section 4.1.

It should be emphasized that, if we do not assume quantificationalism, this is far from a trivial task. In the linguistic case, the notion of sensitivity is characterized in terms of how manipulating the ranges of the quantifiers occurring in the definition can change what the latter defines. But the notions of a quantifier, of an occurrence of a quantifier, and of the range of a quantifier are all linguistic notions that do not have straightforward worldly analogs. We could, of course, find worldly analogs of these linguistic notions by relying on the quasi-syntactic picture of metaphysical structure we discussed in Section 3.3. But we already know that paradox threatens. We thus have reasons to explore approaches to de-syntactifying the notion of quantificational sensitivity that do not rely on the quasi-syntactic picture.

Quantificationalism has the resources to do just that. Roughly, we define an entity M to be *quantificationally sensitive* to another entity N when, at some domain, M is not equivalent to the result of adding N to, or removing N from, the domain of quantification within M. Here, 'equivalent' is a placeholder. Any equivalence relation, from coextensiveness to identity, can be plugged in, yielding different notions of quantificational sensitivity.

The operation of adding or removing something from the domain of quantification is made sense of through domain specifiers. First, using standard higher-order resources alone, we can define, for any property F, a property $F^{\pm N}$ that differs from F only in whether or not it applies to a given entity N. When F does apply to N, $F^{\pm N}$ does not, and vice versa.[37] We then regiment the metaphysical version of quantificational sensitivity as follows:[38]

$$\mathsf{Sens}_{\sim} MN^{\sigma} := \hat{\exists}^{\top} X^{\sigma \to t}(\mathsf{At}(X)(M) \not\sim \mathsf{At}(X^{\pm N})(M)).$$

[37] Formally,

$$(\cdot)^{\pm(\cdot)} := \lambda Xyz.(Xy \to (Xz \wedge z \not\equiv y)) \wedge (\neg Xy \to (Xz \vee z \equiv y)).$$

[38] It is important to distinguish between the definition given below and the following:

$$\mathsf{Sens}^{*}_{\sim} MN^{\sigma} := \blacklozenge(M \not\sim \mathsf{At}(\mathsf{E!}^{\pm N})(M)).$$

Intuitively, the definition in the main text says that for some domain F, the result of setting F as the domain of quantification within M is not equivalent to the result of setting F plus or minus N as the domain of quantification within M. On the other hand, the definition just given says that, domain-occasionally, M is not equivalent to the result of adding or removing N from the *current* domain of quantification within M. The difference is that, at a domain F, the latter adds or removes not N itself but $\mathsf{At}(F)(N)$, by (Comp) and (App). The two come apart when N is not stable.

Our official definition of metaphysical predicativity is: M is *metaphysically predicative* when it is quantificationally insensitive to itself, i.e., at all domains of the relevant type, M is equivalent to the result of adding or removing M itself from the domain of quantification:

$$\mathsf{Pred}_{\sim}M := \neg\mathsf{Sens}_{\sim}MM.$$

Again, $\sim$ here is a placeholder for any equivalence relation. We say that M is *metaphysically impredicative* when it is not metaphysically predicative.

I want to emphasize a crucial distinction between my approach and the way predicativity is generally used in existing predicativist theories. Predicativity *shapes* these theories, but generally does not itself feature in them. For example, Russell's ramified type theory is purposefully designed to make the kind of quantification involved in impredicative definitions meaningless (Russell, 1908, p. 226).[39] The idea that impredicative definitions are illegitimate shapes how he designs the type hierarchy, but the theory itself cannot express the distinction between predicative and impredicative definitions. Later predicativists, such as Weyl (1994), work in non-ramified settings and restrict the acceptable instances of comprehension to those given by formulas that are predicative in a syntactic sense (for Weyl, roughly, formulas that do not contain quantifiers ranging over sets). Again, predicativity is used to shape the theory without directly featuring in it. My approach differs precisely in this respect. Metaphysical predicativity is expressible in the object language, making predicativist constraints on existence expressible within the predicativist's own theory.

To get a feel for how this all works, let us explore the notion of metaphysical predicativity that we get when we think of equivalence as coextensiveness. Take the property of being a natural number. In our higher-order language, we can regiment it in a way that matches its definition given in (15), as follows:

$$\mathsf{Nat} := \lambda x.\forall Y(\mathsf{Hereditary}Y \to Yx);$$
$$\mathsf{Hereditary} := \lambda X.X0 \wedge \forall y(Xy \to X(sy)),$$

where 0 and s are stable constants standing respectively for 0 and the successor function. Our background logic proves that Nat is metaphysically impredicative in the presence of modest assumptions: for example, adequate arithmetical axioms guaranteeing that Nat forms a natural number structure (roughly, something isomorphic to a standard model of Peano arithmetic) plus the assumption that something is not a natural number—for concreteness, Julius Caesar. Under these assumptions, we can prove that Nat still does not apply to Julius Caesar relative to the domain of properties identical to Nat. However, if we remove Nat from this domain, the result is the empty domain, relative to which Nat vacuously applies to everything, including Julius Caesar.

Contrast this with the set of even numbers. Expand the language from the previous example by adding three stable constants, $\in$, $+$, and $\leq$, standing respectively for membership, addition, and the less-than-or-equal relation. Assume we have enough set theory to derive extensionality for $\in$, stated with the classical quantifier. We regiment the set of even numbers using another constant ev of type e,

[39] For an extended discussion of Russell's reason for theorizing in a ramified type theory, see Hodes (2015).

subject to a constraint mirroring the definition given in (14):

$$\blacksquare\hat{\forall}^{\top}x(x \in \mathsf{ev} \leftrightarrow \mathsf{E!}x \wedge \mathsf{Nat}\,x \wedge \hat{\exists}^{\mathsf{Nat}}y(y \leq x \wedge x \equiv y + y)) \tag{18}$$

Here Nat can be defined as before or taken as an additional primitive.

Clearly, ev is not stable, but given natural auxiliary assumptions we can show it is nonetheless metaphysically predicative.[40] The key observation is that for **any** first-order property F, being a natural number among the Fs is equivalent to being a natural number among the $F^{\pm\mathsf{ev}}$s, and ev is neither a natural number among the Fs nor a natural number among the $F^{\pm\mathsf{ev}}$s. So, applying $\mathsf{At}(F)$ or $\mathsf{At}(F^{\pm\mathsf{ev}})$ to $\lambda x.\mathsf{Nat}\,x \wedge \hat{\exists}^{\mathsf{Nat}}y(y \leq x \wedge x \equiv y + y)$ in (18) yields coextensive properties, whence $\mathsf{At}(F)(\mathsf{ev})$ and $\mathsf{At}(F^{\pm\mathsf{ev}})(\mathsf{ev})$ must agree on which natural numbers among the Fs—or, equivalently, natural numbers among the $F^{\pm\mathsf{ev}}$s—they contain, and since membership in either requires being a natural number relative to the relevant domain, neither contains anything else. In particular, neither can contain ev, since the latter is not a natural number relative to either domain. By extensionality, then, $\mathsf{At}(F)(\mathsf{ev})$ and $\mathsf{At}(F^{\pm\mathsf{ev}})(\mathsf{ev})$ are the same set, and ev is metaphysically predicative.[41]

4.3. **Metaphysical Predicativism.** I close by discussing the philosophical fruitfulness of my notion of metaphysical predicativity. The sort of predicativist I have sketched in Section 4.2, recall, wants to impose predicativity as a necessary condition on existence. We can straightforwardly do so through the notion of metaphysical predicativity:

Metaphysical Predicativism$_\sim$: $\forall x\mathsf{Pred}_\sim x$.

This principle is schematic in two ways: different instances are obtained by fixing the type of x and selecting a particular equivalence relation. In what follows, I will call *metaphysical predicativism* the theory that results from adding every instance of the schema above where $\sim$ is coextensiveness and x is of any type.

We have seen that an important application of predicativist theses like Russell's VCP is blocking certain paradoxes. It is fair to wonder whether metaphysical predicativism can do the same. I suggest it can. Metaphysical predicativism can be used as an add-on to solutions to certain paradoxes that work by denying the existence of certain problematic entities. It is used to *explain* why these problematic entities do not exist: they are metaphysically impredicative, and there are no such things as metaphysically impredicative entities. As a proof of concept, I discuss solutions along these lines to two paradoxes: the *Fine–Krämer puzzle of ground* and *Prior's paradox*.

---

[40] These assumptions are: (i) that ev is neither 0 nor the successor of **anything**, (ii) that whenever F exists, so does $\lambda x.Fx \wedge x \not\equiv \mathsf{ev}$, and (iii) that the property $\lambda x.x \not\equiv \mathsf{ev}$ exists.

[41] By a similar argument, a constant n for the *set* of natural numbers, whose members are domain-invariably the existing things that satisfy Nat, must pick out a metaphysically predicative entity, given analogous assumptions (e.g., that n is neither 0 nor a successor). So a set can be predicative even though it is constructed by separation from an impredicative property. I take this to be a feature: self-sensitivity sharpens one way in which an entity can quantificationally involve itself, and n does not quantificationally involve itself this way. Still, n involves itself in another sense: provided there is something other than a natural number, it is quantificationally sensitive to the property of *being a member of* n, by an argument mirroring the one for Nat. This circularity could be regarded as an alternative notion of metaphysical impredicativity. And it may well be that different notions of metaphysical predicativity are better suited for different purposes: the present one for the intensional paradoxes, sensitivity to one's own membership property for a quantificationalist predicativist set theory.

The Fine–Krämer puzzle of ground (Fine, 2010; Krämer, 2013) is a paradox about the relationship between grounding, truth and quantification. Here is a sketch of Krämer's simpler version of the puzzle. Grounding theorists widely hold that, much like true disjunctions are grounded by each of their true disjuncts, true existential generalizations are grounded by each of their true instances (Fine, 2012). If Fx is true and x exists, then the proposition that something is F is grounded by Fx. Formally, with $\prec\colon t \to t \to t$ expressing the grounding relation,

$$P \wedge \mathsf{E!}a \to P \prec \exists x\, P[x/a]. \tag{$\exists$Gr}$$

Grounding theorists also widely hold that grounding is irreflexive:

$$P \not\prec P. \tag{Irr}$$

Consider now the truth-teller proposition, $\exists p\, p$. Suppose it is true and it exists. Being a true existential generalization, the truth-teller is grounded by each of its true instances as per ($\exists$Gr). But the truth-teller is its own instance: the truth-teller is $\beta\eta$-equivalent, hence identical, to the proposition that it is true that there is a truth, $(\lambda p.p)(\exists p\, p)$. So, the truth-teller is self-grounded, contradicting (Irr).

We can block the argument just given by denying that the truth-teller exists. This move should come with a story about why this is so, lest it seem ad hoc. Both Fine (2010, p. 104) and Krämer (2013, p. 88) flirt with the idea that considerations of predicativity could be brought to bear in telling such a story, but neither develops the idea in detail. Metaphysical predicativism can do just that. The truth-teller is metaphysically impredicative, and there are no such things as metaphysically impredicative propositions.

The truth-teller is metaphysically impredicative under the assumption that it is true. Consider the domain of falsehoods, $\lambda p.\neg p$. The truth-teller is false at this domain, which only contains falsehoods. But adding the truth-teller itself to the domain of falsehoods yields a domain that verifies the truth-teller. For the resulting domain does contain a truth, namely the truth-teller itself. The truth-teller, then, is sensitive to itself, hence metaphysically impredicative. Metaphysical predicativism prohibits its existence, so that ($\exists$Gr) cannot be used to derive the paradoxical self-grounding claim. The non-existence of the truth-teller is not simply stipulated to block the argument. Under metaphysical predicativism, it is derived from and explained by its metaphysical impredicativity.

Consider next *Prior's paradox* (Prior, 1961). At midnight, Alice sincerely utters the sentence 'everything I say at midnight is false.' Intuitively, by doing so Alice succeeds in saying exactly one thing at midnight, namely, that everything she says at midnight is false. Curiously, this cannot be.

Formally, regiment the proposition that everything Alice says at midnight is false as

$$\forall p(Sp \to \neg p),$$

where S stands for 'is said by Alice at midnight.' Call this proposition the *Prior proposition.* Assume the Prior proposition is said by Alice at midnight: $S(\forall p(Sp \to \neg p))$. Under this assumption, the Prior proposition cannot be true. If it were, i.e., if it were true that everything Alice said at midnight is false, then the Prior proposition itself would be false, being one of the things Alice said at midnight. So the Prior proposition is false, meaning its negation $\exists p(Sp \wedge p)$ is true: there must be some *true* proposition that Alice said at midnight. This true proposition must be different from the Prior proposition, which is false. Alice must have said

at least two distinct things at midnight. Nothing in the argument depends on the interpretation of S. As is well known, the paradox can be run with any operator S in place of 'Alice says at midnight that,' for instance 'it seems as though Alice, at midnight, says that' or 'Alice, at midnight, fears that.'

The argument just given relies implicitly on the assumption that the Prior proposition exists, which ensures that we can instantiate the universal quantifier in $\forall p(Sp \to \neg p)$ with the Prior proposition itself. The argument can be blocked by rejecting this assumption.[42] Solutions that do so have some advantages. In a positive free logic, denying the existence of the Prior proposition allows one to hold onto the intuitive description of what Alice said at midnight: Alice said at midnight that everything she said at midnight is false, and she said only that:

$$S(\forall p(Sp \to \neg p)) \wedge \forall q(Sq \to q \equiv \forall p(Sp \to \neg p)). \tag{19}$$

This is because, even though she did say at midnight that everything she said at midnight is false, she said nothing at all at midnight.

As before, one should tell a story about why $\forall p(Sp \to \neg p)$ does not exist. It is natural to appeal to impredicativity. The proposition that everything Alice said at midnight is false does intuitively exhibit some form of impredicativity: it seems to quantify over the things that Alice says at midnight, one of which is that proposition itself.[43]

Metaphysical predicativism vindicates this intuition. In the presence of (19) above and the assumption that S is stable, we can prove that the Prior proposition is metaphysically impredicative. If (19) holds, the Prior proposition must be (vacuously) true and must not exist. But if we add the Prior proposition to the current domain, the result is a domain that falsifies the Prior proposition: the resulting domain does contain a true proposition that Alice said at midnight, namely the Prior proposition itself. The proposed predicativist solution to Prior's paradox, then, begins by endorsing the intuitive description of what went on at midnight, formalized by (19). It then shows that the Prior proposition cannot exist, and explains its non-existence via its metaphysical impredicativity.

A distinguishing feature of Prior's paradox is that the Prior proposition is only contingently paradoxical. There is nothing paradoxical in my truly saying that everything that Alice said at midnight is false, for example if all she said at midnight was that snow is green and she did not herself say the Prior proposition at midnight. Metaphysical predicativism can explain this contingency as contingency in the impredicativity of the Prior proposition. There are models of quantificationalism and metaphysical predicativism in which the Prior proposition exists, is true, and is not said by Alice at midnight.

I have given but sketches of solutions to the Fine–Krämer puzzle of ground and Prior's paradox based on metaphysical predicativism. There is a lot more work to be done in fleshing out the details of these solutions and exploring how they compare to alternative approaches. What I hope to have shown with these sketches is that metaphysical predicativity has philosophical potential. Quantificationalism, by providing a framework in which metaphysical predicativity can be formulated, inherits this potential.

[42] Bacon et al. (2016) discuss solutions of this sort at length.

[43] Russell (1908) blames a very similar paradox on impredicativity. The point that the Prior proposition itself exhibits some kind of impredicativity is made explicitly, e.g., in Kaplan (1995); Bacon et al. (2016); Bacon (2021).

5. Further Applications

In this section I sketch a few more applications of quantificationalism. My aim here is to showcase the *breadth* of the quantificationalist research program rather than to explore any particular application in depth. So my discussion will be more superficial and programmatic than in the previous section.

5.1. **Ontological Disputes.** A first application is to *ontological disputes*, namely disputes about what exists. The paradigm is the dispute between the *universalist*, who holds that any things whatsoever compose something, and the *nihilist*, who holds that nothing is composed of anything, so that there are no tables, only simples arranged tablewise (Lewis, 1986; van Inwagen, 1990).

Witnessing their prolonged intractability, some philosophers have thought we should be *deflationists* about many ontological disputes. There is no real disagreement between the universalist and the nihilist. Their dispute is *merely verbal*, much like the dispute about whether a hot dog is a taco. Others, the *realists*, have resisted this conclusion. The disagreement between universalists and nihilists is substantive. The fact that we have not found out which side is right just shows that the question is hard, not that it has no right answer. Quantificationalism has something to offer to both the deflationist and the realist, though I think that, perhaps surprisingly, the most interesting upshots are for the realists.

The best known version of deflationism about ontology is *quantifier variance* (Hirsch, 2002, 2010). The universalist and the nihilist are talking past each other because they mean different things by the quantifier 'there is.' Each of them speaks truly in their own language. You might think quantificationalism is an especially useful framework for quantifier variantists, since it makes generating alternative quantifier meanings very easy: for any property F, $\mathsf{At}(F)(\hat{\forall})$ is a quantifier ranging over precisely the Fs. There is something to this thought. Sider (2009, §5) argues that the quantifier variantist owes us translations between the universalist's language and the nihilist's that preserve logical form, mapping quantifiers to quantifiers. This is easily done in one direction: the universalist simply interprets the nihilist's quantifier as a restriction of their own to the simples. The other direction, Sider claims, is harder, since standard logical frameworks do not provide the tools for explicitly defining quantifier *expansions*. The best the nihilist can do is paraphrase universalist quantification in terms of plural quantification over simples, but this at best yields a truth-preserving sentence-to-sentence mapping, not one from the universalist's quantifier to another quantifier. Quantificationalism supplies exactly the missing tools. The nihilist can define the universalist quantifier by specifying the property *existing or being composite* as the new domain of quantification within their own primitive quantifier. According to the new quantifiers, there are tables, even though in fact there are none.

That being said, I think the most interesting way that quantificationalism reframes the debate over substantivity is on the realist's side. Realists of a certain stripe argue that we can engage in substantive ontological disputes because, even though different quantifier meanings verify and falsify different ontological theses, we can single out a *special* quantifier meaning and decide to conduct ontology in terms of it. Ontology deals with what exists in the sense of the special quantifier.

The realist should say something about what makes the special quantifier special. A particularly attractive idea is that the special quantifier can be characterized

purely by its inferential role: the special quantifier is just the *classical* quantifier (Williamson, 1988, 2003; Dorr, 2014; Bacon, 2024c).[44] We can speak of "the" classical quantifier because any two expressions that satisfy the inferential role of a classical quantifier must be logically equivalent (Harris, 1982). Insofar as the two sides of an ontological dispute in fact are using their quantifiers classically, they must be employing the same notion of existence up to logical equivalence.

But what if either side is not using their quantifiers classically? What if, for example, the universalist and the nihilist engaging in the dispute are both contingentists (or actualists) when it comes to modal metaphysics? As we saw in Section 3.2, contingentists articulate ontological questions in terms of free quantifiers. Two contingentists who debate whether there are any composite objects, paradigmatically an ontological question, will therefore regard 'there are' as a free quantifier. There are compelling reasons to think contingentists are not simply mistaken about what the correct logic of their ontologically committing quantifiers is: contingentists also help themselves to classical "outer" quantifiers to state stronger generalizations, but do not regard these quantifiers as ontologically committing.

The idea that the special quantifier is pinned down by its classical inferential role, therefore, cannot help the realist characterize the dispute between contingentist universalists and contingentist nihilists as substantive. Nor is there an obvious generalization of that idea to the case of merely free quantifiers, for the inferential role standardly associated with free quantifiers does not suffice to pin down their meanings uniquely. For example, '**some** concrete objects' and '**some** spatiotemporally located objects' satisfy the standard inferential role of free quantifiers, but are not obviously equivalent. Which, if either, does the contingentist pick out with their primitive free quantifier?

Quantificationalism can provide an answer to this question. While the *standard* inferential role of free quantifiers does not single them out uniquely, the *quantificationalist* inferential role of free quantifiers does. The quantificationalist can say that the special *free* quantifier is distinguished in virtue of interacting with domain specifiers the right way.

This comes down to two simple ideas about domain specification. First, that *to exist among the Fs* is just to be an F (Gen). And second, that there is a trivial way of specifying the domain of quantification, namely by specifying *the existing things* (Id). Whenever, in addition to the standard inferential role of free quantifiers, 'exists' is used this way, its meaning is uniquely pinned down up to logical equivalence relative to that of domain specifiers.[45]

The dispute between the contingentist universalist and the contingentist nihilist, therefore, can be deemed substantive if both parties also embrace quantificationalism and agree to use their free, ontologically committing quantifiers in accordance with (Gen) and (Id) relative to the same domain specifiers. Williamson (1988) has

[44] A perhaps better known approach is that of Sider (2009, 2011), who suggests instead that the special quantifier is the *fundamental* one, the one that "carves reality at its joints." The worry I am about to raise does not apply to this approach.

[45] If E_1 and E_2 both satisfy (Gen) and (Id), then $E_1 \equiv \mathsf{At}(E_2)(E_1)$ by (Id) for E_2, and $\mathsf{At}(E_2)(E_1) \equiv E_2$ by (Gen) for E_1; so $E_1 \equiv E_2$. Note this argument only goes through if E_1 and E_2 satisfy (Gen) and (Id) with respect to the same syncategorematic operation At, so the argument is not as general as Harris's result for classical quantifiers. Another prominent way of pinning down free quantifiers via a richer inferential role due to Williamson (1988), discussed in the next paragraph, faces a similar worry.

a similar strategy. He shows that a uniqueness result for free quantifiers can be proved if we add the *being constraint* to their inferential role. This is a schema to the effect that each primitive predicate of the language expresses an existence-entailing property. Free logicians can then enter into substantive ontological disputes if they both accept the being constraint *and* are in enough agreement about which predicates they take as primitive. Whether a quantifier is quantificationalist is independent of whether it satisfies the being constraint in a given language, so the two strategies can be mixed and matched. The classical approach to substantivity can coexist with the quantificationalist. The classical approach characterizes one way for a quantifier to be special, namely being *maximally general* (i.e., classical). The quantificationalist approach characterizes another: a quantifier can be special in the sense of being *maximally flexible*, with a domain that is arbitrarily re-specifiable by domain specifiers. No quantifier can be special in both ways, but each way of being special pins down the meaning of a quantifier uniquely up to logical equivalence so that disputes conducted in terms of it are substantive.

5.2. **Quantificationalist Modal Realism.** A second application concerns the reduction of modality to quantification.

Lewis (1968, 1986) famously proposes to reduce modality to quantification over possible worlds: possibly, P just in case P is true at w for some possible world w. Lewis suggests that the modifier 'at w' works

> mainly by restricting the domains of quantifiers in its scope, in much the same way that the restricting modifier 'in Australia' does. In Australia, all swans are black—all swans are indeed black, if we ignore everything not in Australia. [...] At some strange world [w], all swans are blue—all swans are indeed blue, if we ignore everything not part of w. (Lewis, 1986, p. 5)

Lewis does not ultimately *analyze* the modal operator 'possibly' as a combination of a quantifier over worlds and a quantifier domain restrictor. He instead proposes that modal operators and phrases of the form 'at w' themselves be *paraphrased away* using quantification over worlds and syntactic restrictions on quantifiers.[46] To illustrate, Lewis paraphrases (20) as (21):

(20) Possibly, all swans are black.

(21) There is a possible world w such that every part of w that is a swan is black.

In the paraphrase, there is no operator expression 'at w' that functions as a quantifier domain restrictor. The idea that modal operators work partly as quantifier domain restrictors is instead implemented by translating sentences that occur within the scope of a modal operator in the target language by syntactically restricting quantifiers to parts of a specific world. In Sider's (2009) terminology from the previous subsection, Lewis's translation does not preserve logical form.

In quantificationalism, however, Lewis's informal suggestion that modal operators are the composition of a quantifier domain restrictor with a quantifier over worlds can be taken literally. We can articulate a quantificationalist version of modal realism where possibility is equivalent to classical existential quantification

[46]This is the proposal that emerges from Lewis (1968), though Lewis (1986) is skeptical that the paraphrase project can succeed and proposes instead to theorize directly in terms of modal realist ideology rather than treating it as a device to regiment ordinary modalizing.

over possible worlds followed by a domain specifier that restricts the domain to parts of the witnessing world:

$$\Diamond P \leftrightarrow \hat{\exists}^{\top} w(\mathsf{World}\, w \wedge \mathsf{At}(\lambda x.x \preceq w)(P)). \tag{22}$$

Here the phrase 'at w' is rendered as the domain specifier $\mathsf{At}(\lambda x.x \preceq w)$, where $\lambda x.x \preceq w$ is the property of being a part of w. This expression is not a notational device for a paraphrase schema. Though syncategorematic, it is part of the language and really occurs in operator position. The equivalence above can be strengthened to a domain-invariably strict equivalence or even a higher-order identity.

The quantificationalist approach allows us to make literal sense of other remarks of Lewis that, arguably, do not survive the counterpart-theoretic paraphrase implementation. Lewis explicitly suggests that the phrase 'at w' modifies not only the interpretations of quantifiers, but also those of phrases not overtly quantificational but which "reveal quantification under analysis," such as verbs like 'to invent' or 'to eradicate,' and even of proper names (Lewis, 1986, p. 6). But the counterpart-theoretic translation of Lewis (1968) fixes all primitive predicates and only applies to sentences of *pure* quantified modal logic, which lacks individual constants entirely. The only way of making sense of these suggestions in counterpart theory is to say that one should first analyze ordinary language to make implicit quantification syntactically explicit and translate proper names along descriptivist lines, then apply the counterpart-theoretic paraphrase to the result (Divers and Parry, 2018). But in quantificationalism domain specifiers can directly combine with predicates and names and make a non-trivial semantic contribution to the result even if said predicates and names contain no overt occurrences of quantifiers.

Besides taking these suggestive Lewisian remarks at face value, quantificationalist modal realism can also be expected to fare well with respect to the problem of *advanced modalizing* that plagues the Lewisian paraphrase approach. This consists in a collection of puzzles arising when the Lewisian paraphrase is applied to sentences of a language containing both ordinary modal operators and 'advanced' modal vocabulary such as 'possible world,' 'counterpart,' and so on. A simple example: Lewis is committed to the claim that there are at least two worlds, so he should also accept that possibly, there are at least two worlds. But the Lewisian translation of the latter claim is that some world contains two worlds as parts, which Lewis rejects.

This problem arises because Lewis thinks of quantifier domain restriction as a syntactic matter. The quantifier in 'there are at least two possible worlds' is intended to have an unrestricted reading, but translating the modal operator forces a syntactic restriction. In quantificationalism, on the other hand, we understand domain specification semantically. Unrestricted quantification over worlds is rendered via the defined classical quantifier $\hat{\exists}^{\top}$, whose domain cannot be re-specified because $\lambda x.\top$ is a stable property. Thus, the claim that possibly, there are at least two worlds is rendered as

$$\Diamond \hat{\exists}^{\top} wv(\mathsf{World}\, w \wedge \mathsf{World}\, v \wedge w \not\equiv v). \tag{23}$$

In the presence of suitable stability assumptions this turns out to be equivalent to the non-modal claim that there are at least two worlds:

$$\hat{\exists}^{\top} wv(\mathsf{World}\, w \wedge \mathsf{World}\, v \wedge w \not\equiv v). \tag{24}$$

In other words, the possibility operator ends up making a redundant semantic contribution.

This is reminiscent of the solution to the problem of advanced modalizing offered by Divers (1999, 2014), which he calls the *redundancy* solution. Divers suggests that the Lewisian paraphrase should only be applied to sentences where "quantification is world-restricted." Otherwise, modal operators should be treated as *redundant*. An advantage of the quantificationalist approach over Divers's is that it *derives* redundancy from independently motivated principles in the logic of domain specifiers, rather than stipulating it.

5.3. **Further Topics.** Finally, I briefly discuss two additional areas where quantificationalism might apply, though I should emphasize that its prospects in these domains are more speculative. The first is the semantics of *exceptive constructions* in natural language. I introduced domain specifiers by analogy with phrases like 'with the exception of the United States.' It is natural to wonder whether this analogy could rise to the level of an *analysis*.[47] Linguists distinguish *connected* exceptives, such as 'every student but Kim,' from *free* exceptives, such as 'except for Kim.' There is some evidence that free exceptives are syntactically sentence adjuncts (Hoeksema, 1987; von Fintel, 1993), which could perhaps be used to motivate an analysis that treats free exceptives as behaving semantically like sentential operators. Von Fintel (1993) insists that free exceptives cannot be interpreted "at the sentence level" because that would predict that a sentence like 'Except for John, everybody likes John' is uninterpretable, which it is not. To a quantificationalist who interprets free exceptives as domain specifiers, however, this sentence is perfectly interpretable: we just need to say that *liking* is not existence-entailing in its second argument place.

The second area, more promissory still, is *potentialism* in the philosophy of mathematics. Potentialists hold that the universe of sets is never complete: any domain of sets can be expanded to a larger one (Fine, 2006b; Linnebo, 2013). They typically articulate this idea using modal operators, which raises the question of what sort of modality is at issue. Quantificationalism suggests an answer: the potentialist's modality might just be identified with a restriction of domain-occasional truth to domains that expand the current one only by adding "new" mathematical objects. Developing this idea would presumably require a generalization of the framework sketched in Section 3, one on which not every property determines a domain of quantification. Otherwise, the classical quantifier $\hat{\exists}^{\top}$ would provide a domain that cannot be expanded, which is exactly what the potentialist denies.

[47] I want to stress again that neither the intelligibility of domain specifiers nor the truth of quantificationalism rests on this. The meaning of domain specifiers is not given by that of any natural language expression: it is fixed by their theoretical role, namely the logic described in Section 3 and its interpretation in terms of quantificational substitutions. This is why I insisted that pronouncing $\mathsf{At}(F)$ as 'among the Fs' is merely a convention. If exceptives turned out to be best analyzed as, say, modifiers of determiners rather than sentence adjuncts, that would not decide the question whether quantificationalism is true, just as the question whether temporalism is true would not be decided by the discovery that the English expressions 'always' and 'sometimes' are not sentential operators (Sections 1 and 2).

6. Conclusion

I take myself to have shown that quantificationalism is philosophically promising. The seemingly simple idea that truth can vary across domains in much the same way it can vary across time and worlds has a surprisingly wide range of potential applications, of which we have barely begun to scratch the surface. At the same time, the project of making precise sense of domain-occasional truth is far from straightforward, and the non-standard choices I made in carrying it out can be shown to have revisionary metaphysical implications (Appendix A).

Quantificationalism should be judged by weighing the benefits of its philosophical applications against the costs incurred in consistently formulating it. I do not think this is a judgment we are in a position to make at present. But I do believe that the prospects are promising enough to warrant further investigation.

Appendix A. Genuine Syncategorematicity

My logic for quantificationalism treats domain specifiers and ■ as *genuinely syncategorematic* expressions: syncategorematic expressions that generate forbidden positions within their scope. In Section 3 I claimed without argument that this design choice is a natural consequence of the substitutional interpretation of domain specifiers sketched in Section 3.3 and that it leads to a revisionary but ultimately illuminating metaphysics of identity. In this appendix, I elaborate on both claims.

A.1. **Motivations for Genuine Syncategorematicity.** Recall the substitutional interpretation of domain specifiers. Every property determines a *quantificational substitution* that moves interpretations of existence predicates in accordance with (Gen) and fixes everything that has nothing to do with quantification. A term $\mathsf{At}(F)(M)$ denotes the result of performing the quantificational substitution determined by F on M.

The main philosophical reason why domain specifiers are genuinely syncategorematic is that performing a quantificational substitution is a way of combining higher-order entities that irreducibly differs from application. If domain specifiers were categorematic terms, then to specify the domain of quantification would amount to *applying* a higher-order entity *to* another. But that is not how we should think of it. Rather, to specify the domain is to *substitute* a higher-order entity *for* another. Quantificational substitutions are transformations that *preserve* applicative structure, not entities that can themselves be applied to others.

On the quasi-syntactic picture, the difference between applying and substituting is easy to see. The proposition that Alice runs has two constituents: Alice and the *running* relation. Applying negation to the proposition that Alice runs yields a proposition with three constituents: the original two, plus negation. By contrast, substituting Bob for Alice in that proposition yields a proposition with exactly as many constituents, in which neither Alice nor anything playing the role of a "substituter" occurs. This is not to say that substitution never increases constituent count: if we substitute *walking fast* for *running* in the proposition that Alice runs, the result has three constituents, with the extra one being the property of relations expressed by the adverb 'fast.' But application *always* preserves the original constituents, and increases constituent count whenever the entity that is being applied has constituents itself. Substitution, by contrast, need not preserve the original

constituents and does not increase constituent count when the substituted entity and the entity being substituted for have the same number of constituents.

I have urged, however, that the quasi-syntactic picture not be taken literally, so this argument, while suggestive, is certainly not decisive. It is not entirely useless either: quasi-syntactic accounts of metaphysical structure have been shown to be consistent when appropriately restricted, for example to *fundamental* entities only (Bacon, 2023). The argument above is effective for quantificationalists who do buy such restricted accounts.

There is a stronger, if less suggestive, argument to the effect that if we think of domain specifier composition in accordance with (Comp), as the substitutional picture suggests, then we should not think that domain specification is a matter of applying a higher-order entity to another. I will call this the *argument from composition*. Here is the rough idea. Suppose that specifying the Fs as the domain within M simply amounted to applying some entity s_F to M. Then s_F is an entity like any other, and further domain specifiers must act on it too. Either s_F is stable or it is not. If it is, then by (App) specifying the Gs as the domain within $s_F M$ yields s_F applied to $\mathsf{At}(G)(M)$: the two specifications would *commute*, so that resetting the domain to the lovers and then to the unloved would be the same as doing so in the reverse order. But they plainly differ: the first should give the domain of those who are unloved among the lovers, the second that of those who love someone unloved. If s_F is not stable, then specifying the Gs within $s_F M$ acts on M *twice over*: once directly, and once again through its effect on s_F, which is itself sensitive to the domain. Neither option gives the verdict of (Comp), on which the outermost specifier acts on M only once, indirectly, by acting on F.

The argument from composition gives us strong reason to deny that quantificational substitutions combine with their arguments through application, and in turn to embrace the syncategorematic treatment of domain specifiers. But this syncategorematic treatment only does justice to the substitutional picture when *genuine*, i.e., in the presence of the restrictions on λ-abstraction introduced in Section 3.

Here is why. Let $\mathcal{L}_+$ be the higher-order language we obtain from the present one by lifting said restrictions, so that we can abstract variables falling in forbidden positions. In $\mathcal{L}_+$, domain specifiers are still technically syncategorematic, but only *superficially* so. Because we can now abstract into forbidden positions, in any context, we can always find a categorematic term $\lambda X.\mathsf{At}(F)(X)$ that can be used to *stand in* for the syncategorematic domain specifier $\mathsf{At}(F)$, in the following sense: any term $\mathsf{At}(F)(M)$ constructed via the term formation rule for At can always be rewritten as a term of the form $(\lambda X.\mathsf{At}(F)(X))M$, which is constructed by *applying* a categorematic surrogate of the domain specifier to M. The two terms are $\beta\eta$-equivalent, hence intersubstitutable in all contexts *salva veritate*. So, even though we did not set out to treat domain specifiers categorematically, $\mathcal{L}_+$ supplies infinitely many categorematic expressions that do exactly the same job that primitive categorematic domain specifiers would.

Unlike syncategorematic domain specifiers, their categorematic surrogates do combine with their argument through application. And because the ambient logic proves that $\beta\eta$-equivalents are intersubstitutable in all contexts, this means that every instance of domain specification does correspond to the application of a higher-order entity to another. The entity s_F mentioned in the argument from composition is just the denotation of $\lambda X.\mathsf{At}(F)(X)$.

The worry raised by the argument from composition can be stated more precisely in the form of an inconsistency result. Let $\mathtt{QFC}_+$ be the logic axiomatized by the same axioms and rules as $\mathtt{QFC}_\dagger$ but formulated in the language $\mathcal{L}_+$ with unrestricted λ-abstraction.

Theorem A.1. $\mathtt{QFC}_+$ *is inconsistent.*

The proof uses an adaptation of the argument from composition to establish that $\mathtt{QFC}_+$ proves that **there is at most one** stable property at any type.[48] However, $\mathtt{QFC}_+$ also proves that there is more than one stable property at some types, for example the type of sentential operators: negation and the identity combinator $\lambda p.p$ are both stable, but provably non-coextensive. A detailed proof can be found in Section B.3.

If the syncategorematic treatment of domain specifiers is to do justice to the substitutional picture, domain specifiers must be *genuinely* syncategorematic, not merely superficially so. This is what the prohibition on λ-abstraction into forbidden positions achieves. Without the ability to abstract into the second argument place of At, categorematic surrogates of domain specifiers cannot be formed. (App) and (Comp) do not come into conflict and the distinction between substitution and application is preserved, as the consistency of $\mathtt{QFC}_\dagger$ shows.[49]

Theorem A.2. $\mathtt{QFC}_\dagger$ *is consistent.*

The proof of this theorem is beyond the scope of this paper, but can be found in Cleani (Manuscript). Consistency is shown by establishing that $\mathtt{QFC}_\dagger$ is sound with respect to a non-empty class of models built on *quantificational substitution structures*, mathematical structures that precisely model the substitutional interpretation of domain specifiers sketched here.

A.2. **Consequences of Genuine Syncategorematicity.** Throughout the main text I have been conflating applicative indiscernibility ($\equiv$) with identity for simplicity. I now explain why $\equiv$ is actually a weaker relation than identity.

[48] Bacon (2019, pp. 1062–64) shows that similar issues arise in a more general setting when one attempts to express metaphysical substitution-theoretic concepts in the object language. For example, if we had a categorematic term i such that $i(M)$ denotes the result of performing a particular metaphysical substitution on M, then there could be at most two *fundamental* entities at any type, in Bacon's sense of the term. The same conclusion would be reached if i were a syncategorematic term but λ-abstraction were unrestricted.

[49] There is, of course, another possible strategy: instead of restricting λ-abstraction, blame the claim that $\beta\eta$-equivalents are synonymous. For this strategy to work, one must deny not only this claim, but also the much weaker principle that $\beta\eta$-equivalent terms are *coextensive*. The main drawback is that the meaning of λ-terms is no longer uniquely pinned down by the inferential role of λ (see Bacon, 2024c, Prop. 3.1), which is a significant cost if we think of λ as a theoretical term whose meaning is not given by any particular natural language expression or construction. This is why I chose not to pursue this strategy.

Another option, which I consider continuous with my approach, is to introduce a base type i of metaphysical substitutions and regard domain specifiers as terms of type i. Domain specifiers would thus be categorematic but not "applicative" expressions: we would need a primitive operation combining a term of type i with one of some type σ other than i in order to form terms such as $\mathsf{At}(F)(M)$. This operation would have to be regarded as genuinely syncategorematic, on pain of reintroducing surrogates of domain specifiers of functional types $\sigma \to \sigma$. I am intrigued by this approach, but I have not investigated whether it can be consistently worked out.

The core inferential role of the relations $\equiv_\sigma$ is captured by a common schematic regimentation of Leibniz's law in higher-order free logics, which we may call *applicative Leibniz's law*:

$$M \equiv N \to FM \to FN. \tag{ALL}$$

An instance of (ALL) says that given $M \equiv N$, if F *applies* to M, then F also *applies* to N. There is another way of formulating Leibniz's law, which might be called *substitutional Leibniz's law*:

$$M \equiv N \to P[M/x] \to P[N/x]. \tag{SLL}$$

An instance of (SLL) says that if $M \equiv N$, then N is *substitutable for* M *in* P *salva veritate*.

In standard higher-order frameworks with unrestricted abstraction, (ALL) and (SLL) are equivalent. Clearly, any instance of (ALL) is an instance of (SLL). Conversely, a claim of the form $P[M/x]$ can always be rewritten as the result of *applying* the property $\lambda x.P$ to M, so that any instance of (SLL) is $\beta\eta$-equivalent to an instance of (ALL).

However, in a framework such as ours, where abstraction is restricted, there is no guarantee that $\lambda x.P$ is always going to be well formed. For example, from $P \equiv Q$ and $\mathsf{At}(F)(P)$, we cannot use (ALL) to conclude that $\mathsf{At}(F)(Q)$, because $\lambda p.\mathsf{At}(F)(p)$ is ill-formed. What is more, there is no other purely logical argument from $P \equiv Q$ and $\mathsf{At}(F)(P)$ to $\mathsf{At}(F)(Q)$: counter-examples to (SLL) are perfectly consistent in $\mathtt{QFC}_\dagger$, and indeed widespread in the models I use to prove consistency.

One could react to this result by concluding that *identity* does not guarantee intersubstitutivity in all contexts *salva veritate*.[50] That is not my reaction. I take it as granted that whatever identity is, identical things ought to be intersubstitutable in all contexts *salva veritate*, and on that basis conclude that applicative indiscernibility is not identity.

Identity and applicative indiscernibility are relations that appear to coincide from the point of view of standard higher-order frameworks with unrestricted abstraction, but are revealed to be distinct once abstraction is restricted. The theoretical role defining applicative indiscernibility, consisting of (ALL) and a reflexivity axiom, happens to pick out identity in contexts where unrestricted abstraction is assumed, but need not do so if the assumption is dropped. The dialectical situation is analogous to that which arises between the classical and the free logicians with respect to the relations of identity and *Leibniz equivalence*. In classical higher-order logic, it is common to *define* identity as the relation two entities bear when they share all their higher-order properties:

$$\lambda xy.\forall Z(Zx \leftrightarrow Zy).$$

Leibniz equivalence can easily be shown to satisfy a substitutional formulation of Leibniz's law in classical higher-order logics with unrestricted abstraction. It suffices to point out that, given classical logic, Leibniz equivalence satisfies an applicative version of Leibniz's law, and we know that given unrestricted abstraction the difference between applicative and substitutional versions of Leibniz's law is merely notational.

[50] This is the reaction of those who countenance *opacity*: failures of substitutivity of identicals in some contexts, such as attitude reports. See especially Bacon and Russell (2019).

In a free logic, however, universal generalizations fail to imply all their instances. So, two Leibniz-equivalent entities may differ in whether they are F, for F a particular predicate. They must therefore be distinct. Leibniz equivalence and identity are two relations that are taken to coincide by the classical logician, but that the free logician must distinguish.[51]

There are principled quantificationalist metaphysical pictures on which violations of (SLL) are expected and quite pervasive. An example is a picture in which metaphysical necessity does not imply domain-invariable truth: certain domains of quantification are not metaphysically possible. There are several ways of motivating this thought, but an intriguing one comes from the literature on higher-order contingentism (Fritz and Goodman, 2016). Higher-order contingentists typically hold that as a matter of metaphysical necessity, higher-order entities can only exist if all the individuals they are about exist as well. The proposition that Gavin Newsom is tall, for example, is about Gavin Newsom and depends on him for its existence in this sense. On the other hand, the present framework has a feature I have not yet commented on: it takes existence properties at different types to be *freely recombinable*: what exists at one type places no restrictions whatsoever on what exists at other types. So, on this picture, metaphysical necessity is *narrower* than domain-invariable truth: certain claims linking what exists at higher types with what exists at the type of individuals will be metaphysically necessary but not domain-invariably true.

For example, *the proposition that Gavin Newsom is tall exists only if Gavin Newsom exists*—call this proposition P—is metaphysically necessary. Since the Boolean tautology $\top$ is also metaphysically necessary, if metaphysically necessary equivalence implies applicative indiscernibility, then P and $\top$ are applicatively indiscernible. But they cannot be identical. There are domains at which, say, the proposition that Gavin Newsom is tall exists but Gavin Newsom does not. Any such domain verifies $\top$ but must falsify P.

Another very general thesis that generates pervasive counter-examples to (SLL) is a comprehension principle for stable entities. Roughly, such a principle asserts that any entity is applicatively indiscernible from a stable entity called its *stabilizer*. Consider, for example, the non-stable property of *loving someone*. Its stabilizer is a stable property whose extension at any domain is those who love someone *in the current domain*, rigidly specified, which cannot be distinguished from the property of loving someone except by genuinely syncategorematic expressions. The function of stabilizers is somewhat similar to that of actuality operators in modal logic (Crossley and Humberstone, 1977). The property of *actually loving someone* is not the same as the property of loving someone, but the two can only be distinguished via modal vocabulary. A non-stable entity and its stabilizer always generate a counter-example to (SLL), since the former but not the latter can be "moved" by a domain specifier.

[51] If a free logician is in a position to *define* a classical quantifier **all**, then they can reduce identity to *classical* Leibniz equivalence, which is the relation of sharing **all** higher-order properties in the classical sense of "all." My framework uses a primitive free quantifier and has the resources to define classical quantifiers, but we cannot reduce identity to classical Leibniz equivalence either, because the latter does not in general satisfy a substitutional version of Leibniz's law in the framework. However, the system does prove that classical Leibniz equivalence and applicative indiscernibility are logically equivalent.

If the quantificationalist embraces violations of (SLL), how should they theorize about identity? Perhaps we could introduce additional typed constants into the language standing for higher-order identity relations $=_\sigma$ and require them to satisfy appropriate substitutional versions of Leibniz's law. This cannot be consistently done, for it would reintroduce categorematic surrogates of domain specifiers into the system: one could simply write

$$\widetilde{\mathsf{At}}(F) := \lambda Xy.\exists Z(Z = X \wedge \mathsf{At}(F)(Z)y)$$

instead of $\widehat{\mathsf{At}}(F) := \lambda X.\mathsf{At}(F)(X)$. The sense in which $\widetilde{\mathsf{At}}(F)$ is a categorematic surrogate of $\mathsf{At}(F)$ is a little weaker, since $\widetilde{\mathsf{At}}(F)(M)$ and $\mathsf{At}(F)(M)$ are only co-extensive when M exists. But, it turns out, that is enough to recover a version of the argument from composition. Roughly, the trick is to run an analog of the same argument within the scope of a domain specifier that forces a universal domain at the type of M. The conclusion is that within the scope of said domain specifier, there is at most one stable property up to applicative indiscernibility, which again $\mathtt{QFC}_\dagger$ disproves.

Thus, if we want to add identity to the language, it should also be a genuinely syncategorematic expression. Happily, the quantificationalist has a story to tell about why we should *expect* identity to be genuinely syncategorematic. In a slogan: *identity is domain-invariable applicative indiscernibility.* Since identity is *defined* in terms of domain-invariable truth this way, its genuine syncategorematicity is no surprise. It is simply another consequence of the genuine syncategorematicity of domain specifiers.

What I mean by this slogan is that the quantificationalist can consistently endorse a substitutional version of Leibniz's law for the composition of applicative indiscernibility with the operator ■:

$$\blacksquare(M \equiv N) \to P[M/x] \to P[N/x], \tag{SLL$_\blacksquare$}$$

where x must be free for both M and N in P.[52] Entities that are domain-invariably applicatively indiscernible are intersubstitutable in all contexts.

We know that applicatively indiscernible entities need not be domain-invariably applicatively indiscernible. This fact is exactly what allowed us to generate counter-examples to (SLL). On the higher-order contingentist picture sketched earlier, for example, $\top$ and *the proposition that Gavin Newsom is tall exists only if Gavin Newsom exists* were in fact applicatively indiscernible, but became applicatively discernible once we specified the domain to be one where Gavin Newsom does not

[52] Officially, I treat (SLL$_\blacksquare$) as an axiom schema, but it is worth noting that it can be derived as a theorem of the system obtained by replacing (SLL$_\blacksquare$) with appropriate principles of propositional granularity. An example is an axiom schema to the effect that propositions are applicatively indiscernible when domain-invariably materially equivalent:

$$\blacksquare(p \leftrightarrow q) \to p \equiv q. \tag{QFreg}$$

Alternatively, we could replace (SLL$_\blacksquare$) with an inference rule, derivable from (QFreg), asserting that if we can *prove* that p and q are materially equivalent, then they must be applicatively indiscernible:

$$\text{from } \vdash p \leftrightarrow q \text{ infer } p \equiv q. \tag{Equiv}$$

The reason is that (Equiv), and therefore (QFreg), ensures that $\blacksquare(P \equiv Q)$ implies both $\forall xP \equiv \forall xQ$ and $\blacksquare P \equiv \blacksquare Q$. These claims, notice, cannot be straightforwardly derived from (ALL) because both the quantifier and ■ are syncategorematic expressions. Once we have them, we can show that $\blacksquare(A \equiv B)$ implies $\blacksquare(M[A/x] \equiv M[B/x])$ by induction on the structure of M.

exist but the proposition that he is tall does. Relative to such a domain, the two propositions are discernible via the property of being true. What (SLL$_\blacksquare$) tells us is that domain specification is the *only* way we can distinguish between two applicatively indiscernible entities. If two entities are not just in fact applicatively indiscernible, but remain so no matter what we specify the domain to be, then there is no way to distinguish between them whatsoever.

Distinguishing between applicative indiscernibility and identity is an undoubtedly revisionary move. But the distinction is principled: it is downstream of the philosophically motivated commitment to the genuine syncategorematicity of domain specifiers. Moreover, it is illuminating in that it makes room for a distinctly quantificationalist understanding of identity itself.

Appendix B. Technical Details

This appendix states the core of the formal framework as compactly as possible. Proofs and further discussion can be found in Cleani (Manuscript).

B.1. **The Language $\mathcal{L}_\dagger$.** Types are generated from e and t by $\to$ in the usual way. A *nice type sequence* is a finite, non-repeating sequence of types $\bar\sigma$; we write $\sigma \in \bar\sigma$ when σ occurs in $\bar\sigma$. Fix a typed supply of variables Var^σ and a signature Σ extending the logical signature Λ, which contains $\to$: $t \to t \to t$, $\bot : t$, and, for each type σ, an existence predicate $\mathsf{E!}_\sigma : \sigma \to t$ and an applicative indiscernibility constant $\equiv_\sigma$: $\sigma \to \sigma \to t$.

The terms of the full language $\mathcal{L}_+(\Sigma)$ and their types are given by the following clauses, where $M \cdot \sigma$ means that M is a $\mathcal{L}_+(\Sigma)$-term of type σ:

(25) $x \cdot \sigma$ for $x \in Var^\sigma$, and $C \cdot \sigma$ for $C \in \Sigma^\sigma$;
(26) $MN \cdot \tau$ whenever $M \cdot \sigma \to \tau$ and $N \cdot \sigma$;
(27) $\lambda x.M \cdot \sigma \to \tau$ whenever $M \cdot \tau$ and $x \in Var^\sigma$;
(28) $\forall x P \cdot t$ whenever $P \cdot t$ and x is a variable of any type;
(29) $\mathsf{At}(\bar F)(M) \cdot \tau$ whenever $M \cdot \tau$ and $\bar F = (F_\sigma)_{\sigma\in\bar\sigma}$ is a *nice term sequence*, i.e., $\bar\sigma$ is a nice type sequence and $F_\sigma \cdot \sigma \to t$ for each $\sigma \in \bar\sigma$;
(30) $\blacksquare P \cdot t$ whenever $P \cdot t$.

The domain specifiers $\mathsf{At}(F)$ of the main text abbreviate $\mathsf{At}(\bar F)$ for the unit sequence $\bar F = (F)$. We write $\mathsf{FV}_\sigma(M)$ for the free variables of type σ in M, and $\blacklozenge P$ for $\neg\blacksquare\neg P$. $\exists x P$ abbreviates $\neg\forall x\neg P$. The categorematic quantifiers are defined by $\hat\forall_\sigma := \lambda X^{\sigma\to t}.\forall y Xy$ and $\hat\exists_\sigma := \lambda X^{\sigma\to t}.\exists y Xy$. The other Boolean connectives and their liftings to higher types are defined as usual.

The set $\mathsf{AV}_\sigma(M)$ of variables of type σ that are *abstractable* in M is defined by:

$$\begin{aligned}
\mathsf{AV}_\sigma(x) = \mathsf{AV}_\sigma(C) &:= Var^\sigma \\
\mathsf{AV}_\sigma(MN) &:= \mathsf{AV}_\sigma(M) \cap \mathsf{AV}_\sigma(N) \\
\mathsf{AV}_\sigma(\lambda x.M) &:= \mathsf{AV}_\sigma(M) \cup (\{x\} \cap Var^\sigma) \\
\mathsf{AV}_\sigma(\forall x P) &:= \mathsf{AV}_\sigma(P) \cup (\{x\} \cap Var^\sigma) \\
\mathsf{AV}_\sigma(\mathsf{At}(\bar F)(M)) &:= \big(\textstyle\bigcap_{\tau\in\bar\sigma} \mathsf{AV}_\sigma(F_\tau)\big) \smallsetminus \mathsf{FV}_\sigma(M) \\
\mathsf{AV}_\sigma(\blacksquare P) &:= Var^\sigma \smallsetminus \mathsf{FV}_\sigma(P)
\end{aligned}$$

In other words, the abstractable variables of M are those that do not have free occurrences in forbidden positions in M. We drop the type subscripts on FV and

AV when the type is clear. The language $\mathcal{L}_\dagger(\Sigma)$ is defined like $\mathcal{L}_+(\Sigma)$, except that clause (27) is replaced by

(27*) $\lambda x.M : \sigma \to \tau$ whenever $M : \tau$ and $x \in \mathsf{AV}_\sigma(M)$,

where $M : \sigma$ means that M is a $\mathcal{L}_\dagger(\Sigma)$-term of type σ.

B.2. **The Logic $\mathtt{QFC}_\dagger$.** Given nice term sequences $\bar{F}$ and $\bar{G}$, indexed by $\bar{\sigma}$ and $\bar{\tau}$ respectively, their *substitutional composition* $\bar{F}\bullet\bar{G}$ is the nice term sequence indexed by the types occurring in either $\bar{\sigma}$ or $\bar{\tau}$, with

$$(\bar{F} \bullet \bar{G})_\rho := \begin{cases} \mathsf{At}(\bar{F})(G_\rho) & \text{if } \rho \in \bar{\tau}, \\ F_\rho & \text{if } \rho \in \bar{\sigma} \text{ but } \rho \notin \bar{\tau}. \end{cases}$$

A *combinator* is a closed term built from variables using only application and λ-abstraction. The logic $\mathtt{QFC}_\dagger$ is the set of $\mathcal{L}_\dagger$-terms $P : t$ such that $\vdash P$, where $\vdash$ is the least relation between sets of $\mathcal{L}_\dagger$-formulas and $\mathcal{L}_\dagger$-formulas satisfying the conditions in Figures 2 to 5. All schemas are restricted to $\mathcal{L}_\dagger$-terms of the appropriate types, and a, M, N are assumed to be free for the relevant variables.

$\vdash P$	P a classical tautology	(Taut)
$\vdash \forall x\, \mathsf{E}!x$		($\forall\mathsf{E}!$)
$\vdash \mathsf{E}!a \to (\forall x P \to P[a/x])$		(FrUI)
$\vdash \forall x(P \to Q) \to (\forall x P \to \forall x Q)$		(Norm)
$\vdash M \equiv M$		(Refl)
$\vdash M \equiv N \to (FM \to FN)$		(ALL)
$\vdash M \equiv N$	M, N $\beta\eta$-equivalent	($\beta\eta$)
$P,\ P \to Q \vdash Q$		(MP)
If $\Gamma \vdash P \to Q$, then $\Gamma \vdash P \to \forall x Q$	x fresh	(UG)

FIGURE 2. Free higher-order logic.

$\vdash \mathsf{At}(\bar{\mathsf{E}}!)(M) \equiv M$	$\bar{\mathsf{E}}! = (\mathsf{E}!_\sigma)_{\sigma\in\bar{\sigma}}$	(Id)
$\vdash \mathsf{At}(\bar{F})(MN) \equiv \mathsf{At}(\bar{F})(M)(\mathsf{At}(\bar{F})(N))$		(App)
$\vdash \mathsf{At}(\bar{F})(\mathsf{At}(\bar{G})(M)) \equiv \mathsf{At}(\bar{F} \bullet \bar{G})(M)$		(Comp)
$\vdash \mathsf{At}(\bar{F})(M) \equiv \mathsf{At}(\bar{G})(M)$	$\bar{G}$ permutes $\bar{F}$	(Perm)
$\vdash \mathsf{At}(\bar{F})(\mathsf{E}!_\sigma) \equiv F_\sigma$	$\sigma \in \bar{\sigma}$	(Gen)
$\vdash \mathsf{At}(\bar{F})(\mathsf{E}!_\sigma) \equiv \mathsf{E}!_\sigma$	$\sigma \notin \bar{\sigma}$	(Mod)
$\vdash \mathsf{At}(\bar{F})(\heartsuit) \equiv \heartsuit$	$\heartsuit \in \{\to, \bot, \equiv_\sigma\}$	(Stab$\heartsuit$)
$\vdash \mathsf{At}(\bar{F})(M) \equiv M$	M a combinator	(Stabλ)
If $\vdash P$, then $\vdash \mathsf{At}(\bar{F})(P)$		(At-Nec)
If $\Gamma \vdash P \to \mathsf{At}(\bar{F})(Q)$, then $\Gamma \vdash P \to \mathsf{At}(\bar{F})(\forall x Q)$	x fresh	(At-UG)

FIGURE 3. Domain specifiers.

$$\vdash \blacksquare P \to \mathsf{At}(\bar{F})(P) \qquad \text{(Master)}$$

$$\vdash \blacksquare(M \equiv N) \to (P[M/x] \to P[N/x]) \qquad (\text{SLL}_{\blacksquare})$$

$$\text{If } \Gamma \vdash P \to \mathsf{At}(\bar{F})(\mathsf{At}(\bar{X}^{\bar{\sigma}_i})(Q)) \text{ for all } i \in \omega, \text{ then } \Gamma \vdash P \to \mathsf{At}(\bar{F})(\blacksquare Q) \quad \{\bar{X}^{\bar{\sigma}_i} : i \in \omega\} \text{ coherent, fresh} \qquad \text{(Surj)}$$

FIGURE 4. Domain invariability.

$$\text{If } \vdash P \leftrightarrow Q, \text{ then } \vdash \blacksquare(\lambda \bar{x}.P \equiv \lambda \bar{x}.Q) \qquad (\text{Equiv}^{+})$$

FIGURE 5. Granularity.

In general, a variable is *fresh* for a rule when it does not occur free in any of the assumptions nor in the conclusion of the rule. E.g., for (At-UG), x cannot be free in any formula in Γ, nor in P, nor in $\bar{F}$. The notion of a collection of sequences of variables being fresh is defined similarly. In (Surj), $\{\bar{X}^{\bar{\sigma}_i} : i \in \omega\}$ is called *coherent* when each $\bar{X}^{\bar{\sigma}_i}$ is a nice term sequence of variables defined exactly on $\bar{\sigma}_i$ and for each type σ the same variable is used across all sequences defined on that type. The rule (Surj) is infinitary, since it generalizes over all nice type sequences. Where $\bar{F}$ is a sequence of existence predicates, it reduces by (Id) to the simpler rule mentioned in Section 3: if $\mathsf{At}(\bar{X}^{\bar{\sigma}_i})(P)$ is derivable for every $i \in \omega$, for some coherent collection of fresh $\bar{X}^{\bar{\sigma}_i}$, then so is $\blacksquare P$. In $\mathtt{QFC}_{\dagger}$, $\blacksquare$ is a normal modality satisfying $\mathtt{T}$ and $\mathtt{4}$ and closed under necessitation. In the main text, I stated (Equiv^{+}) without $\blacksquare$ in the consequent, but clearly the two formulations are equivalent given that theorems can be $\blacksquare$-necessitated (by (At-Nec) and (Surj)) and that $\blacksquare$ satisfies $\mathtt{T}$.

All three of the following principles governing the interaction of domain specifiers with categorematic quantifiers are derivable in $\mathtt{QFC}_{\dagger}$:

$$\hat{\forall}^{F} x\, Fx \tag{31}$$

$$\hat{\forall}^{F} x\, Gx \to (Fy \to Gy) \tag{32}$$

$$\hat{\forall}^{F} x(Gx \to Hx) \to (\hat{\forall}^{F} x\, Gx \to \hat{\forall}^{F} x\, Hx) \tag{33}$$

Moreover, $\mathtt{QFC}_{\dagger}$ is closed under universal generalization for categorematic quantifiers:

$$\text{if } \vdash Fx \to P, \text{ then } \vdash \hat{\forall}^{F} x\, P \qquad x \notin \mathsf{FV}(F), \text{ and } x \in \mathsf{AV}(P). \qquad (\hat{\text{UG}})$$

For (31), use (At-Nec) on a $\beta\eta$-equivalent of ($\forall$E!) and distribute. For (32), note that (FrUI), (Taut) and (Equiv^{+}) give

$$\blacksquare((\lambda Xy.(\hat{\forall} X \to (\mathsf{E}!y \to Xy))) \equiv \lambda Xy.\top). \tag{34}$$

The left-hand side of this equation is the result of applying a combinator to $\mathsf{E}!$, $\hat{\forall}$ and $\to$, and the right-hand side is stable by (Stabλ) and (Stab$\heartsuit$). So applying (Master) and distributing with (App), (Gen), (Stab$\heartsuit$) and (Stabλ) gives

$$(\lambda Xy.(\hat{\forall}^{F} X \to (Fy \to Xy))) \equiv \lambda Xy.\top. \tag{35}$$

Since $(\lambda Xy.\top)Gy$ is β-equivalent to $\top$, (ALL) and ($\beta\eta$) give (32). A similar argument establishes (33).

For (UĜ), first note that both of the following are derivable in $\mathtt{QFC}_\dagger$:

$$\hat{\forall}^F x\, P \leftrightarrow \hat{\forall}^F x\, (Fx \to P); \tag{36}$$

$$\hat{\forall}^F (\lambda x.\top). \tag{37}$$

The former follows using a combinator argument similar to the one used for (32); the latter follows from the free-logical truth $\hat{\forall}\lambda x.\top$ and (At-Nec). Now suppose $\vdash Fx \to P$ with x not free in F and $x \in \mathsf{AV}(P)$. Then $\vdash (Fx \to P) \leftrightarrow \top$, so by (Equiv$^+$) we get $\vdash \blacksquare((\lambda x.Fx \to P) \equiv \lambda x.\top)$. So, (SLL$_\blacksquare$) and (37) give $\vdash \hat{\forall}^F x\, (Fx \to P)$. By (36) this implies $\vdash \hat{\forall}^F x\, P$, as desired.

B.3. **The Inconsistency of $\mathtt{QFC}_+$.** Let $\mathtt{QFC}_+$ be the logic axiomatized the same way as $\mathtt{QFC}_\dagger$, but in the language $\mathcal{L}_+$.

Theorem B.1. $\mathtt{QFC}_+$ *is inconsistent.*

Proof. Let us abbreviate the categorematic surrogate $\lambda x.\mathsf{At}(F)(x)$ of a domain specifier as $\widehat{\mathsf{At}}(F)$. Suppose that F and G are stable properties. Since $\mathsf{At}(G)(M)$ and $\widehat{\mathsf{At}}(G)M$ are $\beta\eta$-equivalent, ($\beta\eta$) and (ALL) give $\mathsf{At}(F)(\mathsf{At}(G)(M)) \equiv \mathsf{At}(F)(\widehat{\mathsf{At}}(G)M)$, and (App) then yields

$$\mathsf{At}(F)(\mathsf{At}(G)(M)) \equiv \mathsf{At}(F)(\widehat{\mathsf{At}}(G))(\mathsf{At}(F)(M)). \tag{$\star$}$$

Begin with an instance of (Gen):

$$\mathsf{At}(F)(\mathsf{E!}) \equiv F. \tag{38}$$

In $\mathtt{QFC}_+$ we may specify the Gs as the domain on both sides, by (ALL) applied to the property $\lambda x.\, \mathsf{At}(G)(\mathsf{At}(F)(\mathsf{E!})) \equiv \mathsf{At}(G)(x)$, which is well formed only because abstraction is unrestricted. This yields:

$$\mathsf{At}(G)(\mathsf{At}(F)(\mathsf{E!})) \equiv \mathsf{At}(G)(F). \tag{39}$$

Because F is stable, the right-hand side reduces to F. But, using ($\star$) and ($\beta\eta$), we can show that the left-hand side reduces to G. One application of ($\star$) gives us

$$\mathsf{At}(G)(\mathsf{At}(F)(\mathsf{E!})) \equiv \mathsf{At}(G)(\widehat{\mathsf{At}}(F))(\mathsf{At}(G)(\mathsf{E!})). \tag{40}$$

By (Gen) we have $\mathsf{At}(G)(\mathsf{E!}) \equiv G$ and by the stability of G we have $G \equiv \mathsf{At}(G)(G)$. So, we may substitute $\mathsf{At}(G)(G)$ for $\mathsf{At}(G)(\mathsf{E!})$ in the equation above, yielding

$$\mathsf{At}(G)(\mathsf{At}(F)(\mathsf{E!})) \equiv \mathsf{At}(G)(\widehat{\mathsf{At}}(F))(\mathsf{At}(G)(G)). \tag{41}$$

Applying ($\star$) once more to "pull out" the domain specifier $\mathsf{At}(G)$ we see that the right-hand side equals $\mathsf{At}(G)(\mathsf{At}(F)(G))$. Since G is stable, this reduces to G, so we may conclude that $G \equiv F$. So, **there is at most one** stable property up to applicative indiscernibility at any type. However, some types, for example $t \to t$, contain **more than one**: $\lambda p.p$ is stable by (Stabλ) and $\neg$ is stable by (Stab$\heartsuit$), (Stabλ) and (App), but clearly $(\lambda p.p) \not\equiv \neg$. □

B.4. **Quantificational Substitution Structures.**

Definition B.2. An *applicative structure* is a pair (A, App), where A is a typed family of sets A^σ and App a typed family of functions $App^{\sigma\tau} : A^{\sigma\to\tau} \times A^\sigma \to A^\tau$. A *substitution structure* is a triple $((I, \circ), (A, App), Sub)$ where $(I, \circ)$ is a monoid with identity 1, (A, App) is an applicative structure, and Sub is a typed family of maps $I \times A^\sigma \to A^\sigma$, written $(i, \mathbf{a}) \mapsto i\mathbf{a}$, such that:

(42) $1\mathbf{a} = \mathbf{a}$;
(43) $i(j\mathbf{a}) = (i \circ j)\mathbf{a}$;
(44) $i\, App(\mathbf{f}, \mathbf{a}) = App(i\mathbf{f}, i\mathbf{a})$.

An element $\mathbf{a}$ is *stable* when $i\mathbf{a} = \mathbf{a}$ for all $i \in I$.

Definition B.3. Let (A, App) be an applicative structure. A *domain* is a partial map $\bar{\mathbf{f}}$ defined on finitely many types, with $\bar{\mathbf{f}}(\sigma) \in A^{\sigma\to t}$ whenever defined; *Dom* is the set of domains. A *quantificational substitution structure* (QSS) is a substitution structure equipped with a surjection $\bar{\mathbf{f}} \mapsto i_{\bar{\mathbf{f}}}$ from *Dom* onto I such that:

(45) For each type σ there is some $\mathbf{e}_\sigma \in A^{\sigma\to t}$ such that $i_{\mathbf{e}_\sigma} = 1$;
(46) (Composition) $i_{\bar{\mathbf{f}}} \circ i_{\bar{\mathbf{g}}} = i_{\bar{\mathbf{f}}\bullet\bar{\mathbf{g}}}$ for all domains $\bar{\mathbf{f}}, \bar{\mathbf{g}}$;
(47) (Identity) $\bar{\mathbf{e}} \bullet \bar{\mathbf{f}} = \bar{\mathbf{f}} \bullet \bar{\mathbf{e}}$ for every domain $\bar{\mathbf{f}}$ and every domain $\bar{\mathbf{e}}$ with $i_{\bar{\mathbf{e}}} = 1$.

Here $(\bar{\mathbf{f}} \bullet \bar{\mathbf{g}})(\rho)$ is $i_{\bar{\mathbf{f}}}\, \bar{\mathbf{g}}(\rho)$ if $\bar{\mathbf{g}}(\rho)$ is defined, $\bar{\mathbf{f}}(\rho)$ if only $\bar{\mathbf{f}}(\rho)$ is defined, and undefined otherwise.

Given the other conditions, (Identity) is equivalent to requiring that $i_{\bar{\mathbf{f}}}\mathbf{e}_\sigma = \bar{\mathbf{f}}(\sigma)$ whenever $\bar{\mathbf{f}}(\sigma)$ is defined, and $i_{\bar{\mathbf{f}}}\mathbf{e}_\sigma = \mathbf{e}_\sigma$ otherwise. These are the semantic counterparts of (Gen) and (Mod).

B.5. **Models.** A variable assignment g maps each $x \in Var^\sigma$ to an element of A^σ. For a nice term sequence $\bar{F}$, write $[\![\bar{F}]\!]^g$ for the domain mapping σ to $[\![F_\sigma]\!]^g$ for each $\sigma \in \bar{\sigma}$. A *pointed QSS* is a QSS $\mathfrak{A}$ together with a relation $\models$ between $\mathfrak{A}$ and A^t. Given a pointed QSS $\mathfrak{A}$, for $\mathbf{a}, \mathbf{b} \in A^\sigma$ write $\mathbf{a} \equiv_{\mathfrak{A}} \mathbf{b}$ when $\mathfrak{A} \models App(\mathbf{h}, \mathbf{a})$ iff $\mathfrak{A} \models App(\mathbf{h}, \mathbf{b})$ for every $\mathbf{h} \in A^{\sigma\to t}$, and, for a domain $\bar{\mathbf{f}}$, write $\mathbf{a} \cong^{\bar{\mathbf{f}}}_{\mathfrak{A}} \mathbf{b}$ when $i_{\bar{\mathbf{f}}} i_{\bar{\mathbf{g}}}\mathbf{a} \equiv_{\mathfrak{A}} i_{\bar{\mathbf{f}}} i_{\bar{\mathbf{g}}}\mathbf{b}$ for every domain $\bar{\mathbf{g}}$.

Definition B.4. An *interpretation* on a pointed QSS $\mathfrak{A}$ is a map $[\![\cdot]\!]^{(\cdot)}$ taking each $\mathcal{L}_\dagger$-term $M : \sigma$ and assignment g to an element of A^σ such that the following hold for every assignment g and every domain $\bar{\mathbf{f}}$:

(48) $[\![x]\!]^g = g(x)$;
(49) $[\![MN]\!]^g = App([\![M]\!]^g, [\![N]\!]^g)$;
(50) $[\![M]\!]^g = [\![N]\!]^h$ whenever M, N are $\beta\eta$-equivalent and g, h agree on their common free variables;
(51) $i_{[\![\mathsf{E}!_\sigma]\!]^g} = 1$;
(52) $[\![\mathsf{At}(\bar{F})(M)]\!]^g = i_{[\![\bar{F}]\!]^g}[\![M]\!]^g$;
(53) $[\![M]\!]^g \cong^{\bar{\mathbf{f}}} [\![M]\!]^h$ whenever $g(y) \cong^{\bar{\mathbf{f}}} h(y)$ for every $y \in \mathsf{FV}(M)$;
(54) $[\![M]\!]^{g[x\mapsto [\![N]\!]^g]} = [\![M[N/x]]\!]^g$ whenever N is free for x in M.

The interpretation is *stable* when $[\![\to]\!]^g$, $[\![\bot]\!]^g$, $[\![\equiv]\!]^g$ and the denotations of all combinators are stable.

Definition B.5. A *$\mathcal{L}_\dagger$-model* consists of a pointed QSS $\mathfrak{A}$ together with a stable interpretation $[\![\cdot]\!]^{(\cdot)}$ on it, such that the following hold for every assignment g and every domain $\bar{\mathbf{f}}$:

(55) $\mathfrak{A} \models [\![P \to Q]\!]^g$ iff $\mathfrak{A} \not\models [\![P]\!]^g$ or $\mathfrak{A} \models [\![Q]\!]^g$;
(56) $\mathfrak{A} \not\models [\![\bot]\!]^g$;
(57) $\mathfrak{A} \models [\![M \equiv N]\!]^g$ iff $[\![M]\!]^g \equiv_{\mathfrak{A}} [\![N]\!]^g$;
(58) $\mathfrak{A} \models i_{\bar{\mathbf{f}}}[\![\forall x P]\!]^g$ iff $\mathfrak{A} \models i_{\bar{\mathbf{f}}}[\![P]\!]^{g[x \mapsto \mathbf{a}]}$ for every $\mathbf{a} \in A^\sigma$ with $\mathfrak{A} \models App(i_{\bar{\mathbf{f}}}[\![\mathsf{E}!_\sigma]\!]^g, i_{\bar{\mathbf{f}}}\mathbf{a})$,
(59) $\mathfrak{A} \models i_{\bar{\mathbf{f}}}[\![\blacksquare P]\!]^g$ iff $\mathfrak{A} \models (i_{\bar{\mathbf{f}}} \circ i_{\bar{\mathbf{h}}})[\![P]\!]^g$ for every $\bar{\mathbf{h}} \in Dom$.

A formula P is *valid* in a model if $\mathfrak{A} \models [\![P]\!]^g$ for every assignment g, and valid *simpliciter* if valid in every $\mathcal{L}_\dagger$-model.

Clauses (58) and (59) build substitutions into the satisfaction conditions because quantifiers and $\blacksquare$, being syncategorematic, receive no independent denotations. The quantificational clause says that the standard free-logical satisfaction conditions hold not only at the current domain but after every quantificational substitution; the clause for $\blacksquare$ treats it as a generalization over all such substitutions.

B.6. **Consistency.** A model is called *quasi-extensional* if propositions that are materially equivalent under every substitution are identical and its underlying QSS is quasi-functional (properties that agree on every argument under every substitution are identical).

Theorem B.6. $\mathtt{QFC}_\dagger$ *is sound with respect to quasi-extensional* $\mathcal{L}_\dagger$*-models. Consequently it is consistent, and so is* $\mathtt{QFC}_\dagger$ *extended with quantificationalism.*

Quasi-extensional models can be constructed, for example, by taking suitable direct powers of extensional models of classical higher-order logic. $\mathtt{QFC}_\dagger$ is not sound with respect to arbitrary $\mathcal{L}_\dagger$-models, but its (Equiv$^+$)-free fragment is, and is indeed complete with respect to that class. For details, see Cleani (Manuscript). For the second part, $\mathtt{QFC}_\dagger$ proves $\mathsf{At}(\bot_{t\to t})(\neg \mathsf{E}!\top)$ using (Gen), (App) and the stability of $\top$: at the empty domain of propositions, $\top$ does not exist. So $\mathsf{E}!\top \wedge \mathsf{E}!(\mathsf{E}!\top)$ implies $\mathsf{E}!\top \wedge \mathsf{E}!(\mathsf{E}!\top) \wedge \blacklozenge\neg\mathsf{E}!\top$, and hence $\exists p(p \wedge \blacklozenge\neg p)$ by (FrUI). Since $\blacksquare$ satisfies $\mathtt{T}$, this yields quantificationalism. As $\mathsf{E}!\top \wedge \mathsf{E}!(\mathsf{E}!\top)$ is satisfiable in some quasi-extensional $\mathcal{L}_\dagger$-model, quantificationalism is consistent with $\mathtt{QFC}_\dagger$.

References

Bacon, A. [2013]. Quantificational Logic and Empty Names. *Philosophers' Imprint*, 13.

——— [2018]. Tense and Relativity. *Noûs*, 52(3):667–696.

——— [2019]. Substitution Structures. *Journal of Philosophical Logic*, 48(6):1017–1075.

——— [2020]. Logical Combinatorialism. *Philosophical Review*, 129(4):537–589.

——— [2021]. Opacity and Paradox. In C. Nicolai and J. Stern, eds., *Modes of Truth: The Unified Approach to Truth, Modality, and Paradox*, pp. 231–265. New York: Routledge.

——— [2023]. A Theory of Structured Propositions. *Philosophical Review*, 132(2):173–238.

——— [2024a]. A case for higher-order metaphysics. In P. Fritz and N. K. Jones, eds., *Higher-Order Metaphysics*, pp. 47–72. Oxford: Oxford University Press.

——— [2024b]. Mathematical Modality: An Investigation in Higher-order Logic. *Journal of Philosophical Logic*, 53(1):131–179.

——— [2024c]. *A Philosophical Introduction to Higher-order Logics*. New York: Routledge.

Bacon, A., Hawthorne, J., and Uzquiano, G. [2016]. Higher-Order Free Logic and the Prior-Kaplan Paradox. *Canadian Journal of Philosophy*, 46(4-5):493–541.

Bacon, A. and Russell, J. S. [2019]. The Logic of Opacity. *Philosophy and Phenomenological Research*, 99(1):81–114.

Brogaard, B. [2012]. *Transient Truths: An Essay in the Metaphysics of Propositions*. Oxford University Press.

Buridan, J. [1976]. *Iohannis Buridani Tractatus de Consequentiis*, *Philosophes médiévaux*, vol. 16. Louvain: Publications universitaires. Ed. Hubert Hubien.

Chihara, C. [1973]. *Ontology and the Vicious-Circle Principle*. Cornell University Press.

Cleani, A. M. [Manuscript]. The Logic of Quantificationalism: Foundations.

Crossley, J. N. and Humberstone, L. [1977]. The logic of "actually". *Reports on Mathematical Logic*, 8:11–29.

Divers, J. [1999]. A Genuine Realist Theory of Advanced Modalizing. *Mind*, 108(430):217–239.

——— [2014]. The Modal Status of the Lewisian Analysis of Modality. *Mind*, 123(491):861–872.

Divers, J. and Parry, J. J. [2018]. Advanced modalizing de dicto and de re. *Analysis*, 78(3):415–425.

Dorr, C. [2004]. Non-Symmetric Relations. *Oxford Studies in Metaphysics*, 1:155–192.

——— [2014]. Quantifier Variance and the Collapse Theorems. *The Monist*, 97(4):503–570.

——— [2016]. To Be F Is To Be G. *Philosophical Perspectives*, 30(1):39–134.

Dorr, C. and Goodman, J. [2020]. Diamonds Are Forever. *Noûs*, 54(3):632–665.

Feferman, S. [1964]. Systems of Predicative Analysis. *Journal of Symbolic Logic*, 29(1):1–30.

——— [2005]. Predicativity. In S. Shapiro, ed., *Oxford Handbook of Philosophy of Mathematics and Logic*, pp. 590–624. Oxford: Oxford University Press.

Fine, K. [1977]. Properties, Propositions and Sets. *Journal of Philosophical Logic*, 6(1):135–191.

——— [1994]. Essence and Modality. *Philosophical Perspectives*, 8:1–16.

——— [2000]. Neutral Relations. *Philosophical Review*, 109(1):1–33.

——— [2005]. Tense and reality. In *Modality and Tense: Philosophical Papers*, pp. 261–320. Oxford: Oxford University Press.

——— [2006a]. The reality of tense. *Synthese*, 150(3):399–414.

——— [2006b]. Relatively unrestricted quantification. In A. Rayo and G. Uzquiano, eds., *Absolute Generality*, pp. 20–44. Oxford University Press.

——— [2010]. Some Puzzles of Ground. *Notre Dame Journal of Formal Logic*, 51(1):97–118.

——— [2012]. Guide to ground. In F. Correia and B. Schnieder, eds., *Metaphysical Grounding*, pp. 37–80. Cambridge University Press.

——— [2016]. Williamson on Fine on Prior on the reduction of possibilist discourse. *Canadian Journal of Philosophy*, 46(4-5):548–570.

Fritz, P. and Goodman, J. [2016]. Higher-Order Contingentism, Part 1: Closure and Generation. *Journal of Philosophical Logic*, 45(6):645–695.

Fritz, P., Lederman, H., and Uzquiano, G. [2021]. Closed Structure. *Journal of Philosophical Logic*, 50(6):1249–1291.

Goodman, J. [2017]. Reality is Not Structured. *Analysis*, 77(1):43–53.

Goodman, N. [1951]. *The Structure of Appearance*. Cambridge, MA, USA: Harvard University Press.

Gödel, K. [1944]. Russell's Mathematical Logic. In P. A. Schilpp, ed., *The Philosophy of Bertrand Russell*, pp. 123–153. Evanston and Chicago: Northwestern University.

Harris, J. H. [1982]. What's so Logical About the "Logical" Axioms? *Studia Logica*, 41(2-3):159–171.

Hirsch, E. [2002]. Quantifier variance and realism. *Philosophical Issues*, 12:51–73.

——— [2010]. *Quantifier Variance and Realism: Essays in Metaontology*. New York: Oxford University Press.

Hodes, H. T. [2015]. Why Ramify? *Notre Dame Journal of Formal Logic*, 56(2):379–415.

Hoeksema, J. [1987]. The logic of exception. In *Proceedings of the Fourth Eastern States Conference on Linguistics (ESCOL 4)*, pp. 100–113.

Kaplan, D. [1989]. Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals. In J. Almog, J. Perry, and H. Wettstein, eds., *Themes From Kaplan*, pp. 481–563. Oxford University Press.

——— [1995]. A Problem in Possible Worlds Semantics. In W. Sinnott-Armstrong, D. Raffman, and N. Asher, eds., *Modality, Morality and Belief: Essays in Honor of Ruth Barcan Marcus*, pp. 41–52. Cambridge University Press.

Krämer, S. [2013]. A Simpler Puzzle of Ground. *Thought: A Journal of Philosophy*, 2(2):85–89.

Lewis, D. K. [1968]. Counterpart Theory and Quantified Modal Logic. *Journal of Philosophy*, 65(5):113–126.

——— [1980]. Index, Context, and Content. In S. Kanger and S. Öhman, eds., *Philosophy and Grammar*, pp. 79–100. Reidel.

——— [1986]. *On the Plurality of Worlds*. Oxford: Blackwell.

Linnebo, Ø. [2013]. The potential hierarchy of sets. *The Review of Symbolic Logic*, 6(2):205–228.

McTaggart, J. E. [1908]. The Unreality of Time. *Mind*, 17(68):457–474.

Mellor, D. H. [1981]. *Real Time*. Cambridge University Press.

Myhill, J. [1958]. Problems Arising in the Formalization of Intensional Logic. *Logique et Analyse*, 1(1):78–83.

Poincaré, H. [1906]. Les mathématiques et la logique, II, III. *Revue de métaphysique et de morale*, 14:17–34, 294–317.

Prior, A. N. [1957]. *Time and Modality*. Westport, Conn.: Oxford University Press.

——— [1961]. On a Family of Paradoxes. *Notre Dame Journal of Formal Logic*, 2(1):16–32.

——— [1967]. *Past, Present and Future*. Oxford, GB: Oxford University Press.

——— [1968]. Now. *Noûs*, 2(2):101–119.

Prior, A. N. and Fine, K. [1977]. *Worlds, Times and Selves*. Duckworth.

Quine, W. V. [1953]. Three Grades of Modal Involvement. *Proceedings of the XIth International Congress of Philosophy*, 14:65–81.

Quine, W. V. O. [1960]. *Word and Object*. Cambridge, MA, USA: MIT Press.

Richard, M. [1981]. Temporalism and Eternalism. *Philosophical Studies*, 39(1):1–13.
Russell, B. [1903]. *Principles of Mathematics*. Cambridge: Cambridge University Press.
——— [1906a]. Les paradoxes de la logique. *Revue de métaphysique et de morale*, 14:627–650.
——— [1906b]. On some difficulties in the theory of transfinite numbers and order types. *Proceedings of the London Mathematical Society*, s2-4(1):29–53.
——— [1908]. Mathematical Logic as Based on the Theory of Types. *American Journal of Mathematics*, 30(3):222–262.
——— [1915]. On the Experience of Time. *The Monist*, 25(2):212–233.
——— [1973]. *Essays in Analysis*. London, England: Allen & Unwin.
——— [2010]. *The Philosophy of Logical Atomism*. Routledge Classics. London: Routledge.
Russo, E. [Manuscript]. Absolute Generality as Higher-Order Identity.
Schaffer, J. [2012]. Necessitarian Propositions. *Synthese*, 189(1):119–162.
Sider, T. [2001]. *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
——— [2009]. Ontological Realism. In D. J. Chalmers, D. Manley, and R. Wasserman, eds., *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 384–423. Oxford University Press.
——— [2011]. *Writing the Book of the World*. New York: Oxford University Press.
Skolem, T. [1967]. Some Remarks on Axiomatized Set Theory. In J. van Heijenoort, ed., *From Frege to Gödel: a Sourcebook in Mathematical Logic*, pp. 290–301. Harvard University Press.
Spinoza, B. d. [1985]. *The Collected Works of Spinoza*. Princeton, N.J.: Princeton University Press. Ed. E. M. Curley.
Stanley, J. [2002]. Nominal restriction. In G. Preyer and G. Peter, eds., *Logical Form and Language*, pp. 365–388. Oxford University Press.
Stanley, J. and Szabó, Z. G. [2000]. On Quantifier Domain Restriction. *Mind and Language*, 15(2-3):219–261.
Stojanovic, I. [2012]. Domain-sensitivity. *Synthese*, 184(2):137–155.
Uzquiano, G. [2015]. A Neglected Resolution of Russell's Paradox of Propositions. *The Review of Symbolic Logic*, 8(2):328–344.
——— [2023]. Ramified structure. *Philosophical Studies*, 180(5–6):1651–1674.
van Inwagen, P. [1990]. *Material Beings*. Cornell University Press.
von Fintel, K. [1993]. Exceptive constructions. *Natural Language Semantics*, 1(2):123–148.
Westerståhl, D. [1985]. Determiners and context sets. In J. van Benthem and A. ter Meulen, eds., *Generalized Quantifiers in Natural Language*, pp. 45–71. Dordrecht: Foris.
Weyl, H. [1994]. *The Continuum: A Critical Examination of the Foundation of Analysis*. New York: Dover Publications.
Whitehead, A. N. and Russell, B. [1910]. *Principia Mathematica*. Cambridge University Press.
Williamson, T. [1985]. Converse Relations. *Philosophical Review*, 94(2):249–262.
——— [1988]. Equivocation and Existence. *Proceedings of the Aristotelian Society*, 88(1):109–128.
——— [2003]. Everything. *Philosophical Perspectives*, 17(1):415–465.

——— [2013]. *Modal Logic as Metaphysics*. Oxford University Press.

Wittgenstein, L. [1921]. *Tractatus Logico-Philosophicus*. New York: Routledge.

Yagisawa, T. [2010]. *Worlds and Individuals, Possible and Otherwise*. Oxford: Oxford University Press.

Yli-Vakkuri, J., Hawthorne, J., and Fritz, P. [2019]. Operator Arguments Revisited. *Philosophical Studies*, 176(11):2933–2959.

Zimmerman, D. W. [2005]. The A-Theory of Time, the B-Theory of Time, and 'Taking Tense Seriously'. *Dialectica*, 59(4):401–457.