Unrestricted Quantification and the Structure of Type Theory

SALVATORE FLORIO
University of Birmingham
s.florio@bham.ac.uk

NICHOLAS K. JONES
University of Birmingham
n.k.jones@bham.ac.uk

Draft of 22 November 2018

Abstract: Semantic theories based on a hierarchy of types have prominently been used to defend the possibility of unrestricted quantification. However, they also pose a *prima facie* problem for it: each quantifier ranges over at most one level of the hierarchy and is therefore not unrestricted. It is difficult to evaluate this problem without a principled account of what it is for a quantifier to be unrestricted. Drawing on an insight of Russell's about the relationship between quantification and the structure of predication, we offer such an account. We use this account to examine the problem in three different type-theoretic settings, which are increasingly permissive with respect to predication. We conclude that unrestricted quantification is available in all but the most permissive kind of type theory.

1 Introduction

The result of theoretical inquiry often depends on the framework in which theorising is conducted. Central features of any such framework are the linguistic resources regarded as legitimate, and the notions taken as primitive. Philosophical inquiry is no exception. To take a prominent example, one's ontological views may depend on whether primitive modal vocabulary is regarded as legitimate. Eliminating modal vocabulary naturally leads one to postulate a vast plenitude of concrete worlds (Lewis 1986), which can be avoided by permitting primitive modals (Prior and Fine 1977; Stalnaker 1976; Stalnaker 2012, especially Chapter 1). For another example, one's views about the existence and nature of properties may depend on whether higher-order quantification is regarded as legitimate. The Quinean view that only first-order quantification is legitimate naturally leads one to postulate the existence of immanent universals (Armstrong 1980; Lewis 1983, 20–25), which can be avoided by permit-

ting primitive quantification into predicate position (van Cleve 1994; Jones 2017). Permitting such higher-order quantification will also affect the result of one's semantic theorising. This paper concerns the influence of primitive higher-order quantification on the semantic thesis that unrestricted quantification is possible.

The possibility of unrestricted quantification appears to be presupposed by the absolutely general modes of inquiry arguably characteristic of metaphysics, logic, and mathematics. A framework capable of accommodating unrestricted quantification is thus needed to make reflective sense of these theoretical pursuits. One prominent candidate is provided by higher-order logic, and type theory more generally. It is well known that standard model-theoretic semantics, couched in set-theoretic terms, cannot accommodate unrestricted interpretations of the quantifiers because there is no universal set available to serve as an unrestricted domain of quantification. In contrast, the framework of higher-order logic makes available a novel notion of property compatible with the existence of a universal property of objects. This universal property can serve as a seemingly unrestricted domain of quantification (Williamson 2003). However, it is unclear whether this framework really is compatible with unrestricted quantification. The problem is that the universal property of objects may not be truly universal. For surely a truly universal property would apply not only to all the objects, but also to any other entities recognised by the framework. Yet no property is universal in that sense. Call this the *intuitive problem for unrestricted quantification within type theory*, or the *intuitive problem* for short.

One can see this as a revenge problem affecting type-theoretic semantics for unrestricted quantification. Like other revenge problems, it arises because the theoretical resources used to resolve a problem allow us to formulate a new instance of that very problem. In the present case, the original problem was that no set-theoretic domain of quantification contains every object. The new problem is that no type-theoretic domain contains every type-theoretic entity. So one wonders whether anything has been gained by adopting the type-theoretic framework.

Is the intuitive problem sound? Our aim in this paper is to clarify and answer this question. We argue that the answer depends on the precise structure of predication. To this end, we consider three kinds of type theory, which are increasingly permissive with respect to predication. It turns out that the intuitive problem is sound in only the most permissive kind of theory. Our argument develops an insight of Bertrand Russell's about the structure of quantification. We use this insight to explicate unrestricted quantification, and to offer a systematic and principled analysis of the relationship between the intuitive problem and the structure of predication.

2 Unrestricted quantification within type theory

This section introduces the type-theoretic framework with which we are concerned, and its role in the semantics of unrestricted quantification.

We will be considering languages with the following kind of type-theoretic structure. There is an infinite hierarchy of syntactic categories, or *types*. Type 0 is the type of singular terms. Type 1 is the type of predicates that form sentences when combined with expressions of type 0. Type 2 is the type of predicates that form sentences when combined with expressions of type 1. And so on. Importantly for us, there will be variables of each type, all of which can be bound by quantifiers.

These languages extend the familiar language of first-order logic. Singular terms and variables like 'a' and 'x' are of type 0. Predicates like 'F' are of type 1, since they form sentences when combined with expressions of type 0, e.g. '$F(a)$'. The type-theoretic languages we are interested in go beyond first-order logic by introducing new kinds of *predicative* expressions (i.e. predicates of predicates, predicates of those predicates, and so on) and the means of quantifying on them. The result is a system of *higher-order logic*.

A word on notation. In first-order logic, it is convenient to let the upper-case/lower-case distinction indicate type: upper-case for predicates, lower-case for terms. We will occasionally employ this convention. However, since we will consider languages with many more than two types, we need something more systematic. So we use superscripts to indicate type. For example, 'a^1' is a constant of type 1, whereas 'x^2' is a variable of type 2, and '$x^2(a^1)$' is an open sentence which combines them.

First-order logic, understood as a regimentation of a fragment of English, is widely regarded as a legitimate language of theorising. According to the type-theoretic framework we are interested in, languages with the structure just described can also be legitimate languages of theorising and do not require reductive explanation. That is, taken at face value, they are in perfectly good standing. This primitivist attitude towards type-theoretic languages is a rejection of W.V.O. Quine's (1970, 66–68) view that higher-order languages must be explained in set-theoretic terms (see, e.g., Prior 1971, Chapter 3; Boolos 1975; Rayo and Yablo 2001; Williamson 2003; Hale 2013, Chapter 8).

Why would one adopt this type-theoretic framework? First and foremost, because it is useful. While the framework has been widely applied in the foundations of mathematics, recent developments have shown that it holds great promise elsewhere. Notably, it has found fruitful applications in the metaphysics of properties, relations, propositions, identity, and modality (e.g. Dunaway 2013; Williamson 2013; Goodman 2017; Dorr 2016; Jones 2017). It has also proven fruitful in the foundations of semantics, where it has been used to to disentangle logic from mathematics (Florio and Incurvati 2018a; Florio and Incurvati 2018b), and to characterise the intended model of the language of set theory (e.g. Boolos 1985; Shapiro 1991, Chapter 5). This last application is particularly relevant to the semantics of unrestricted quantification, to which we turn shortly.

Given a type-theoretic language understood as we have described, one can use it to *introduce* a metaphysical hierarchy of entities. At the base of the hierarchy,

level 0, are the objects. To say that there are objects, one uses existential quantifiers binding variables of type 0, e.g. $\exists x^0(\ldots x^0 \ldots)$.[1] At the next level up, level 1, are the properties of objects. To say that they exist, one uses existential quantifiers binding variables of type 1, e.g. $\exists x^1(\ldots x^1 \ldots)$. Level 2 contains properties of level 1 properties. To say that they exist, one uses existential quantifiers binding variables of type 2, e.g. $\exists x^2(\ldots x^2 \ldots)$. And so on. Finally, one says that an entity has a property by predicating the property of the entity, e.g. $a^1(b^0)$ or $a^2(b^1)$.

Importantly for us, this description of a metaphysical hierarchy of objects and properties should not be understood as relying on an antecedent conception of such entities. In line with the primitivist interpretation of type-theoretic languages assumed here, our talk about different levels should be understood in terms of the different orders of quantification available in this setting, rather than conversely. We thereby legitimise a familiar and natural way of talking, making our exposition go more smoothly. But it is important to remember the sense we have given to talk about objects and properties, which may not always legitimise familiar ways of using these notions.

Canonical set-theoretic semantics embodies a Quinean attitude towards this metaphysical hierarchy and the type-theoretic language used to introduce it. Taking the first-order language of set theory as given, this approach interprets quantifiers of all types as ranging over sets. As a result, type-theoretic languages provide no new conceptual resources beyond the background set theory used to interpret them. In effect, the metaphysical hierarchy is revealed as merely a hierarchy of sets. Our rejection of the Quinean attitude makes room for a different approach. Taking a type-theoretic language and its corresponding metaphysical hierarchy as given makes a novel suite of conceptual resources available. These resources can then be used in our semantic theorising. In particular, they can be used to develop a semantic theory that seems able to accommodate the possibility of unrestricted quantification.

To make reflective sense of this possibility, we require a semantic theory in which the quantifiers can have an unrestricted interpretation. For simplicity, we let our object language be a first-order language of the usual sort. An interpretation of this language has two components. One interprets the non-logical vocabulary. We assume this component to be fixed. The other provides a domain of quantification, which we allow to vary. So our goal will be to characterise the truth conditions of quantified sentences relative to arbitrary domains of quantification: sentence s is true in domain d. To avoid complications, we focus on the simplest generalizations possible in the language, namely sentences like $\forall v^0 F^1(v^0)$, which we call *basic generalisations*.[2]

[1] Henceforth, we typically omit quotation marks in the interests of readability. We will also indulge in some harmless use/mention confusion when it simplifies exposition. Context will serve to disambiguate.

[2] Our focus on basic generalisations does not narrow the scope of our investigation. This is because

(Henceforth, we omit type indices in sentences of the first-order object language.)

To a first approximation, the truth conditions for basic generalisations can be stated as follows:

> $\forall v F(v)$ is true in d if and only if, for each x in d, F applies to x.

The precise regimentation of this clause, including type indices, depends on the theoretical framework in which one is working. For example, in standard model-theoretic semantics all metalinguistic quantifiers are first-order (i.e. they bind variables of type 0). The truth conditions for basic generalisations are then regimented as:

> $\forall v F(v)$ is true in d^0 if and only if, for each object x^0 in d^0, F applies to x^0.

Notice that this sort of regimentation treats domains as objects (level 0), usually sets, though different implementations may employ different kinds of objects.

Is unrestricted quantification possible in this framework? A necessary condition is the existence of a domain containing all objects. Yet, if domains are objects, the following domain-theoretic principle of separation entails that no domain contains all objects:

> For any domain d, there is a domain d^* comprising all and only the objects in d that satisfy $\varphi(x^0)$, where $\varphi(x^0)$ is any formula of one's language of theorising.[3]

The incompatibility of separation with a domain of level 0 containing all objects is one lesson of Russell's paradox. If there is such a d^0, then there is, by separation, another domain r^0 containing all and only the objects that do not contain themselves. But since r^0 is itself an object, its existence is classically inconsistent: it contains itself if and only if it does not.

Adopting a type-theoretic framework resolves this problem by making available a novel notion of property compatible with the existence of a universal property of objects. The key move is to use expressions of type 1 to regiment domain talk, as in this regimentation of the truth conditions for basic generalisations:

---

complex generalisations can be expressed by means of basic ones. Consider a complex generalisation of the form $\forall v^0 \varphi(v^0)$, where f^1 is the property determined by $\varphi(v^0)$. By interpreting F as standing for f^1, we can express the complex generalisation $\forall v^0 \varphi(v^0)$ as the basic generalisation $\forall v^0 F(v^0)$.

[3] A full justification of the the domain-theoretic principle of separation is beyond the scope of this article. However, one may start from the following simple observation. In ordinary conversation, we sometimes interpret our interlocutor's domain of quantification as progressively narrowing. The domain-theoretic principle of separation merely encodes the following plausible generalisation about how this narrowing may proceed. We begin by interpreting our interlocutor as employing a given domain. When a certain condition becomes salient, we then interpret our interlocutor as employing a restriction of that domain to things that satisfy the condition.

> $\forall v F(v)$ is true in d^1 if and only if, for each object x^0 such that $d^1(x^0)$, F applies to x^0.[4]

Above, we identified a necessary condition on the possibility of unrestricted quantification, namely the existence of a domain containing all objects. Given this regimentation of domain talk, it follows from a standard principle of comprehension that there is such a domain, i.e. a property u^1 such that $u^1(x^0)$ for every object x^0.[5] By replacing d^1 with u^1, we obtain from the present regimentation an account of the truth conditions for basic generalisations in this seemingly unrestricted domain of objects. Because the paradoxical reasoning based on separation requires that domains are objects, and u^1 is not an object, its existence is compatible with separation. More specifically, that reasoning requires asking whether a certain domain contains itself, which one cannot meaningfully do if domains are properties of level 1. To ask whether a domain of level 1 contains itself, one would need to use a formula of the form $d^1(d^1)$, which is ungrammatical within standard forms of type theory (but see §6 for an alternative).

We appear to have vindicated the possibility of unrestricted quantification. At this point, however, the intuitive problem rears its ugly head. Although there is a property u^1 possessed by all the objects, surely a truly universal domain would also be possessed by all properties. But, within the type-theoretic framework, one cannot grammatically say that a property possesses another property (of the same or higher level), never mind show that some property is possessed by *all* properties.[6]

So, is quantification over u^1 really unrestricted? Some have claimed so (e.g. Rayo and Williamson 2003; Williamson 2003; Rayo 2006). Others, however, have claimed not. For example, Leonard Linsky (1992, 262) makes the following general claim:

> the theory of logical types [...], in all of its many varieties, has at its center a rejection of this unrestricted notion of *everything*.

A similar view is suggested by remarks of Øystein Linnebo and Agustin Rayo (2012, especially 292–293), who consider arguments from the open-endedness of the type-theoretic hierarchy to open-endedness of the set-theoretic hierarchy. More recently, Stephan Krämer (2017) argues that a specific kind of type theory is incompatible with unrestricted quantification; we discuss this argument in §5.

---

[4]Note that, since the type of the domain has changed from 0 to 1, the type of the predicate 'true in' will have to change too, to accept expressions of type 1 in its second argument.

[5]The relevant instance of comprehension is: $\exists y^1 \forall x^0 (y^1(x^0) \leftrightarrow x^0 = x^0)$.

[6]Within the type-theoretic framework, to say that a property has a property of the same or higher level, one would need to use a formula like $a^1(b^1)$ or $a^1(b^2)$ where the subject has the same or higher type than the predicate. Those formulae are ungrammatical in standard forms of type theory, though we discuss an alternative later (§6).

It is difficult to make progress on this issue without a principled account of what it is for a domain to be unrestricted. In particular, we need an account that is applicable in a range of different settings, including alternative developments of the type-theoretic framework introduced above. Important as it is to our understanding of unrestricted quantification, such an account is presently absent from the literature. In the next section, we address this lacuna by identifying the core theoretical role governing unrestricted quantification and offering a positive account of its occupier.

Our account will reveal an intimate connection between unrestricted quantification and the structure of meaningful predicability. We will employ the following methodological assumption in our investigation. According to the theoretical framework assumed here, type-theoretic languages are legitimate languages of theorising and not in need of further reductive explanation. In keeping with this approach, we take syntactic restrictions on predication to line up with meaningful predicability. This will allow us to draw conclusions about meaningful predicability, and hence unrestricted quantification, from premises about what kinds of predication count as well-formed in different developments of the framework. In short, we interpret the type-theoretic structure as encoding the limits of meaningful thought and talk.

3 Russellian domains

Unrestricted quantification is quantification over an unrestricted domain. So we can explicate unrestricted quantification by explicating unrestricted domains. We do so by identifying the theoretical role governing unrestricted domains and, drawing on an insight of Russell's, offering *Russellian domains* as occupants of this role. Subsequent sections apply our proposal within three different kinds of type theory.

We begin with the following natural conception of universal quantification as preclusion of counterexamples: for a universal generalisation to be true in a domain is for that domain to lack counterexamples. For example, for the generalisation 'everything is F' to be true in a domain d is for d to lack counterexamples, i.e. to be free from non-Fs. When the domain is unrestricted, true universal quantification over it precludes there from being absolutely any counterexamples whatsoever. We take this to be *the* distinctive feature of unrestrictedness, which delivers the following theoretical role: an unrestricted domain is a domain such that true universal quantification over it precludes there from being absolutely any counterexamples whatsoever.

Moreover, we can argue that this role is correct by appeal to the initial motivations for unrestricted quantification, namely making reflective sense of the absolutely general modes inquiry arguably characteristic of metaphysics, logic, and mathematics.[7] By way of example, consider the physicalist thesis that everything is physical. The

[7] Williamson (2003, 415-416) employs this kind of motivation for unrestricted quantification.

metaphysically interesting version of this thesis is refuted by anything non-physical whatsoever, which requires quantification over a domain d that is unrestricted in the sense of our theoretical role: true universal quantification over d precludes there from being any counterexamples whatsoever.

What must a domain be like to occupy this theoretical role? Our key idea is inspired by this passage from Russell 1908 (236–237):

> Every proposition containing *all* asserts that some propositional function is always true; and this means that all values of the said function are true, not that the function is true for all arguments, since there are arguments for which any given function is meaningless, *i.e.*, has no value. Hence we can speak of *all* of a collection when and only when the collection forms part or the whole of the *range of significance* of some propositional function, the range of significance being defined as the collection of those arguments for which the function in question is significant, *i.e.*, has a value. [...] A *type* is defined as the range of significance of a propositional function, *i.e.*, as the collection of arguments for which the said function has values.

For Russell, quantifiers are operators on certain worldly items he calls propositional functions, i.e. functions from entities to propositions. It is useful in the present context to adopt a more linguistic interpretation.[8] Focussing on the case of basic generalisations, we can identify propositional functions with predicates. The range of significance of a predicate F then comprises the things of which F can be meaningfully predicated, i.e. the things that can be meaningfully said to be F.

In the quoted passage Russell claims that, for every propositional function, some entities lie outside its range of significance. The range of significance of F comprises whatever can be meaningfully said to be F. So Russell's claim amounts to the following: for every F, some entity cannot be meaningfully said to be F, i.e. there is no such proposition as the proposition that it is F.

We can now state Russell's key insight about the structure of quantification: sensible quantification never goes beyond the range of significance. As he puts it, "we can speak of *all* of a collection when and only when the collection forms part or the whole of the *range of significance* of some propositional function."

Russell describes this as a form of restricted quantification. Unlike more familiar forms of quantifier restriction, however, "this is an *internal* limitation upon x, given by the nature of the function; and it is a limitation which does not require explicit statement, since it is impossible for a function to be true more generally than for all its values." (Russell 1908, 234). In present terminology, it is impossible for a

[8] We also thereby bypass certain exegetical worries concerning the compatibility of propositional functions with Russell's multiple relation theory of judgment.

predicate F to be true or false of things that cannot be meaningfully said to be F. So the domain of quantification of $\forall v F(v)$ must be included in the range of significance of F.[9] This is Russell's internal limitation on quantification.

We embrace this Russellian insight about the structure of quantification, but deny that his internal limitation always yields a restricted form of quantification. Properly understood, what Russell described as a restriction on quantification need really be no restriction at all. When the domain coincides with the whole range of significance, the domain should not be seen as restricted.

We can make this precise as follows. Define a domain as *Russellian* for a basic generalisation $\forall v F(v)$ if, and only if, it coincides with the range of significance of F. We contend that Russellian domains play the unrestricted domain role characterised above. So our main thesis is:

(R=U) Unrestricted domains are all and only Russellian domains.

According to (R=U), a domain d is unrestricted for $\forall v F(v)$ if and only if d is Russellian for $\forall v F(v)$.

Before continuing, let us pause to consider two concerns that might arise. According to the first, this is a terminological issue. After all, unrestrictedness is a theoretical notion. Provided we are clear about what we mean by unrestrictedness, we can use the notion as we please. Consequently, (R=U) is at most definitional, not a substantive thesis.

This worry is misplaced. The issue is substantive because of the theoretical role for unrestrictedness independently argued for above. Our claim is that all and only Russellian domains satisfy that role, which is clearly not a definitional matter.

The second concern is that our main thesis cannot be expressed within the type-theoretic framework. For example, recall how we characterised ranges of significance: "the range of significance of F comprises whatever can be meaningfully said to be F". To have its intended effect, the quantifier "whatever" here must range over entities of all levels, not just those of some particular level. But such quantification across types is not permitted within the framework. The concern is thus that we cannot properly characterise ranges of significance, and so we cannot express our main thesis (R=U).

This is an old and familiar problem facing type-theoretic approaches. A version of the problem was driving Frege's (1892) notorious paradox of the concept *horse*, and was also recognised by Wittgenstein (1922, e.g. Propositions 3.331 and 4.1241) and Gödel (1944, 466). The problem arises from a mismatch of perspectives. The

[9] One way to see this is as follows. Given a domain d, the basic generalisation $\forall v F(v)$ expresses a proposition whose truth requires, for each x in d, the truth of the singular proposition that x is F. When d extends beyond the range of significance of F, it contains some thing for which there is no such singular proposition. So $\forall v F(v)$ does not express a proposition when "interpreted" over d: sensible quantification never goes beyond range of significance.

first is the external perspective of someone attempting to describe the type-theoretic framework from the outside, aiming to understand its workings. The second is the internal perspective of someone who thinks, speaks, and works within the system. The problem arises from attempting to reformulate the external perspective on the system within the system itself. The expressive limitations currently at issue constrain the extent to which the external perspective is available within the system.

Although this is clearly a deep and difficult issue, it is not obvious what to make of this mismatch between perspectives. On the one hand, external theorising is central to making sense of, and evaluating, theoretical options based on alternative conceptual schemes. On the other hand, it would be nice if the resultant theoretical considerations could be appreciated by all parties, regardless of perspective. Alas, situations often arise in philosophy where these desiderata cannot be jointly satisfied. The present case is one such situation.

We cannot address this issue fully here, and so we simply pause for two observations. Firstly, this is a very general problem for those who wish to make reflective sense of the type-theoretic framework. In this respect, our proposal is no worse off than many others in the area. Secondly, various techniques have been proposed for reducing the system's expressive limitations, many of which could be employed here.[10] With these issues set aside, we now argue for (R=U).

First, we show that every Russellian domain is unrestricted. Suppose that d is Russellian for $\forall v F(v)$. By definition of Russellian, F is not meaningfully predicable of anything outside d. So nothing outside d can even count as a counterexample to $\forall v F(v)$. An entity can count as a counterexample only if it can be meaningfully said not to be F. But what cannot be meaningfully said to be F cannot be meaningfully said not to be F either. Then, because d is Russellian for $\forall v F(v)$, it contains everything that can count as a counterexample to the generalisation. So the truth of $\forall v F(v)$ in d precludes there from being any counterexamples whatosever. By the theoretical role for unrestrictedness, therefore, d is unrestricted for $\forall v F(v)$.

Now we argue for the converse: every non-Rusellian domain is restricted. Suppose d is non-Russellian for $\forall v F(v)$. Then, by definition of Russellian, either (a) F is not meaningfully predicable of something in d, or (b) F is meaningfully predicable of something x outside d. In case (a), the domain is not included in the range of significance of F, which, as Russell observed, is impossible for meaningful quantification. In case (b), x can count as a counterexample to $\forall v F(v)$ and hence the truth of $\forall v F(v)$ in d doesn't preclude x from being a counterexample. It follows by the

[10] In place of generalisation across levels, one can employ schematic generality and typical ambiguity, or infinitary conjunction and disjunction. One can also adopt a metalinguistic stance on the type-theoretic language and invoke ungrammaticality of certain strings (and those used to "interpret" them) in place of talk about meaninglessness. For recent discussions of related issues, see e.g. Wright 1998, Glanzberg 2004, Lavine 2006, Linnebo 2006, Hale 2013, Krämer 2013, Hale and Wright 2012, Jones 2016, and Florio and Linnebo 2019 (Chapters 11 and 12).

theoretical role for unrestrictedness that d is not unrestricted for $\forall v F(v)$.

Returning to the intuitive problem, consider u^1, the universal property of objects. Is quantification over u^1 really unrestricted? Assuming (R=U), this is a question about whether u^1 is a Russellian domain for $\forall v F(v)$: does u^1 coincide with the range of significance of F? The following sections examine this question within three different kinds of type theory. As we will argue, the answer depends on the precise structure of meaningful predicability.

4 Strict type theory

The most familiar form of type theory embodies a very restrictive conception of predication: a predication is well-formed if and only if the type of its predicate immediately succeeds the type of its subject. This section argues that, in this setting, u^1—a property of level 1 possessed by all and only objects—is indeed an unrestricted domain for basic generalisations in a first-order object-language. We begin with a more precise description of this *strict type theory*, keeping it as simple as present purposes require.

Elaborating on the outline in §2, the language has an infinite hierarchy of types indexed by the natural numbers. There are *monadic* variables and constants of each type, alongside the usual logical operators. There is also an intra-type identity predicate for each type. Quantifiers can bind variables of any type. The distinctive feature of strict type theory is its treatment of predication: $s^i(t^j)$ is well-formed if and only if $i = j + 1$.

We assume the standard, classical rules for the logical operators. In addition, we have a principle of comprehension for each type $i \geq 1$:

$$\exists x^i \forall y^{i-1} \big(x^i(y^{i-1}) \leftrightarrow \varphi(y^{i-1}) \big)$$

subject to the usual proviso that x^i does not occur in φ. We also have a form of Leibniz's Law for each type i:

$$\forall x^i \forall y^i \big(x^i =_{i+1} y^i \leftrightarrow \forall z^{i+1} \big(z^{i+1}(x^i) \leftrightarrow z^{i+1}(y^i) \big) \big)$$

As presented, the system lacks relations. This will greatly simplify the discussion of cumulativity in the next section. Nothing of substance turns on this omission of relations because one can simulate them within the present system, given appropriate assumptions. In particular, binary relations (including cross-level relations) can be identified with properties of pairs.[11]

[11] For definitions of cross-level pairing in an extensional setting, see Linnebo and Rayo 2012, Appendix B2, and Florio and Linnebo 2019, Appendix 11.B. It is worth emphasising that these complications are required only because we have omitted relational types in order to simplify the discussion

Let us now examine the intuitive problem within this strict system. The idea driving the problem was that a truly universal property would apply not only to all the objects, but also to any other entities recognised by the framework. So, in particular, (a) a property is an unrestricted domain only if it is possessed by all properties of level 1, and (b) although u^1 is possessed by all objects, it is not possessed by any property of level 1. Therefore, u^1 is not an unrestricted domain.

The strict system avoids this problem because neither (a) nor (b), let alone the more general idea driving the problem, is expressible without violating the strict type restrictions on well-formed predication. For example, one would express the claim that u^1 is not possessed by any property of level 1 with the formula $\forall x^1 \neg u^1(x^1)$. Yet this violates the requirement that the type of the subject be immediately lower than the type of the predicate. Given our assumption that syntactic restrictions on predication align with meaningful predicability, it might therefore appear that the intuitive problem is merely an expressive illusion.

However, matters are not so simple. A closely related problem is expressible within the strict system. As observed above, although the system lacks relations, one can simulate them within it. This allows us to make sense of the claim that the properties of level 1 outnumber the objects. Moreover, the regimentation of this claim is provable, yielding a type-theoretic version of Cantor's theorem.[12] At a first pass, this appears to be incompatible with u^1 being unrestricted. We now use (R=U) to argue that this appearance is misleading.

According to (R=U), a domain d is unrestricted for a basic generalisation $\forall v F(v)$ if and only if d is Russellian for that generalisation, i.e. coincides with the range of significance of F. What is this range? It depends how our first-order object language is interpreted. Earlier, we assumed a fixed interpretation of its non-logical vocabulary. This is an interpretation of the usual sort, namely one in which the quantifiers range over objects and the predicates express properties of level 1. So let f^1 be the property expressed by F. By the strict type restriction on predication, f^1 is grammatically predicable of all and only objects. Since we are taking syntactic restrictions on predication to line up with meaningful predicability, F is meaningfully predicable of all and only objects. So the range of significance of F comprises the objects. Since all and only objects possess u^1, u^1 is Russellian, hence unrestricted, for $\forall v F(v)$.

Cantor's theorem provides a sense in which the properties outnumber the objects, hence also a sense in which some entities are outside u^1. But the theorem does not entail anything about meaningful predicability. So by (R=U), it does not entail anything about unrestrictedness. In fact, in the strict type theory, u^1 coincides with

of cumulativity in §5.

[12]The claim that the properties of level 1 outnumber the objects is regimented as the claim that no relation from the objects to the properties is both functional and onto. For a proof, see (Shapiro 1991, 103–104).

the range of significance of every predicate of our first-order object language. So, as argued just above, u^1 is an unrestricted domain for every basic generalisation $\forall v F(v)$ of our object language.

The point extends to basic generalisations in higher-order object languages. An object language predicate of type i is interpreted by a property of level i, and is therefore meaningfully predicable of all and only entities of level $i-1$. Let u^i be the property possessed by all entities of level $i-1$. Then u^i is Russellian, hence unrestricted, for $\forall v^{i-1} F^i(v^{i-1})$. The upshot is that the strict type theory is hospitable to unrestricted quantification of every order.

5 Cumulative type theory

Strict type theory encodes a very restrictive conception of predication. However, one can make good mathematical sense of a more permissive approach: a predication is well-formed if and only if the type of the subject is lower than the type of the predicate.[13] This means that predication of all lower types is now permitted, not just of the immediately preceding type. So one can say, for example, that a property is possessed by entities of different levels, e.g. $F^2(a^1) \wedge F^2(b^0)$. This kind of *cumulative type theory* has been recently discussed in favorable terms by Wolfgang Degen and Jan Johanssen (2000), and Linnebo and Rayo (2012). Timothy Williamson (2013, 237–238) also employs cumulativity in the semantics of higher-order modal languages.

The cumulative system considered here is just like the strict system except for its treatments of predication and comprehension. As just mentioned, a predication $s^i(t^j)$ is well-formed if, and only if, $i > j$ (rather than $i = j+1$). As for comprehension, there are various options. A natural idea is that one can define a property of a given level i by specifying, for each level below i, which entities of that level possess it. Formally,

$$\exists x^i \big(\forall y^{i-1}(x^i(y^{i-1}) \leftrightarrow \varphi_{i-1}(y^{i-1})) \wedge \ldots \wedge \forall y^0(x^i(y^0) \leftrightarrow \varphi_0(y^0))\big)$$

Note that this allows one to choose different formulae for different levels. A weaker variant requires greater uniformity in the choice of formulae (e.g. Linnebo and Rayo 2012, 288).

Why adopt this version of type theory? First, there is no obvious reason why the type theory *has* to be strict, given that we can make good mathematical sense of cumulativity. Second, it has been argued that we should embrace infinite types (Linnebo and Rayo 2012), which naturally delivers cumulativity. Since no type im-

[13] Michael Potter (2002, §5.12) points out that Ralph Hawtrey suggested this approach to Russell in a letter of 1908.

mediately precedes the first infinite type ω, the strict system permits no predications of the form $s^{\omega}(t^{i})$, whereas the cumulative system permits them whenever i is finite. Although we remain officially neutral about the ultimate success of these arguments, they clearly motivate taking cumulativity seriously.

Let us now examine the intuitive problem within this cumulative system. The simplest version we have seen so far rested on two claims: (a) a property is an unrestricted domain only if it is possessed by all properties of level 1; (b) although u^1 is possessed by all objects, it is not possessed by any property of level 1. It follows that u^1 is not an unrestricted domain. Just like in the strict system, however, neither claim is well-formed in the cumulative system, and so there is no problem here to respond to.

Another version of the intuitive problem has been endorsed in a recent paper by Krämer (2017). In the context of a debate between defenders of unrestricted quantification (Absolutists) and their opponents (Relativists), Krämer writes that, within cumulative type theory:

> there is more than is dreamt of by [the first-order object language's] quantifiers—what sets the picture apart from the Relativist's is only that what is more is not in the range of the semantics' first-order quantifiers but their second-order cousins. So the [type-theoretic] version of Absolutism is not really worthy of the name [...] it preserves the letter of Absolutism, but [...] gives up on its spirit. (Krämer 2017, 13)

Krämer's argument proceeds by considering a certain sentence which says that some properties of level 1 are not objects; he calls this sentence (MORE). Krämer then argues as follows:

(1) In the cumulative system, there are cross-type identity relations.

(2) If there are cross-type identity relations, then (MORE) is expressible and provable (hence true).

(3) If (MORE) is true, then the quantifiers of the first-order object language are not unrestricted (because (MORE) is contrary to "the spirit of Absolutism").

As we will explain below, one can see this as a version of the Cantorian argument discussed in the previous section, where the role previously played by a cardinality comparison between objects and properties is now played by (MORE). We will argue that our response to the Cantorian argument carries over to this case. Specifically, (MORE) is consistent with unrestricted quantification, assuming R=U. We begin by explaining Krämer's premises (1) and (2).

To see why the cumulative system contains cross-type identity relations, note that they are definable using the resources of the system.[14] Both the strict and cumulative systems have intra-type identity relations characterized by indistinguishability within the next level up:

$$x^i =_{i+1} y^i \leftrightarrow \forall z^{i+1}(z^{i+1}(x^i) \leftrightarrow z^{i+1}(y^i))$$

The cumulative system can also define cross-type identity relations characterized by indistinguishability within the first type predicable of the relata.[15] Where $k = \max(i, j) + 1$, we can regiment this as:

$$x^i \equiv_k y^j \leftrightarrow \forall z^k(z^k(x^i) \leftrightarrow z^k(y^j))$$

This generalises intra-type identity because, when x and y are of the same type i, the first type predicable of both is the next one up, i.e. $k = i + 1$.

We can use cross-type identity to formally express Krämer's (MORE) within the cumulative system:

(MORE) $$\exists x^1 \forall y^0 (x^1 \not\equiv_2 y^0)$$

This says that some property of level 1 is cross-type distinct from every object, i.e. some property is not an object. This is not only expressible but provable as a theorem of the cumulative system. Perhaps the easiest way to see why is as follows. Suppose for *reductio* that (MORE) is false. Then every property (of level 1) is cross-type identical to an object. So cross-type identity is a relation from level 0 to level 1 that is both functional and onto. But this contradicts the type-theoretic version of Cantor's theorem discussed in the previous section. Since the cumulative system extends the strict one, the type-theoretic version of Cantor's theorem is available there too. So by *reductio*, (MORE) is true. Given that this reasoning employs only logical rules of the cumulative system, (MORE) is a theorem of that system.

This suggests that (MORE) raises no new problems for unrestricted quantification beyond the type-theoretic version of Cantor's theorem. And this is borne out by

[14]There would be no need to define cross-type identity explicitly if we were working in a relational system that treated all predicates cumulatively, including identity. Because identity has a privileged relational status in our chosen monadic system, we did not assign a type to the identity sign. Making the types cumulative therefore does not force identity to become cumulative too.

[15]See Degen and Johannsen 2000, 149–150, and Linnebo and Rayo 2012, 282. Krämer (2017, 17–18) has a slightly more complex formulation because his system has primitive relational types.

It is not obvious that any of these definitions captures a form of identity. Identity should satisfy this condition: if x is identical to y, then anything meaningfully predicable of x is meaningfully predicable of y. But it is unclear whether that condition is satisfied by indistinguishability within the first level predicable of both x and y, hence also unclear whether the definition captures a genuine form of identity. For argument's sake, however, we grant that cross-type identity as defined really is a form of identity. This concession only strengthens Krämer's case.

closer examination: a version of our response to that problem within the strict system is also available within the cumulative one.

Recall, the key issue is whether u^1—the property of level 1 possessed by all and only objects—is an unrestricted domain for a basic generalisation of the first-order object language, such as $\forall v F(v)$. It follows from (MORE) that some property is not in this domain because that property is not identical to any object. But given R=U, this is silent about whether u^1 is unrestricted. According to R=U, a domain is unrestricted for $\forall v F(v)$ if, and only if, it coincides with the range of significance of F. Although (MORE) entails that some entity lies outside u^1, it does not entail that F is meaningfully predicable of that entity. So (MORE) does not entail that the range of significance of F extends beyond u^1. It therefore does not entail that u^1 is a restricted domain for $\forall v F(v)$. In a nutshell, (MORE) provides new entities outside u^1, but not new candidate counterexamples to $\forall v F(v)$. It therefore does not conflict with u^1 being unrestricted.

In fact, it follows from R=U that u^1 is unrestricted for $\forall v F(v)$. The reason is just the same as in the strict system. In brief, since F expresses a property of level 1, it is meaningfully predicable of all and only objects. Since all and only objects possess u^1, u^1 is Russellian, hence unrestricted, for $\forall v F(v)$.[16]

The point extends to basic generalisations in higher-order object languages. An object language predicate of type i is interpreted by a property of level i, and is therefore meaningfully predicable of all and only entities of levels below i. As before, let u^i be the property possessed by all entities of level $i-1$. Then a Russellian, hence unrestricted, domain for $\forall v^{i-1} F^i(v^{i-1})$ is a property that combines every property u^j where j is below i.[17]

The upshot is that cumulative type theory is hospitable to unrestricted quantification of every order, and (MORE) is compatible with Absolutism, namely the thesis that unrestricted quantification is possible. This undermines the following premise of Krämer's version of the intuitive problem:

(3) If (MORE) is true, then the quantifiers of the first-order object language are not unrestricted (because (MORE) is contrary to "the spirit of Absolutism").

Given the theoretical role of unrestrictedness identified in §3, the spirit of Absolutism

[16]The argument assumes that F expresses a property of level 1. Different interpretations are available. For example, suppose F expresses a property of level i. Then F is meaningfully predicable of all and only entities from levels below i. So an unrestricted domain for $\forall v F(v)$ is not u^1 but a property possessed by every entity whose level is below i (see footnote 17). In short, the existence of an unrestricted domain for $\forall v F(v)$ is independent of how F is interpreted.

[17]More precisely, a Russellian domain for $\forall v^{i-1} F^i(v^{i-1})$ is any property a^i such that

$$\forall x^0 a^i(x^0) \wedge \forall x^1 a^i(x^1) \wedge \ldots \wedge \forall x^{i-1} a^i(x^{i-1})$$

The existence of such a property follows from comprehension.

requires a domain containing all the candidate counterexamples to any given generalisation. We have just seen that, within the cumulative system, there are domains containing all candidate counterexamples in the strong sense of containing everything that can be meaningfully said to be a counterexample.

Let us take stock. The strict and cumulative systems both permit unrestricted quantification. They differ over which domains count as unrestricted for which generalisations. And they do so because they differ over the range of significance they assign to predicates of a given type i: all levels below i in the cumulative system, but only level $i-1$ in the strict system. An unrestricted domain for $\forall v^{i-1} F^i(v^{i-1})$ is thus u^i in the strict system, whereas it is a property combining all of $u^1, \ldots, u^i$ in the cumulative system.

Within the cumulative system, a version of (MORE) is provable for any types i, j where $i > j$:

$$\exists x^i \forall y^j (x^i \not\equiv_{i+1} y^j)$$

This shows, in effect, that every domain can be expanded. The mistake underlying Krämer's argument was to treat this kind of expandability as incompatible with unrestrictedness. However, the question isn't whether the domain is expandable. The question is whether it's expandable with new candidate counterexamples to a given generalisation, that is, whether the expanded domain is contained within the appropriate range of significance. The domain expansion arising from (MORE) is irrelevant to that question. By combining R=U with cumulativity, we see that unrestricted domains may also be expandable.

6 Liberal type theory

The final form of type theory considered here adopts a maximally permissive approach to predication: $s^i(t^j)$ is well-formed for all i, j. In this kind of *liberal type theory*, one can now say that a property is possessed by entities of the same or higher level than it, e.g. $F^1(a^2)$. Following Kurt Gödel (1933, 46), we assume for the time being that these newly permitted predications are always false, though we will revisit this issue later. We will focus on a liberal system that is in all other respects just like the cumulative one.

Why take this framework seriously? Firstly, there is no obvious reason why the type theory has to be just strict or cumulative, especially given that one can make good mathematical sense of liberalism (Gödel 1933, 46). Secondly, an argument from cumulativity to liberalism has recently been offered by Linnebo and Rayo (2012, 281–283), and we have already seen good reasons to take cumulativity seriously. We examine their argument in §7.

Recall the simplest version of the intuitive problem from earlier: (a) a property is an unrestricted domain only if it is possessed by all properties of level 1; (b)

although u^1 is possessed by all objects, it is not possessed by any property of level 1; therefore u^1 is not an unrestricted domain. Both claims are expressible within the liberal system. Claim (a) follows from R=U together with our assumed alignment of syntactic restrictions on predication with meaningful predicability (§3). Claim (b) is a consequence of the assumption that predications of the form $s^i(t^i)$ are always false. It follows that u^1 is not an unrestricted domain for basic generalisations in our first-order object language. This does not yet show that there are no such domains. However, given any domain d^i, one can use the following argument to refute the hypothesis that d^i is unrestricted for $\forall v F(v)$.

Consider any candidate domain d^i. We now show that d^i is not unrestricted for $\forall v F(v)$. Note first that not every entity possesses d^i. For example, since predications of the form $s^i(t^i)$ are false, no property of level i possesses d^i. Let f^j be the property expressed by F. Within the liberal system, f^j is grammatically predicable of properties of level i, i.e. $f^j(x^i)$ is well-formed for any i and j. So by our assumed alignment of syntactic structure with meaningful predicability, F is meaningfully predicable of properties of level i. Since no property of level i possesses d^i, it follows that the range of significance of F is not included in d^i. By R=U, d^i is not unrestricted for $\forall v F(v)$.

Because the levels i and j were arbitrary, this argument can be reproduced for higher-order object languages. Given any domain and basic generalisation in a higher-order language, one can use the argument to show that the domain is not unrestricted for the generalisation. In short, the intuitive problem is both expressible and sound in the liberal system.

In fact, a more informative result is available too. Define a type i as *bounded* if and only if, for some level j, expressions of type i are not meaningfully predicable of entities of level j or above. Consider a basic generalisation such as $\forall v F^i(v)$. We can show that this generalisation admits an unrestricted domain just in case its predicate's type i is bounded.

First, suppose that type i is unbounded. Adapting the argument of the previous paragraph, it follows that no domain is unrestricted for $\forall v F^i(v)$. Now suppose instead that i is bounded. Then the range of significance of F^i is confined below some level j. Let C be the collection of levels within which F^i can be meaningful predicated. We can find a domain coinciding with F^i's range of significance by combining the universal properties u^{j+1} where $j \in C$. For example, if F^i can be meaningfully predicated only within levels 0 and 2, then its range of significance coincides with a property that combines u^1 with u^3. Putting these results together, a basic generalisation admits an unrestricted domain just in case its predicate is of a bounded type.

The preceding arguments assume that, when $j \geq i$, predications of the form $s^i(t^j)$ are always false. But that assumption is not forced on us solely by liberalism's maximally permissive approach to predication. This suggests the possibility of reinstating unrestricted quantification by rejecting the assumption.

We followed Gödel in making this assumption because liberalism opens up new questions about the behaviour of properties predicated within their own level or above. The assumption embodies a radical yet simple answer to those questions. Other answers may be available. The effect of any such answer will be to allow properties that are more inclusive in the sense of being possessed by entities from their own level or above. Some answers will even be consistent with a property possessed by every entity, regardless of level. That would require suitable restrictions on comprehension and separation, but such a property would be a Russellian domain for any generalisation. A detailed investigation of the options here would take us too far afield from our primary concerns. Instead, we simply note that the impossibility of unrestricted quantification requires more than liberalism's maximally permissive conception of predication; it also depends on how properties behave when predicated within their own level and above.

7 Does cumulativity entail liberalism?

We have seen that unrestricted quantification is available in the strict and cumulative systems, but not in the liberal one. However, Linnebo and Rayo (2012, 281–283) have recently argued from cumulativity to liberalism. If sound, their argument shows, in effect, that the cumulative system is unstable, and hence that unrestricted quantification requires a strict type theory.

Recall that relations of cross-type identity are available in the cumulative setting. Linnebo and Rayo claim that such relations suffice for liberalism's type-unrestricted approach to predication:

> Once this type-unrestricted version of identity is in place, defining a type-unrestricted notion of predication is straightforward. One simply uses type-unrestricted identity to raise the type of the predicate enough to ensure that the predication is legitimate. Linnebo and Rayo 2012, 282

We can spell out the challenge emerging from this passage thus:

(1) In the cumulative system, there is a cross-type identity relation.

(2) If there is a cross-type identity relation, then there is a liberal notion of predication.

(3) If there is a liberal notion of predication, then one should adopt a liberal type theory.

We have already seen that (1) is true (§5). Moreover, we may grant (3) on the basis of the following considerations. If a liberal notion of predication is available, then

our assumed alignment of syntactic structure with meaningful predicability militates in favour of a similarly liberal type theory. For adopting a cumulative system would involve adopting what one regards as an unduly restrictive approach to predication. Although (1) and (3) are true, we will argue that (2) is false.

To see why one might believe premise (2), consider a predication that is ill-formed within the cumulative system but well-formed within the liberal system, e.g. '$a^1(b^2)$'. According to Linnebo and Rayo, one can leverage cross-type identity to show that any such predication is meaningful. Here is one way of doing so. Let a^1 be the property expressed by 'a^1', and b^2 be the property expressed by 'b^2'. A predication like '$a^1(b^2)$' is meaningful just in case it makes sense to predicate a^1 of b^2. So suppose there is an entity a^3 such that $a^1 \equiv a^3$. It makes sense to predicate a^3 of b^2. Since $a^1 \equiv a^3$, predicating a^3 of b^2 is just the same as predicating a^1 of b^2. So it makes sense to predicate a^1 of b^2, and '$a^1(b^2)$' is meaningful after all, despite violating the cumulative system's type restrictions.

We contend that the precise notion of predication delivered by this argument depends on the extension of the cross-type identity relation. In particular, one obtains a liberal notion of predication only if every entity is cross-type identical with an entity of any higher level. In symbols,

$$\forall x^i \exists y^j (x^i \equiv y^j)$$

for all types i, j with $i < j$. But this *axiom scheme of type-raising* may be contested. In fact, it follows from the version of comprehension presented in §5 that cross-type identity never holds between entities from different levels. For example, the following consequence of comprehension says that there is a property which distinguishes between the entities a^1 and a^3 employed in the argument:

$$\exists x^4 \big(\forall y^3 (x^4(y^3) \leftrightarrow y^3 = a^3) \wedge \forall y^1 (x^4(y^1) \leftrightarrow y^1 \neq a^1)\big)$$

It follows from this together with the definition of cross-type identity that $a^1 \not\equiv a^3$. Similar principles are available to distinguish any other pair of entities drawn from different levels.

The situation points to this conclusion: cumulativity alone does not suffice for a liberal notion of predication. For there is a consistent development of the cumulative system in which cross-type identity is empty. It is only when the cumulative system is augmented with substantive metaphysical assumptions—namely, all instances of the axiom scheme of type-raising—that the preceding argument delivers a liberal form of predication.[18]

[18]In this connection, notice that Linnebo and Rayo (2012, 288) include the relevant assumptions as additional axioms of their system. The assumption is also embedded in Degen and Johannsen's flexible rules for the quantifiers, which allow universal instantiation with respect to variable of any lower type (Degen and Johannsen 2000, 149).

Linnebo and Rayo avoid this problem by offering an alternative argument from cross-type identity to a liberal notion of predication. As we will see, however, it is doubtful whether this alternative argument really delivers a form of *predication*.

They first define the following cross-type relation:

$$x^i \varepsilon y^j \quad =_{\text{def}} \quad \begin{cases} y^j(x^i) & \text{if } i < j \\ \exists y^{i+1}(y^j \equiv y^{i+1} \wedge y^{i+1}(x^i)) & \text{if } i \geq j \end{cases}$$

They then use this relation to interpret predication; that is, they take '$y^j(x^i)$' to mean the same as '$x^i \varepsilon y^j$'. This provides meaning for predications that are ill-formed in the cumulative system. When '$y^j(x^i)$' is well-formed within the cumulative system, its meaning is unaffected by this interpretation. But when '$y^j(x^i)$' is ill-formed within the cumulative system, i.e. when $i \geq j$, this interpretation supplies it with the following meaning: something from immediately above x^i's level is both identical to y^j and possessed by x^i.

From the cumulative system's perspective, there are reasons to deny that ε is a form of predication. It is really a combination of quantification, cross-type identity, predication, and logical operations. To see this, consider the variant relation ε^* defined thus:

$$x^i \varepsilon^* y^j \quad =_{\text{def}} \quad \begin{cases} y^j(x^i) & \text{if } i < j \\ \forall y^{i+1}(y^j \equiv y^{i+1} \rightarrow y^{i+1}(x^i)) & \text{if } i \geq j \end{cases}$$

Like ε, this relation extends the cumulative notion of predication. It differs from ε only when $i \geq j$. In that case, it requires that *everything* immediately above x^i's level and identical to y^j is also possessed by x^i (whereas ε requires only that *something* immediately above x^i's level is). Moreover, ε^* has as good a claim as ε to be a form of predication. So we have two options for interpreting '$y^j(x^i)$': as synonymous either with '$x^i \varepsilon^* y^j$' or with '$x^i \varepsilon y^j$'. However, these interpretations are not equivalent. For example, if $i \geq j$ and y^j is not identical to anything of level $i+1$ or higher, then '$x^i \varepsilon y^j$' is false while '$x^i \varepsilon^* y^j$' is true. The quantifiers and other logical operators in the definitions are therefore not semantically inert. This suggests that neither relation is a form of predication. Both are logically complex combinations of quantification, cross-type identity, and genuine predication.[19]

To summarize, we have seen two potential ways of using cross-type identity to transform a cumulative notion of predication into a liberal one. The first delivered a genuinely predicational notion but fell short of liberalism. The second did not

[19]One might observe that ε and ε^* become coextensive if we assume the axiom scheme of type raising mentioned above. Even so, this shows only that the second argument from cumulativity to liberalism is no better than the first: both depend on substantive metaphysical assumptions.

deliver a genuinely predicational notion. Either way, the cumulative system appears to provide a stable theoretical framework hospitable to unrestricted quantification.

8 Conclusion

One central difficulty for unrestricted quantification is that, if domains are objects, no domain contains all objects. The intuitive problem we have considered is a kind of revenge problem facing type-theoretic responses to this difficulty: if domains are type-theoretic entities, no domain contains all type-theoretic entities. In light of this problem, one wonders whether anything is gained by adopting a type-theoretic framework.

Different authors have offered different views of this issue. Their arguments tend to invoke features peculiar to a specific type-theoretic framework. Evaluation of these arguments has been hampered by the absence of a principled account of unrestrictedness, which can be applied consistently across different frameworks.

We have provided precisely such an account, by arguing for an explication of unrestricted domains as Russellian domains. This reduces the question of whether a domain is unrestricted to a question about meaningful predicability: does the domain coincide with the appropriate range of significance? Because different type-theoretic frameworks embody different conceptions of meaningful predicability, they can diverge over this question. On the one hand, the strict and cumulative systems constrain meaningful predicability so that there is always a domain coincident with each range of significance. This makes both systems hospitable to unrestricted quantification. On the other hand, the liberal system imposes no constraints on meaningful predicability. In this respect, liberalism is like more familiar semantic frameworks where one can only quantify over objects: every entity recognised by the framework lies within every predicate's range of significance. Suitable principles governing the existence of domains (e.g. forms of separation and comprehension) then entail that the range of significance is never confined within a domain, which precludes unrestricted quantification. This shows that the intuitive problem requires substantive assumptions about the structure of predication. We conclude that something is gained by going type-theoretic just in case type-theoretic structure constrains meaningful thought and talk in an appropriate way.

References

Armstrong, D. (1980). Against "ostrich nominalism": a reply to Michael Devitt. *Pacific Philosophical Quarterly*, 61:440–449. Reprinted as (Mellor and Oliver, 1997, ch9).

Beaney, M., editor (1997). *The Frege Reader*. Blackwell Publishers.

Benacerraf, P. and Putnam, H., editors (1984). *Philosophy of Mathematics: Selected Readings*. Cambridge University Press, 2 edition.

Boolos, G. (1975). On second-order logic. *Journal of Philosophy*, 72(16):509–527. Reprinted as (Boolos, 1998, ch3).

Boolos, G. (1985). Nominalist platonism. *The Philosophical Review*, 94(3):327–344. Reprinted as (Boolos, 1998, ch5).

Boolos, G. (1998). *Logic, Logic and Logic*. Cambridge, Mass.: Harvard University Press.

Degen, W. and Johannsen, J. (2000). Cumulative higher-order logic as a foundation for set theory. *Mathematical Logic Quarterly*, 46(2):147–170.

Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, 30.

Dunaway, B. (2013). Modal quantification without worlds. In Bennett, K. and Zimmerman, D. W., editors, *Oxford Studies in Metaphysics Volume 8*, volume 8, chapter 4. OUP.

Feferman, S., Jr. J. W. D., Goldfarb, W., Parsons, C., and Solovay, R. M., editors (1995). *Kurt Gödel: Collected Works Volume III*. OUP.

Florio, S. and Incurvati, L. (2018a). Metalogic and the overgeneration argument. Manuscript under review.

Florio, S. and Incurvati, L. (2018b). Overgeneration in the higher infinite. Manuscript in preparation.

Florio, S. and Linnebo, Ø. (2019?). *The Many and the One: A Philosophical Study*. OUP. Manuscript in progress.

Frege, G. (1892). On concept and object. *Vietreljahrsschrift fur wissenschaftliche Philosophie*, 16:192–205. Reprinted in (Beaney, 1997, pp181–193).

Glanzberg, M. (2004). Quantification and realism. *Philosophy and Phenomenological Research*, 69:541–571.

Gödel, K. (1933). The present situation in the foundations of mathematics. Lecture for the Mathematical Association of America, 29-30 Deecember 1933. Published as (Feferman et al., 1995, 45–53).

Gödel, K. (1944). Russell's mathematical logic. In Schlipp, P. A., editor, *The Philosophy of Bertrand Russell*, The Library of Living Philosophers Volume V, chapter 3. Evanston and Chicago, Northwestern University. Reprinted as (Benacerraf and Putnam, 1984, 447–470).

Goodman, J. (2017). Reality is not structured. *Analysis*, 77(1):43–53.

Hale, B. (2013). *Necessary Beings: An Essay on Ontology, Modality, and the Relations Between Them*. OUP.

Hale, B. and Wright, C. (2001). *The Reason's Proper Study*. OUP.

Hale, B. and Wright, C. (2012). Horse sense. *The Journal of Philosophy*, 109:85–131.

Jones, N. K. (2016). A higher-order solution to the problem of the concept *horse*. *Ergo: An Open Access Journal of Philosophy*, 3:1–35.

Jones, N. K. (2017). Nominalist realism. *Noûs*. DOI 10.1111/nous.12193.

Krämer, S. (2013). Semantic values in higher-order semantics. *Philosophical Studies*, 168(3):709–724.

Krämer, S. (2017). Everything, and then some. *Mind*, pages 187–30.

Lavine, S. (2006). Something about everything. In Rayo, A. and Uzquiano, G., editors, *Absolute Generality*, pages 99–148. Oxford University Press.

Lewis, D. (1983). New work for a theory of universals. *The Australasian Journal of Philosophy*, 61(4):343–77. Reprinted as Lewis (1999, ch1).

Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell Publishing.

Lewis, D. (1999). *Papers in Metaphysics and Epistemology*. Cambridge University Press.

Linnebo, Ø. (2006). Sets, properties, and unrestricted quantification. In Rayo, A. and Uzquiano, G., editors, *Absolute Generality*, chapter 6. OUP.

Linnebo, Ø. and Rayo, A. (2012). Hierarchies ontological and ideological. *Mind*, 121(482):269–308.

Linsky, L. (1992). The unity of the proposition. *Journal of the History of Philosophy*, 30(2):243–273.

Mellor, D. H. and Oliver, A., editors (1997). *Properties*. OUP.

Potter, M. (2002). *Reason's Nearest Kin*. OUP.

Prior, A. N. (1971). *Objects of Thought*. OUP. Edited by P. T. Geach and A. J. P. Kenny.

Prior, A. N. and Fine, K. (1977). *Worlds, Times and Selves*. London: Duckworth.

Quine, W. V. O. (1970). *Philosophy of Logic*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition.

Rayo, A. (2006). Beyond plurals. In Rayo, A. and Uzquiano, G., editors, *Absolute Generality*, pages 220–254. OUP.

Rayo, A. and Williamson, T. (2003). A completeness theorem for unrestricted first-order languages. In Beall, J. C., editor, *Liars and Heaps: New Essays on Paradox*, chapter 15. OUP.

Rayo, A. and Yablo, S. (2001). Nominalism through de-nominalization. *Noûs*, 35(1):74–92.

Russell, B. (1908). Mathematical logic as based on the theory of types. *American Journal of Mathematics*, 30(3):222–262.

Shapiro, S. (1991). *Foundations Without Foundationalism: a Case for Second-Order Logic*. Clarendon Press, Oxford: OUP.

Stalnaker, R. C. (1976). Possible worlds. *Noûs*, 10(1):65–75.

Stalnaker, R. C. (2012). *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton University Press.

van Cleve, J. (1994). Predication without universals? a fling with ostrich nominalism. *Philosophy and Phenomenological Research*, 54(3):577–590.

Williamson, T. (2003). Everything. *Philosophical Perspectives*, 17(1):415–65.

Williamson, T. (2013). *Modal Logic as Metaphysics*. OUP.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. Routledge & Kegan Paul. Translated by C. K. Ogden with assistance from G. E. Moore, F. P. Ramsey, and Wittgenstein.

Wright, C. (1998). Why Frege did not deserve his *granum salis*: a note on the paradox of 'the concept horse' and the ascription of *bedeutungen* to predicates. *Grazer Philosophische Studien*, 55(239–263). Reprinted in (Hale and Wright, 2001, ch3).