

Two conceptions of absolute generality

Florio, Salvatore; Jones, Nicholas K

DOI:

[10.1007/s11098-022-01908-0](https://doi.org/10.1007/s11098-022-01908-0)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Florio, S & Jones, NK 2023, 'Two conceptions of absolute generality', *Philosophical Studies*.
<https://doi.org/10.1007/s11098-022-01908-0>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.



Two conceptions of absolute generality

Salvatore Florio^{1,2} · Nicholas K. Jones^{3,4}

Accepted: 4 December 2022
© The Author(s) 2023

Abstract

What is absolutely unrestricted quantification? We distinguish two theoretical roles and identify two conceptions of absolute generality: maximally strong generality and maximally inclusive generality. We also distinguish two corresponding kinds of absolute domain. A maximally strong domain contains every potential counterexample to a generalisation. A maximally inclusive domain is such that no domain extends it. We argue that both conceptions of absolute generality are legitimate and investigate the relations between them. Although these conceptions coincide in standard settings, we show how they diverge under more complex assumptions about the structure of meaningful predication, such as cumulative type theory. We conclude by arguing that maximally strong generality is the more theoretically valuable conception.

Keywords Absolute generality · Unrestricted quantification · Type theory · Higher-order logic · Property theory · Range of significance

Quantification is usually restricted. When one looks in the empty fridge and says “there’s no milk”, one is not refuted by the milk in the shop down the road. Different modes of inquiry aspire to different levels of generality. Generalisations in physics, for example, may concern the whole of physical reality, whereas generalisations in the special sciences concern only restricted portions of physical reality such as biological or social systems. Some modes of inquiry even aspire to maximal generality; many generalisations in metaphysics, logic, and mathematics, for example, appear to concern absolutely everything whatsoever, without restriction, physical or otherwise.

✉ Salvatore Florio
s.florio@bham.ac.uk

Nicholas K. Jones
nicholas.jones@philosophy.ox.ac.uk

¹ University of Birmingham, Birmingham, UK

² University of Oslo, Oslo, Norway

³ University of Oxford, Oxford, UK

⁴ St. John’s College, Oxford, UK

What is it to generalise without restriction? We argue that there is no univocal answer to this question. There are different legitimate conceptions of this phenomenon of *absolute generality*. We will distinguish two such conceptions, investigate the relation between them, and argue that one is more theoretically valuable than the other.¹ This will yield a new perspective from which one might reevaluate some prominent arguments against absolute generality, and the possibility of absolute generality in richer and more complex settings than first-order logic, for example cumulative and other forms of type theory.

1 Absolute generality and absolute domains

We begin from the idea that absolute generality is generalisation over a certain kind of domain, which we will call an *absolute domain*.² Using this terminology, we can distinguish two important questions, highlighted also by Agustín Rayo and Gabriel Uzquiano (2006, 2):

Metaphysical Question Is there an absolute domain?

Availability Question Could an absolute domain be available for human inquiry?

However, to properly understand these questions, we require an answer to the following, more fundamental question:

Analytical Question What is it for a domain to be absolute?

Every account of absolute generality requires an answer to this question. Absent such an answer, any account of absolute generality remains incomplete. Yet this foundational question has received insufficient attention in the prior literature. It will be our focus below.

The Analytical Question might appear straightforward: if a domain contains absolutely everything whatsoever, then it's absolute; otherwise it's not absolute. This appearance is misleading because it is not always clear what counts as "absolutely everything whatsoever". We now provide an example.

One prominent strategy for accommodating absolute domains employs primitive higher-order quantification and identifies domains with certain higher-order

¹ We focus on these conceptions because they seem especially natural and plausible targets of earlier discussions. Other conceptions are also available; see Button & Trueman 2021 and Pickel forthcoming for examples. Limitations of space prevent a more systematic investigation of the options.

² Our talk of domains is intended to codify talk about the entities relevant to the truth of a generalisation. We do not assume that each domain is an object. Indeed, many of the views we discuss later reject that assumption. See Cartwright 1994 for more.

“entities”.³ More concretely, the values of second-order variables—call them *properties of objects*—can be used as domains for first-order quantifiers. A property possessed by all objects, such as the property of being self-identical, seems a good candidate for an absolute domain. There is a clear sense in which this domain contains every object: every object instantiates it. But in standard systems of higher-order logic, it does not contain every property of objects. More precisely, one cannot even express in those systems that it contains every property of objects. What should we make of this fact? Must an absolute domain contain not just every object, but every property of objects, or even every higher-order entity whatsoever?

Our way of making progress here is to focus on the theoretical role that absolute generality plays. In fact, we will identify two such roles, *maximal strength* and *maximal inclusivity*. These roles give rise to two conceptions of absolute generality—as maximally strong generality, and as maximally inclusive generality—and two corresponding answers to the Analytical Question. Our primary concern is to develop an answer appropriate to maximally strong generality (Sect. 3) and discuss two objections to it (Sect. 4). As our responses to those objections will reveal, the two conceptions coincide in some settings. However, the conceptions diverge in other settings, and are therefore distinct (Sects. 5 and 6). Although we regard both conceptions as legitimate, we will argue that maximal strength is more theoretically valuable than maximal inclusivity (Sect. 7). Our next task is to introduce these two conceptions more fully. For simplicity, we focus on universal generalisation throughout the rest of the paper. Our discussion can be adapted to existential and other forms of generalisation.

We began with the idea that different modes of inquiry involve different levels of generality, and that some generalisations express greater generality than others. This brings with it the idea of increasing levels of generality, where greater levels of generality expand the supply of potential counterexamples. Absolute generality is a limit level, a kind of maximal generality. We will discuss two ways of understanding this limit.

According to the first understanding, the limit is reached when no expansion would include more potential counterexamples. We call this *maximally strong generality*. Note that ‘potential counterexample’ here and throughout has the most permissive sense, i.e. anything that can be meaningfully said to be a counterexample.

According to the second understanding, the limit is reached when no expansion would include more entities, potential counterexamples or otherwise. We call this *maximally inclusive generality*.

Given our starting assumption that absolute generality is generalisation over a certain kind of domain, we want to know what kinds of domain give rise to these two forms of maximal generality. Maximally strong generality arises from interpreting a generalisation over a domain that contains every potential counterexample to the generalisation. We say that any such domain is *maximally strong* for the generalisation in question. Note that this is a relational property of domains because it makes essential reference to a generalisation: different generalisations may have different supplies of

³ For relevant discussion, see Prior 1971, Chapter 3; Boolos 1975; Rayo & Yablo 2001; Williamson 2003; Hale 2013, Chapter 8.

potential counterexamples. A generalisation expresses maximally strong generality when interpreted over a domain that is maximally strong for it.

Maximally inclusive generality arises from interpreting a generalisation over a domain d such that no other domain contains everything in d as well as something else not in d . We say that any such domain d is *maximally inclusive*. Note that this is a monadic property of domains: maximal inclusivity is not dependent on any generalisation. A generalisation expresses maximally inclusive generality when interpreted over a maximally inclusive domain.⁴

We will argue that these two conceptions diverge because a domain can be maximally strong without being maximally inclusive.⁵ It follows that a generalisation can express maximally strong generality without expressing maximally inclusive generality. How can this be? As we will see, in some contexts a domain can be made more inclusive by including new entities irrelevant to the truth-value of, for example, ‘everything is F ’. The new entities involved in these contexts cannot be meaningfully said to instantiate the property f expressed by F .

2 The analytical question

We have now outlined the two conceptions of absolute generality we want to explore. Our next goal is to clarify the Analytical Question and to differentiate two kinds of answers to it.

What is it for a generalisation to express absolute generality, when interpreted over a domain? There are three individually necessary and jointly sufficient conditions. First, the generalisation must be meaningfully interpretable over the domain; in this case, call the domain *suitable for the generalisation*. Second, semantic evaluation of the generalisation must take account of each element of the domain that can be meaningfully said to be a counterexample. In short, no potential counterexample in the domain is semantically idle. In this case, say that the generalisation *exhausts the domain*.⁶ Exhaustion will play little explicit role below until our discussion of cumulativity (Sect. 6). We typically assume that the generalisations we discuss exhaust their domains, as in usual treatments of quantification. Third, we need some

⁴ We have used ‘maximally strong’ and ‘maximally inclusive’ ambiguously for forms of generality and properties of domains. We will also use them below for properties of generalisations. A generalisation is maximally inclusive (strong) if and only if it is interpreted over a domain that is maximally inclusive (strong for it). We allow context to disambiguate.

⁵ Is the converse true? That is, can a domain can be maximally inclusive without being maximally strong? In certain natural settings (e.g. standard formulations of strict and cumulative type theory) this is not possible. However, other settings are available in which the possibility arises. We do not pursue this topic further. We will argue (Sect. 7) that maximal inclusivity brings no additional expressivity, irrespective of whether maximal inclusivity implies maximal strength.

⁶ To see why this second condition is required, consider the following case. The generalisation ‘everything is F ’ is evaluated in a non-standard way over a domain d : the evaluation takes no account of whether some entity a in d is F , even though a can be meaningfully said to be F . Thus interpreted, the truth of ‘everything is F ’ in d does not preclude a from being a counterexample. This should not count as absolute generality. In effect, the real domain on this interpretation is not d , but d minus a . Requiring that the generalisation exhaust the domain prevents this kind of mismatch, thereby ensuring that no potential counterexample in the domain is semantically idle. As we will see in Sect. 6, this kind of mismatch can arise in cumulative type theory.

further property of the domains that are suitable for the generalisation in question; this is our target notion of *absolute domain*. As we will understand it, the Analytical Question concerns this further property of domains.

For present purposes, it is useful to distinguish two kinds of answers to the Analytical Question. According to one kind of answer, what it is for a domain to be absolute is independent of the generalisation in question. Accordingly, a domain may be absolute *simpliciter*. Formally, we can regard these answers as specifying what it is for a domain d to be absolute using a monadic condition $\phi(d)$. We call these *monadic answers* to the Analytical Question.

A different kind of answer says that what it is for a domain to be absolute depends on the generalisation in question. Accordingly, a domain is never absolute *simpliciter* but only ever absolute relative to, or *for*, a given generalisation g . Formally, these answers specify what it is for a domain d to be absolute using a relational condition $\phi(d, g)$. We call these *relational answers* to the Analytical Question. The present investigation of relational answers was spurred by the critical discussion in Button & Trueman 2021, to which we return below (Sect. 4).

To recapitulate, the difference between monadic and relational answers is as follows. A generalisation g expresses absolute generality if and only if g is interpreted over a domain d that is suitable for g , exhausted by g , and also such that...

$$\begin{aligned} (\textit{monadic}) \dots \phi(d). \\ (\textit{relational}) \dots \phi(d, g). \end{aligned}$$

This distinction is useful because one of the two conceptions of absolute generality introduced in Sect. 1 requires a monadic answer to the Analytical Question, whereas the other requires a relational answer.

Consider first the conception of absolute generality as maximally inclusive generality. On this view, absolute domains are maximally inclusive domains. Recall that a domain is maximally inclusive when no other domain contains everything in it as well as something else not in it. This characterisation concerns only features of domains themselves—specifically, what they contain—making no essential reference to a generalisation. We therefore have a monadic property of domains, and so maximally inclusive generality requires a monadic answer to the Analytical Question.

Now consider the conception of absolute generality as maximally strong generality. On this view, absolute domains are maximally strong domains. Recall that a domain is maximally strong for a generalisation when it contains every potential counterexample to the generalisation, in the most permissive sense of ‘potential counterexample’, i.e. anything that can be meaningfully said to be a counterexample.

Reference to the generalisation is essential here because different generalisations may have different supplies of potential counterexamples. For example, if the property expressed by F can be meaningfully predicated of some entity but the property expressed by H cannot be, then this entity is a potential counterexample to ‘everything is F ’ but not to ‘everything is H ’. So a domain that is maximally strong for one of those generalisations need not be maximally strong for the other. Maximal strength is therefore a relational property of domains, and the corresponding conception of absolute generality requires a relational answer to the Analytical Question. We say more about this in the next section.

3 A relational answer

We have argued that our two conceptions of absolute generality require different kinds of answers to the Analytical Question. Maximally inclusive generality requires a monadic answer. Maximally strong generality requires a relational answer. We now focus on the conception of absolute generality as maximally strong generality, and develop a corresponding relational answer to the Analytical Question, i.e. a specification of $\phi(d, g)$ in the relational schema of the previous section. We already have one such answer available: for d to be absolute for g is for d to contain every potential counterexample to g . However, we can expand this into a more informative answer, one that explains in more detail which domains contain every potential counterexample to a given generalisation. We draw on our proposal in Florio & Jones 2021.

Without loss of generality, we may restrict attention to basic universal generalisations, i.e. universal generalisations of the form ‘everything is F ’.⁷ Our guiding idea, inspired by Russell (1908), is that maximally strong generality is expressed when the domain contains every potential counterexample to the generalisation. All and only potential counterexamples are in principle relevant to the truth-value of the generalisation. As a result, “restricting” ‘everything is F ’ by excluding from the domain things that cannot be meaningfully said not to be F , is really no restriction at all. Our answer to the Analytical Question must therefore specify what kind of domain includes every potential counterexample to a universal generalisation’s truth.

Say that an entity is in the *range of significance* of a property if the property can be meaningfully predicated of the entity, i.e. if it makes sense to say that the entity instantiates the property. Moreover, say that a domain is *Russellian* for a generalisation when it contains the whole range of significance of the property expressed by the predicate in the generalisation. A domain includes every potential counterexample to a universal generalisation just in case it is Russellian for the generalisation. We now explain why.

⁷ Our focus on basic generalisations does not narrow the scope of our investigation because complex generalisations can be expressed using basic ones. Consider a complex generalisation $\forall v\phi(v)$ and let f be the property determined by $\phi(v)$. We can express the complex generalisation as a basic generalisation ‘everything is F ’ by interpreting F as expressing f .

A potential counterexample to ‘everything is F ’ is an entity that can be meaningfully said not to instantiate f , the property expressed by F . (We generally use uppercase letters for predicates and corresponding lowercase letters for the properties they express.) An entity can be meaningfully said not to instantiate f just in case it can be meaningfully said to instantiate f . So the potential counterexamples to ‘everything is F ’ coincide with the range of significance of f .

It follows that a Russellian domain for ‘everything is F ’ contains every potential counterexample to it. This means that maximally strong domains just are Russellian domains. On the conception of absolute generality as maximal strength, therefore, the answer to the Analytical Question is:

R=U Russellian domains are all and only absolute domains.

Following our previous usage (Florio & Jones 2021), we call this thesis R=U since it identifies *Russellian* and *unrestricted*, i.e. absolute, domains.

R=U is a relational answer to the Analytical Question because what counts as an absolute domain depends on the property expressed by the generalisation’s predicate. If different properties have different ranges of significance, then different domains will be absolute for basic generalisations involving them. The result is a close connection between absolute generality and the structure of meaningful predication. Note, however, that R=U itself is neutral about the precise details of this structure.

In Florio and Jones 2021, we considered three hypotheses about the structure of meaningful predication, corresponding to three different forms of type theory: strict, cumulative, and liberal.

Each hypothesis partitions reality into a well-ordered series of levels, or *types*. At the bottom (type 0) are the *objects*. Higher types comprise various kinds of properties: type 1 comprises properties of objects, type 2 comprises properties of properties of objects, and so on. The hypotheses differ over how types constrain meaningful predication, i.e. what can be meaningfully said to instantiate what. Objects cannot be meaningfully predicated of anything, under any hypothesis. So to see the differences between the hypotheses, consider a property f of some type $i > 0$.

In strict type theory, f can be meaningfully predicated of all and only entities from the immediately preceding type $i - 1$. In cumulative type theory, f can be meaningfully predicated of all and only entities from any preceding type $j < i$. In liberal type theory, f can be meaningfully predicated of any entity, irrespective of type. Note that these are metaphysical hypotheses about properties and their ranges of significance, not about the syntax of the language used to express them. From this metaphysical perspective, talk about predication is talk about instantiation. And the claim that an entity is meaningfully predicable of another may be understood as the claim that there is a proposition to the effect that the one instantiates the other.

Under which, if any, of these hypotheses is maximally strong generality possible given R=U? Whereas both strict and cumulative type theory permit absolute generality, liberal type theory does not. To see why, consider a universal generalisation ‘everything is F ’. This generalisation expresses maximally strong generality just in

case it's interpreted over a domain that contains the range of significance of f . So the key question is whether such a domain exists.

In strict type theory, a domain contains the range of significance of f if it contains everything whose type immediately precedes f 's type. It follows from a standard comprehension, or existence, principle that there is such a domain, and hence, by $R=U$, that absolute generality is possible.

In cumulative type theory, a domain contains the range of significance of f if it contains everything whose type precedes f 's type. It again follows from a natural comprehension principle that there is such a domain, and hence by $R=U$, that absolute generality is possible.

In liberal type theory, a domain contains the range of significance of f if it contains everything of every type. The existence of such a domain is incompatible with a standard comprehension principle, by a version of Russell's paradox. So by $R=U$, absolute generality is not possible.

The preceding arguments employ comprehension principles and presuppose classical logic. So one might be able to avoid these conclusions by restricting comprehension or changing the logic. We won't discuss these issues here. We are happy to regard these conclusions as conditional on these features of the type theory in question. The main point is that $R=U$ is neutral about the precise structure of meaningful predication and is therefore applicable under a range of different logico-metaphysical assumptions about that structure.

In the three settings above, the structure of meaningful predication is relatively simple. More complex structures are also possible. For an illustrative example, imagine a view on which there are two basic types, abstracta and concreta; moreover, some properties can be meaningfully predicated only of abstracta, others only of concreta, and yet others of both abstracta and concreta. We can apply $R=U$ here to determine which generalisations express absolute generality over which domains. Suppose *being located* can be meaningfully predicated of all and only concreta. Then 'everything is located' expresses absolute generality just in case the domain contains all concreta. By contrast, suppose *being self-identical* can be meaningfully predicated of all abstracta and all concreta, and nothing else. Then 'everything is self-identical' expresses absolute generality just in case the domain contains all abstracta and all concreta.

One can even apply $R=U$ to more complex views on which meaningful predication marks interesting distinctions within the types. For example, Gödel suggested a view of this kind as a solution to the paradoxes. He explains it as follows:

It should be noted that the theory of types brings in a new idea for the solution of the paradoxes, especially suited to their intensional form. It consists in blaming the paradoxes [...] on the assumption that every concept gives a meaningful proposition, if asserted for any arbitrary object or objects as arguments. [...]

The theory of simple types (in its realistic interpretation) can be considered as a carrying through of this scheme, based, however, on the following additional assumption concerning meaningfulness: "Whenever an object x can replace another object y in one meaningful proposition, it can do so in every mean-

ingful proposition.” This of course has the consequence that the objects are divided into mutually exclusive ranges of significance, each range consisting of those objects which can replace each other [...].

It is not impossible that the idea of limited ranges of significance could be carried out without the above restrictive principle. It might even turn out that it is possible to assume every concept to be significant everywhere except for certain “singular points” or “limiting points”, so that the paradoxes appear as something analogous to dividing by zero. (Gödel 1944, 149–150)

This proposal has recently been developed in Schindler 2019 and in Picenni & Schindler forthcoming.

To see how $R=U$ applies to this kind of view, suppose a property f can be meaningfully predicated only throughout some type i , with just one exception a . Then ‘everything is F ’ expresses absolute generality just in case the domain contains everything of type i with the possible exception of a . Thus $R=U$ again determines what kind of domain is implicated in absolute generality.

4 Two objections to $R = U$

The conception of absolute generality as maximally strong generality requires a relational answer to the Analytical Question. In this relational sense of absoluteness, no domain is absolute *simpliciter*. Rather, a domain is only ever absolute for a generalisation. In the previous section we developed a relational answer appropriate to the conception of absolute generality as maximally strong generality, namely $R=U$. This relational answer also gives rise to a non-trivial form of relativity. Specifically, a domain can be absolute for one generalisation yet not for another, even though the domain is suitable for both generalisations.

To illustrate this relativity, recall our earlier example, on which *being located* can be meaningfully predicated of concreta, whereas *being self-identical* can be meaningfully predicated of both abstracta and concreta, and nothing else. Then a domain is absolute for ‘everything is located’ just in case it contains all concreta. By contrast, a domain is absolute for ‘everything is self-identical’ just in case it contains all abstracta and all concreta. The domain of all and only concreta will therefore be absolute for the former generalisation but not the latter, even though it is suitable for both. The rest of this section considers two arguments against $R=U$ and the associated form of relativity.

Firstly, *the false advertising objection*, as raised by Tim Button and Rob Trueman (2021). According to this objection, relativity is incompatible with absoluteness. For example, absolute location would be location that is not relative to anything. And absolute truth would be truth without relativity to context, time, world, or any other parameter. Absolute generality should likewise be a non-relative form of generality. It is therefore false advertising to present $R=U$ as an account of absolute generality. Button and Trueman put it thus:

[T]he debate here is about absolute generality. It would be false advertising to enter that debate, promising to vindicate unrestricted quantification, and then only deliver relatively unrestricted quantification. (933)

Secondly, *the wrong subject matter objection*. According to this objection, the notion of absolute generality operative in the prior literature is not a relative or relational notion. That literature seems to treat absoluteness as a monadic notion, without explicit relativisation to a generalisation. One might therefore suspect that we have changed the subject by focussing on a relational notion of absoluteness.

To see where these objections go wrong, we should distinguish the notion under analysis from the theoretical tools used to analyse it. The notion presently under analysis is absolute generality. The theoretical tool used in the analysis is that of an absolute domain. This distinction gives rise to two versions of each objection. On the first version, they target absolute generality itself and assume that $R=U$ treats this as a relational or relative notion. On the second version, they target the notion of an absolute domain and assume that $R=U$ treats this as a relational or relative notion. Neither version of the objections is compelling. The first version's assumption is false. The second version's assumption is true but unproblematic.

Let us begin with the first version, which assumes that $R=U$ treats absolute generality as relational or relative. This assumption is false. According to relational answers to the Analytical Question, including $R=U$, a generalisation is absolute just in case it is interpreted over a domain that is absolute for the generalisation. Given a choice of interpretation, including a domain, these views classify some generalisations as absolute and others as not absolute. They do not classify any generalisation as absolute relative to anything other than the chosen interpretation, which should be uncontroversial. Rather, an underlying relation on domains is used to analyse a monadic property of generalisations. So, contrary to the false advertising objection, relational answers to the Analytical Question in general, and $R=U$ in particular, deliver a monadic notion of absolute generality rather than a merely relative one. And because they deliver a monadic notion of absolute generality, relational answers do not represent the notion of absolute generality operative in the prior literature as relational rather than monadic, contrary to the wrong subject matter objection.

Now for the second version, which assumes that $R=U$ treats the notion of absolute domain as a relational or relative notion. This assumption is true but unproblematic. Contrary to the false advertising objection, there is no good pre-theoretical reason to prefer a monadic answer to the Analytical Question over a relational one. Both kinds of answer provide legitimate theoretical options. They should be evaluated on the merits of their resulting accounts of absolute generality. This second version of the wrong subject matter objection relies on two claims:

- (i) the prior literature presupposes a monadic notion of absolute domain;
- (ii) one should treat the notion of absolute domain as monadic, if the prior literature does.

It's simplest to consider these claims in reverse order.

When evaluating competing theoretical analyses, there is no general presumption in favour of one form of analysis rather than another. This holds regardless of whether previous writers have presupposed some one particular form. One's theoretical apparatus should be evaluated on the merits of the account it delivers, not on the basis of accord with prior theories. Claim (ii) is therefore false.

Turning to claim (i), we argue that it is unclear whether the monadic rather than the relational notion of absolute domain is presupposed by the prior literature. The distinction between monadic and relational notions is not discussed explicitly. So the intended notion must be inferred from what was explicitly said. But this is difficult to do because most prior literature primarily concerns a setting in which a relational notion governed by $R=U$ is equivalent in the following sense to any plausible monadic notion:

For any domain d suitable for a generalisation g , d is absolute in the monadic sense just in case d is absolute for g .

We now elaborate on this equivalence.

Most prior literature operates within strict type theory, focusing primarily on interpretations of first-order languages. Recall that a domain is suitable for a generalisation just in case the generalisation can be meaningfully interpreted over the domain. So, when is a domain suitable for a generalisation within strict type theory?

Within strict type theory, properties of type n can be meaningfully predicated only of entities from the immediately preceding type $n - 1$. It is thus natural to assume that a domain d is suitable for a generalisation 'everything is F ' if and only if d has the same type as f . No other domains contain entities that can be meaningfully said to instantiate f .⁸

Given these assumptions about suitability, we can extract an account of which domains are absolute in the monadic sense. To begin, consider a predicate F that expresses a property of type 1. A suitable domain for 'everything is F ' is also a property of type 1, i.e. a property that can be meaningfully predicated of all and only objects (type 0). Now, suppose that 'everything is F ' expresses absolute generality over some domain or other. Then some suitable domain is absolute in the monadic sense. Since a suitable domain must be of type 1, the only plausible candidate is a universal property of type 1, i.e. a property instantiated by every object. Every other domain of type 1 is not absolute because it leaves out some objects and is expanded by such a universal property. Generalising, a domain is absolute in the monadic sense just in case it is a universal property of some type n , i.e. a property instantiated by every entity of type $n - 1$.⁹

⁸ Unless otherwise noted, we assume an *impermissive* interpretation of quantification: a domain d is suitable for a generalisation $\forall v\phi(v)$ only if each element of d can be meaningfully said to satisfy $\phi(v)$. Here we follow Florio & Jones 2021, 50–51.

⁹ This view is in keeping with Button & Trueman 2021, Williamson 2003, and Williamson 2013, 236–240.

We can now prove the desired equivalence, within strict type theory, between monadic accounts of absolute domain and a relational account governed by $R=U$. Let d be a domain suitable for a generalisation g whose predicate F expresses a property f of type n . Since d is suitable for g , then d has the same type n as f . As argued above, d is absolute in the monadic sense just in case it is a universal property of type n . All and only such properties contain the range of significance of f and are therefore Russellian for g . By $R=U$, the domains Russellian for g are all and only the domains absolute for g . So, putting these pieces together, d is absolute in the monadic sense just in case d is absolute for g , as required.

We have just seen that the relational account of absolute domain arising from $R=U$ is extensionally equivalent to any plausible monadic account of absolute domain, if the background type theory is strict. Insofar as the prior literature presupposes a strict type theory, there is no significant difference between these monadic and relational accounts of absolute domain. Contrary to claim (i) of the wrong subject matter objection, it is therefore not obvious that the monadic rather than the relational notion of absolute domain provides the best interpretation of the previous literature. The distinction between those interpretations makes no significant difference within strict type theory.

5 A distinction without a difference?

We have identified two conceptions of absolute generality: maximal inclusivity and maximal strength. We argued that maximal inclusivity requires a monadic answer to the Analytical Question, whereas maximal strength requires a relational answer, specifically $R=U$. We closed the previous section by showing that these answers are equivalent within strict type theory. It follows that the two conceptions of absolute generality are also equivalent within strict type theory. One might therefore wonder whether the distinction between the two conceptions is a distinction without a difference. We now argue that this is not so.

To see why the conceptions differ, consider:

Invariance If d is suitable for generalisations g and g^* each of which exhausts d , then either both or neither of g and g^* express absolute generality over d .

Monadic answers to the Analytical Question entail Invariance. By contrast, relational answers do not entail Invariance. In fact, $R=U$ can provide counterexamples to Invariance if the background type theory is not strict. The equivalence argued for in the previous section holds only because both views entail Invariance within strict type theory. The two conceptions diverge in more permissive settings. The next section explores one such possibility in more detail.

Here's why Invariance follows from monadic accounts of absolute domain. Recall from Sect. 2 that a generalisation g expresses absolute generality over a domain d if and only if (i) d is suitable for g , (ii) g exhausts d , and (iii) d is absolute. According

to monadic but not relational accounts, the notion of absoluteness in (iii) is monadic: d is absolute *simpliciter*.

Suppose domain d is suitable for two generalisations g and g^* each of which exhausts d . Assume without loss of generality that g but not g^* expresses absolute generality over d . Since g expresses absolute generality over d , it follows from monadic accounts that d is absolute *simpliciter*. This contradicts the assumption that g^* does not express absolute generality over d , which requires that d not be absolute *simpliciter*. So if d is suitable for both g and g^* each of which exhausts d , then g expresses absolute generality over d just in case g^* does too. Invariance follows.

This argument breaks down on relational accounts, such as $R=U$. The underlying issue is that d may be absolute for g without being absolute for g^* , even if it's suitable for and exhausted by both. In that case, relational accounts say that g but not g^* expresses absolute generality over d .

To illustrate how Invariance can fail, recall again our earlier example of abstracta and concreta: *being located* can be meaningfully predicated only of concreta, whereas *being self-identical* can be meaningfully predicated of both concreta and abstracta. Let c be a property instantiated by all and only concreta. According to $R=U$, c is absolute for 'everything is located' because only concreta can be meaningfully said not to be located, hence can be meaningfully said to be counterexamples. By contrast, c is not absolute for 'everything is self-identical': c does not include everything that can be meaningfully said to be a counterexample because abstracta can be meaningfully said not to be self-identical. Yet c is suitable for and exhausted by both generalisations. This contradicts Invariance.

Assuming $R=U$, counterexamples to Invariance are impossible in strict type theory. Counterexamples to Invariance essentially involve a domain d that is suitable for two generalisations, yet not absolute for both. We argued in the previous section that, within strict type theory, a domain d is suitable for a generalisation 'everything is F ' just in case d has the same type as the property f . So if d is suitable for two generalisations 'everything is F_1 ' and 'everything is F_2 ', then the properties f_1 and f_2 must have the same type. Therefore f_1 and f_2 have the same range of significance, and the same domains are Russellian for the two generalisations. It follows from $R=U$ that exactly the same domains are absolute for the two generalisations. Counterexamples to Invariance are therefore impossible if the background type theory is strict. Divergence between our two conceptions of absolute generality can arise only outside of this theory, where relational but not monadic accounts of absolute domain permit failures of Invariance.

6 Absolute generality in cumulative type theory

We have just seen that although monadic but not relational answers to the Analytical Question entail Invariance, extensional divergence is impossible within strict type theory. We will now see how they diverge within cumulative type theory. Specifically, maximal strength but not maximal inclusivity is possible within standard forms of cumulative type theory.

Maximally strong domains are possible in cumulative type theory (Florio & Jones 2021, Sect. 5). Consider a property of any type i . This property can be meaningfully predicated only of entities from types below i . So if F expresses a property of type i , any domain combining all types below i will be Russellian and hence also maximally strong for ‘everything is F ’. In standard systems of cumulative type theory, such domains exist for each type i .¹⁰

Matters are more delicate for maximal inclusivity. There is ultimately a strong case for the impossibility of maximally inclusive domains in cumulative type theory. However, it is instructive to begin with a more direct but ultimately less compelling argument against maximally inclusive domains in cumulative type theory.

Recall that a domain is maximally inclusive just in case no domain is more inclusive than it. And a domain d is more inclusive than a domain d^- just in case the following both hold:

- (i) everything in d^- is also in d ;
- (ii) something in d is not in d^- .

It makes sense to compare domains d and d^- for inclusivity only if (i) and (ii) make sense. Different systems of type theory therefore have different consequences about which domains can be meaningfully compared for inclusivity. So a maximally inclusive domain is best understood as a domain that is more inclusive than any other domain *with which it can be meaningfully compared for inclusivity*.

In strict type theory, domains can be meaningfully compared for inclusivity only when they have the same type. A maximally inclusive domain therefore needn’t contain every entity of every type, just every entity of the immediately preceding type. As a result, there is a maximally inclusive domain for each type, for example, the property of being a self-identical entity of that type.

In standard systems of cumulative type theory, one natural thought is that domains of any types can be compared for inclusivity (Krämer 2017, Button & Trueman 2021). These comparisons are enabled by the cross-type identity relations (Degen & Johannsen 2000, Linnebo & Rayo 2012), which cumulative type theory arguably provides:

¹⁰ This argument for the possibility of absolute generality in the sense of maximally strong generality in cumulative type theory presupposes $R=U$. Button and Trueman (2021, Sect. 7.6) contest this presupposition by arguing that cumulative type theory generates counterexamples to $R=U$. We disagree. Let us explain.

Florio & Jones 2021, Sect. 3, contains a proof of $R=U$. The proof assumes an *impermissive* interpretation of quantification: meaningful quantification never goes beyond the range of significance. Correspondingly, we defined a Russellian domain as one that *coincides* with the range of significance. By contrast, Button and Trueman’s counterexamples to $R=U$ require a *permissive* interpretation of quantification: meaningful quantification may go beyond the range of significance. But the definition of Russellian domain is not adjusted correspondingly with this shift in the interpretation of quantification. On a permissive interpretation of quantification, a Russellian domain should *contain* the range of significance, rather than *coincide* with it. This reinstates $R=U$. We return to permissive and impermissive interpretations of quantification in Sect. 7.

- (i*) everything in d^- is cross-type identical to something in d ;
(ii*) something in d is not cross-type identical to anything in d^- .

If cross-type identity makes sense across any pair of types, then (i*) and (ii*) make sense for any domains, irrespective of type. Any domains can then be meaningfully compared for inclusivity. Consequently, a domain is maximally inclusive only if it contains every entity of every type.¹¹ But, as emphasised by Stephan Krämer (2017), this is not possible. It follows that maximally inclusive domains are impossible in cumulative type theory.

This argument is persuasive only if cross-type identity really is a form of identity. This may be disputed (Florio & Jones 2021, 56). The problem is that differences in what can be meaningfully predicated of a and b should entail that a is not identical to b . Yet such differences are compatible with cross-type identity as standardly defined.¹² For example, if a is of type 2 and b is of type 4, then properties of type 3 can be meaningfully predicated of a and not b , even if a and b are cross-type “identical”.

The preceding arguments depend on substantive and controversial assumptions about the particular cumulative system one employs, and its metaphysical interpretation. Rather than examining these assumptions here, we now consider an argument that avoids them. This new argument relies only on structural features of any monadic answer to the Analytical Question alongside a natural assumption about what forms of quantification make sense. We call it *the argument from incorrect predictions*, for reasons that will become clear.

This new argument employs two principles. Firstly, we have already seen that every monadic account verifies:

Invariance If d is suitable for generalisations g and g^* each of which exhausts d , then either both or neither of g and g^* express absolute generality over d .

Secondly, cumulative type theories verify:

Monotonicity Let predicates F and H express properties f and h , where h has higher type than f . If d is suitable for ‘everything is F ’, then d is also suitable for ‘everything is H ’.

¹¹ Suppose that d does not contain some entity x . Standard systems of cumulative type theory include a form of comprehension entailing the existence of a property including everything in d as well as x . This property is a domain. And if all domains can be meaningfully compared for inclusivity, this domain is more inclusive than d . Since x was arbitrary, a maximally inclusive domain must contain every entity, irrespective of type.

¹² Cross-type identity is standardly defined as cross-type indistinguishability, where a is cross-type indistinguishable from b just in case, for every property that can be meaningfully predicated of both a and b , a instantiates it if and only if b does too. Formally, $\forall x^i(x(a) \leftrightarrow x(b))$ where i is the first type above the types of both a and b . These definitions are employed in Degen & Johannsen 2000, Linnebo & Rayo 2012, Krämer 2017, and Button & Trueman 2021.

Intuitively, Monotonicity says that increasing the type of property in a basic generalisation does not undermine suitability. This holds because, in cumulative type theory, increases in type correspond to expansions in range of significance. So if f is meaningfully predicable of everything in d , then so is h , whenever h has higher type than f . Increases in the type of property expressed by a predicate therefore introduce no new obstacles to a basic generalisation's meaningfulness over d .

The argument from incorrect predictions now proceeds thus. The following jointly entail incorrect predictions about which generalisations express absolute generality over which domains: Invariance, Monotonicity, and the assumption that absolute generality is possible in cumulative type theory. We explain this major assumption shortly. As the previous paragraph noted, cumulative type theories verify Monotonicity, and every monadic account of absolute domain verifies Invariance. It follows on every monadic account that absolute generality is not possible in cumulative type theory.

We now explain why our argument's major premise holds, assuming temporarily a monadic account of absolute domain. Suppose that absolute generality is possible in cumulative type theory. Let 'everything is F ' be any generalisation that expresses absolute generality over some domain d . By our earlier (Sect. 2) account of what it is for a generalisation to express absolute generality, it follows that d is suitable for 'everything is F ' and d is absolute *simpliciter*. In standard formulations of cumulative type theory, there is no highest type. So consider any other generalisation 'everything is H ' where H expresses property h of some type higher than the types of both f and d .

Because d is suitable for 'everything is F ' and h has higher type than f , Monotonicity implies that d is also suitable for 'everything is H '. And because d is absolute *simpliciter*, it follows from Invariance that if 'everything is H ' exhausts d , then 'everything is H ' expresses absolute generality over d . But this should not be. Since h has higher type than d , d excludes many potential counterexamples to 'everything is H '. For example, entities of the same type as d are not in d and yet may refute 'everything is H '.

So if absolute generality is possible in cumulative type theory, then an absolute domain suitable for a universal generalisation needn't contain all potential counterexamples to it. A universal generalisation can therefore express absolute generality, even though its truth does not preclude the existence of counterexamples. This is hard to accept. Even one who denies that preclusion of all counterexamples is sufficient for absoluteness should still regard it as necessary. Absolute generality is therefore impossible in cumulative type theory, assuming a monadic account of absolute domain.

Our reasoning two paragraphs ago establishes the following conditional: if 'everything is F ' expresses absolute generality over d , then 'everything is H ' fails to exhaust d whenever h has higher type than d and f . So the possibility of absolute generality depends on whether certain generalisations exhaust d , i.e. whether their semantic evaluation takes account of each potential counterexample in the domain. This depends in turn on what forms of quantification make sense. For example, some cumulative systems include quantifiers that, to put it intuitively, check all entities in the domain that are below a certain type (Degen & Johannsen 2000, Linnebo

& Rayo 2012, Button & Trueman 2021). By appropriate choice of type, these quantifiers will allow exhaustion of any domain and absolute generality will therefore be possible.¹³ By contrast, other systems include only quantifiers restricted to specific types of entities (Florio & Jones 2021). Intuitively, these quantifiers check only those entities in the domain that are of the specific type. If only these quantifiers are available, ‘everything is F ’ will not exhaust d whenever the range of significance of f includes entities in d that belong to different types. Despite the presence of absolute domains, absolute generality will not then be possible.

Now, maximally strong generality is possible in cumulative type theory (Florio & Jones 2021, Sect. 5). So the argument from incorrect predictions must break down on the conception of absolute generality as maximally strong generality. Where does it break down?

The argument breaks down because Invariance can fail on the conception of absolute generality as maximally strong generality. As we saw in Sect. 5, $R=U$ permits Invariance to fail. It fails whenever a domain can be suitable for two generalisations whilst excluding potential counterexamples to one and not the other. The domain is then Russellian for one but not the other, hence by $R=U$ also absolute for one but not the other. This violates Invariance. As we now argue, cumulative type theory generates cases of this kind.

The above argument began by supposing that ‘everything is F ’ expresses absolute generality over a domain d . According to $R=U$, it follows that d is Russellian for ‘everything is F ’, hence contains the range of significance of f . We then considered a generalisation ‘everything is H ’ that exhausts d and where h has higher type than both f and d . By Monotonicity, d is suitable for ‘everything is H ’. However, d is not Russellian for ‘everything is H ’. For the type of h is higher than that of d , and so h is meaningfully predicable of some entities outside d ; examples include any entity from the same type as d . It follows that d does not contain the range of significance of h , hence is not Russellian for ‘everything is H ’. So ‘everything is F ’ but not ‘everything is H ’ expresses maximally strong generality over d . Taking absolute generality to be maximally strong generality, those generalisations constitute a counterexample to Invariance.

7 Inclusivity or strength?

Under some substantive assumptions about the structure of meaningful predication (e.g. strict type theory), maximally strong generality and maximally inclusive generality are equivalent. Under other such assumptions (e.g. cumulative type theory), they are not equivalent. We therefore recognise two corresponding conceptions of absolute generality. Two questions now arise. Are both conceptions operative in the prior literature? Which conception is more theoretically valuable? We briefly discuss the first question, before turning to the second and, for

¹³ More precisely, suppose that f has type i . Then if the quantifier in ‘everything is F ’ checks at least all entities in d with type below i , ‘everything is F ’ will exhaust d .

our purposes, more important question. We will then argue that maximally strong generality is a more theoretically valuable conception, focussing on the theoretical value of expressivity.

The two conceptions of absolute generality open up a new interpretative perspective on the prior literature. On inspection, one can find passages suggestive of each conception. Let us illustrate with some examples. When introducing the concept of absolute generality in his influential paper “Everything”, Timothy Williamson writes:

Consider an example of a more specific metaphysical theory: out-and-out, no-holds-barred ontological naturalism, as in the slogan ‘Everything is part of the natural world’, in brief, ‘Everything is natural’. To interpret such naturalists as leaving it open that there are some contextually irrelevant non-natural things would be to miss their point, by failing to appreciate the radical extent of their claim (whether it is true or false). To understand them properly, one must interpret them as generalizing without any restriction whatsoever [...]. (Williamson 2003, 416)

This is strongly suggestive of maximal strength. Likewise for the opening paragraph of Rayo and Uzquiano’s Introduction to *Absolute Generality*:

When a philosopher asserts [“There are no abstract objects”], for example, we generally take the domain of her inquiry to comprise absolutely everything there is [...]. When presented with a purported counterexample, we do not regard it as open to the philosopher to reply that certain abstract objects are not relevant to her claim because, despite the fact they exist, they lie outside of her domain of inquiry. (Rayo & Uzquiano 2006, 1)

In a similar vein, here’s James Studd:

[I]nterpreting ‘everything’ to range over a less-than-absolutely-comprehensive domain appears to deprive the theorem of its intended generality. With the initial quantifier so restricted, an utterance of [“Everything is self-identical”] or [“Everything is the sole element of its singleton set”] fails to rule out the possibility of non-self-identical things or singletonless items outside the limited domain. To capture these theorems in their intended generality seems to call, on the contrary, for quantification over an absolutely comprehensive domain. (Studd 2019, 7)

In each case, note the connection between absolute generality and a universal generalisation whose truth precludes any counterexamples whatsoever, i.e. a maximally strong generalisation.

By way of contrast, the following passages are suggestive of maximal inclusivity. In each case, the mere existence of entities outside a domain, or the existence of a more inclusive domain, is taken to show that a domain is not absolute. The capacity of these additional entities to serve as potential counterexamples is not explicitly treated as relevant. For example, Vann McGee writes:

The thesis that genuinely universal quantification is not possible, so that, whenever we use quantifiers, even if it looks as if we are using them unrestrictedly, there will always be things that lie outside our universe of discourse, is not an easy doctrine to maintain. (McGee 2006, 185)

Whereas McGee focuses on entities outside the domain, Michael Glanzberg focuses on expansions of domains:

I have argued that all quantifiers must be construed as ranging over contextually provided domains, and that for any context, there is a distinct context which provides a wider domain of quantification. Hence, there is no absolutely unrestricted quantification. (Glanzberg 2006, 45)

In light of the above, it is plausible that both conceptions of absolute generality—as maximal strength and maximal inclusivity—are operative in the earlier literature. However, this is not decisive. Most prior discussion does not explicitly distinguish the two conceptions and presupposes a setting in which they are equivalent, as argued in Sect. 4. This makes it difficult to discern which notion was really intended. One may even question whether there is always a determinate fact of the matter. We won't try to settle these interpretative questions here. From our present perspective, the more important question is which conception is more theoretically valuable. To this we now turn.

We now show that maximal strength is more theoretically valuable than maximal inclusivity. We focus on the theoretical value of expressivity and argue for the following two claims. First, no expressivity is gained but some may be lost by switching from a maximally strong domain to any other domain, maximally inclusive or otherwise. Second, no expressivity is lost and some may be gained by switching from any domain, maximally inclusive or otherwise, to a maximally strong domain. We therefore regard maximal strength as the primary notion and more deserving of the label 'absolute generality'. We argue for the first claim. A parallel argument establishes the second claim.

Suppose we interpret a generalisation 'everything is F ' over a domain d_1 that is maximally strong for this generalisation. What expressivity would be gained or lost by switching to a different domain d_2 ?

Consider what this switch involves. There are two ways d_2 might differ from d_1 . Either something is in d_2 but not in d_1 or vice versa. We discuss each of these ways in turn.

We begin with the case in which d_2 and d_1 differ because something is in d_2 but not in d_1 . Since d_1 is maximally strong for 'everything is F ', d_2 contains something outside f 's range of significance. So by switching to d_2 , we interpret 'everything is F ' over a domain that goes beyond f 's range of significance. Is such an interpretation possible? This gives rise to a dilemma, according to whether d_2 is suitable for 'everything is F ' or not.

On one horn of the dilemma, d_2 is not suitable for 'everything is F '. This follows from what we call an *impermissive* interpretation of quantification: meaningful quantification never goes beyond the range of significance. That is, a domain is suitable for 'everything is F ' only if the domain is contained in f 's range of significance.

There is then no such thing as switching from d_1 to d_2 . For there is no such thing as meaningfully generalising over a domain that extends beyond the range of significance. So let us set this horn aside.

On the other horn of the dilemma, d_2 is suitable for ‘everything is F ’. This requires a *permissive* interpretation of quantification: meaningful quantification can go beyond the range of significance. That is, a domain may be suitable for ‘everything is F ’ even though it extends beyond f ’s range of significance. Switching from d_1 to d_2 is now a genuine possibility. However, the entities in d_2 but not in d_1 are outside f ’s range of significance and so play no role in semantic evaluation of ‘everything is F ’. On the one hand, they are not potential counterexamples to ‘everything is F ’ because they cannot be meaningfully said to instantiate f . (Recall from Sect. 3 that an entity can be meaningfully said not to instantiate a property just in case it can be meaningfully said to instantiate that property.) On the other hand, the truth of ‘everything is F ’ cannot require them to instantiate f because, again, they cannot be meaningfully said to instantiate f . So although d_2 contains entities not in d_1 , those entities are irrelevant to the truth or falsity of ‘everything is F ’ interpreted over d_2 . No expressivity is therefore gained by switching from maximally strong domain d_1 to d_2 .

It follows that no expressivity is gained by switching from a maximally strong domain to any other domain. On impermissive interpretations of quantification, there is no such thing as switching. On permissive interpretations of quantification, any differences between the two domains are semantically idle.

We now turn to the case in which d_2 and maximally strong d_1 differ because something is in d_1 but not in d_2 . Any such entity a is either in f ’s range of significance or not. Suppose that a is in f ’s range of significance. Then some expressivity is lost by switching to d_2 : d_2 excludes a potential counterexample to ‘everything is F ’. Suppose instead that a is not in f ’s range of significance. (This can only occur on permissive interpretations of quantification; otherwise d_1 would not be suitable for ‘everything is F ’.) Then a plays no role in semantic evaluation of ‘everything is F ’, as explained above. So, again, no expressivity is gained by switching from maximally strong domain d_1 to any other domain d_2 . In fact, some expressivity may be lost because potential counterexamples may be excluded from d_2 .

To summarise, we have argued that switching from a maximally strong domain to another, maximally inclusive or otherwise, gains no expressivity and may in fact lose some. What about the opposite switch, from any domain, maximally inclusive or otherwise, to another that is maximally strong? A parallel argument shows that no expressivity is lost and some may in fact be gained.¹⁴ The upshot is that maximal strength is in

¹⁴ A sketch of the argument follows. A domain d_2 can differ from a maximally strong domain d_1 in two ways. First, something is in d_2 but not in d_1 . Any such entity is not in f ’s range of significance and so plays no role in semantic evaluation of ‘everything is F ’. Therefore no expressivity is lost by switching from d_2 to d_1 . Second, something is in d_1 but not in d_2 . Any such entity is either in f ’s range of significance or not. If it is, then some expressivity is gained by switching from d_2 to d_1 : d_1 includes a potential counterexample to ‘everything is F ’ that d_2 excludes. If it is not, then no expressivity is lost by switching from d_2 to d_1 because the entity plays no role in semantic evaluation of ‘everything is F ’. In conclusion, switching from any domain to a maximally strong domain loses no expressivity and may in fact gain some by including further potential counterexamples.

this sense more theoretically valuable than maximal inclusivity. Although both kinds of maximal generality are legitimate, we therefore regard maximally strong generality as a better candidate for the label ‘absolute generality’.

Acknowledgements For helpful comments and discussion, we would like to thank Tim Button, Pierluigi Graziani, Bruno Jacinto, Alex Kaiserman, Øystein Linnebo, Ofra Magidor, Agustín Rayo, Alex Roberts, Chris Scambler, James Studd, Rob Trueman, and Gabriel Uzquiano as well as audiences in Jerusalem, Lisbon, Oxford, and Urbino. We are also grateful to two anonymous referees whose suggestions significantly improved the paper. Salvatore’s work was supported by the Arts and Humanities Research Council [grant number AH/V015516/1].

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Boolos, G. (1975). On second-order logic. *Journal of Philosophy*, 72(16), 509–527.
- Button, T., & Trueman, R. (2021). Against cumulative type theory. *The Review of Symbolic Logic*, 15(4), 907–949.
- Cartwright, R. L. (1994). Speaking of everything. *Noûs*, 28(1), 1–20.
- Degen, W., & Johannsen, J. (2000). Cumulative higher-order logic as a foundation for set theory. *Mathematical Logic Quarterly*, 46(2), 147–170.
- Florio, S., & Jones, N. K. (2021). Unrestricted quantification and the structure of type theory. *Philosophy and Phenomenological Research*, 102(1), 44–64.
- Glanzberg, M. (2006). Context and unrestricted quantification. In Rayo, A., & Uzquiano, G. (eds.), *Absolute Generality* (pp. 20–44). Oxford University Press.
- Gödel, K. (1944). Russell’s mathematical logic. In Schlipp, P. A. (ed.), *The Library of Living Philosophers: The Philosophy of Bertrand Russell* (pp. 123–154). Tudor Publishing Company.
- Hale, B. (2013). *Necessary Beings: An Essay on Ontology, Modality, and the Relations Between Them*. Oxford University Press.
- Krämer, S. (2017). Everything, and then some. *Mind*, 126(502), 499–528.
- Linnebo, Ø., & Rayo, A. (2012). Hierarchies ontological and ideological. *Mind*, 121(482), 269–308.
- McGee, V. (2006). There’s a rule for everything. In Rayo, A., & Uzquiano, G. (eds.), *Absolute Generality* (pp. 179–202). Oxford University Press.
- Picenni, S., & Schindler, T. (forthcoming). Transcending the theory of types. In Petrolo, M., & Venturi, G. (eds.), *Paradoxes Between Truth and Proof*. Springer.
- Pickel, B. (forthcoming). Against second-order primitivism. In Fritz, P., & Jones, N. K. (eds.), *Higher-Order Metaphysics*. Oxford University Press.
- Prior, A. N. (1971). *Objects of Thought*. Geach, P. T., & Kenny, A. J. P. (eds.). Oxford University Press.
- Rayo, A., & Uzquiano, G. (eds.) (2006). *Absolute Generality*. Oxford University Press.
- Rayo, A., & 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.
- Schindler, T. (2019). Classes, why and how. *Philosophical Study*, 176, 407–435.
- Studd, J. P. (2019). *Everything, More or Less*. Oxford University Press.
- Williamson, T. (2003). Everything. *Philosophical Perspectives*, 17, 415–465.
- Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford University Press.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.