

ON TYPE DISTINCTIONS AND EXPRESSIVITY

SALVATORE FLORIO

Quine maintained that philosophical and scientific theorising should be conducted in an untyped language, which has just one style of variables and quantifiers. By contrast, typed languages, such as those advocated by Frege and Russell, include multiple styles of variables and matching kinds of quantification. Which form should our theories take? In this article, I argue that expressivity does not favour typed languages over untyped ones.

I

Introduction. Formal languages have come to play a key role in the pursuit of philosophical and scientific rigour. In choosing a formal language, however, we face an embarrassment of riches, given the wide variety of available options. One of the fundamental differences concerns the basic structure of the language and, in particular, whether or not the language includes type distinctions.

Philosophical analysis often relies on the familiar language of first-order logic, which has a simple type-theoretic structure. There is just one style of variables and one matching kind of quantification—the language is *untyped*. The semantic picture associated with this language is correspondingly simple: there is one domain of entities, which can be characterised using the predicates available in the language.

Several alternative frameworks are based on *typed* languages, which exhibit richer structures of types. A prominent example is second-order logic, which adds to first-order logic a second style of variables and a matching kind of quantification. The semantic picture associated with this system is less straightforward but, on a popular view, the language describes two incommensurable domains: objects and properties of objects (or concepts). Each domain is then characterised using predicates appropriate to that domain. One can introduce more styles of variables and quantification, as is in full higher-order logic, which includes infinitely many beyond the second.

Does a single style of variables suffice for theoretical inquiry? According to the philosophical tradition famously represented by Quine, the answer is positive: an untyped system is enough for our philosophical and scientific needs. The tradition associated with Frege and Russell supports a negative answer: we require typed languages. These languages have received renewed philosophical attention in the past two decades, and recent philosophical work in metaphysics (see Skiba 2021 and Fritz and Jones forthcoming a) as well as philosophy of mathematics (e.g. Linnebo and Rayo 2012, Ladyman and Presnell 2018, and Corfield 2020) has explored new applications of typed systems, including sophisticated frameworks such as type theories with transfinite types and homotopy type theory.

As also suggested by longstanding philosophical disputes, the contrast between these two positions seems to be significant. Moreover, it is of foundational importance, since it concerns the basic form of our theories. What considerations support the adoption of a framework with one style of variables rather than many? It has been argued, for example, that the richer type-theoretic structure of second-order logic is superfluous, illegitimate, or even unintelligible (e.g. Quine 1956 and Quine 1986, pp. 66-8). But one can also find many

arguments defending the theoretical importance, or even the indispensability, of the richer structure (e.g. Williamson 2003, Williamson 2013, Dorr 2016, Jones 2018, and Trueman 2021). For instance, Russell held that avoiding paradoxes involving self-reference *necessitates* ‘the division of objects into types’ (1908, p. 237)

Our focus here is on the important but puzzling role that the notion of *expressivity* plays in this debate. On the one hand, any untyped language can be extended to a typed one and can therefore be seen as part of a broader, typed language. As a result, typed languages may seem to afford *greater expressive power* than untyped ones and, on this basis, be preferable. On the other hand, type-theoretic hierarchies have been thought to have expressive limitations that can be overcome in an untyped language. Wittgenstein and Gödel, among others, expressed views of this kind (Wittgenstein 1979 pp. 106ff.; Gödel 1944, p. 466). Moreover, there are arguments that untyped languages are *too expressive*, since they permit the formalisation of nonsense or force us to make unnecessary theoretical choices, which can be avoided in typed languages.

In this article, I will develop some of these argumentative strategies and show that expressivity does not favour typed languages over untyped ones. Overall, this will provide a partial vindication of untyped languages. I will defend two main claims. The first is that, in this context, appealing to greater expressive power in support of typed languages is not effective. The second is that untyped languages are not too expressive in the sense described above. So, if there are compelling reasons to adopt a typed language, they are not based on expressivity.

II

Many-sorted logic as a framework for the debate. To set the stage for the debate and compare alternative approaches, I believe it is illuminating to use the framework of many-sorted logic. I will start by giving a brief, informal presentation of the framework, before commenting on its benefits here.

The signature of a many-sorted logic includes a set of objects representing the *sorts* admitted by the logic. The set could be finite or infinite. Each sort has an associated style of variables and quantifiers. For example, one may countenance infinitely many sorts and use the natural numbers as superscripts to represent them. Then, for each natural number n , there are variables of sort n (x^n, y^n, z^n, \dots) and corresponding quantifiers ($\forall x^n, \forall y^n, \forall z^n, \dots$). The non-logical terminology comprises constants of each sort (e.g. a^0, a^1, a^2, \dots) as well predicates, and functions, with appropriate sortal restrictions. More specifically, each argument place is restricted to one sort. For example, one may have a binary relation R taking terms (variables or constants) of sort 1 in its first argument place, and terms of sort 0 in its second argument place. This means that $R(x^1, a^0)$ is well formed, but $R(x^0, a^1)$ is not. To mark sortal restrictions, predicates and functions are usually decorated with a sequence representing the sorts of their arguments. So, in our example, the relation R would be presented as $R^{<1,0>}$. (We omit this decoration whenever possible.)

An important choice point has to do with the predicates and functions that are admitted in the language. One could be maximally liberal and, for any finite combination of sorts, allow predicates whose arguments correspond to that combination. That is, for any sequences of sorts $\langle i_1, \dots, i_n \rangle$ and $\langle i_1, \dots, i_n, i_{n+1} \rangle$, one would admit predicates like $R^{<i_1, \dots, i_n>}$ and

functions like $f^{<i_1, \dots, i_n>}$ yielding a term of sort i_{n+1} . Alternatively, one could be more restrictive and ban predicates or functions whose arguments correspond to certain combinations of sorts.¹ For example, one may allow relations like $R^{<1,0>}$ but ban relations like $R^{<1,1>}$. In this case, one allows relations between sort 1 and sort 0, but no relations between terms of sort 1. As we will see later, some debates about the correct form of type theory can be interpreted here as debates about which predicates to allow.

There is a key difference between the framework of higher-order logic and that of many-sorted logic. In higher-order logic, predication is normally treated as a logical notion and is not represented by any symbol in the language. By contrast, predication is introduced as a non-logical symbol in many-sorted logic. So, in order to express in second-order logic that a property applies to an object, one would use a formula like $x^1(y^0)$, or $X(y)$. In many-sorted logic, one would use a formula like $P^{<1,0>}(x^1, y^0)$, relying on the relational symbol $P^{<1,0>}$ to express predication for terms for the relevant sorts.

It is straightforward to recapture the standard systems of first- and second-order logic within many-sorted logic. First-order logic is the limit case in which there is only one sort. Second-order logic corresponds to the case in which there are two sorts and predicates are constrained as follows. All predicates of first-order logic are allowed. In addition, there are relational symbols expressing predication relations holding between a term of the second sort (e.g. 1) and terms of the first sort (e.g. 0). The second-order axiom scheme of comprehension asserts that for any appropriate formula φ , there is a second-order entity that holds of exactly the first-order entities satisfying φ . In symbols:

$$\exists X \forall y_1 \dots \forall y_n (X(y_1, \dots, y_n) \leftrightarrow \varphi(y_1, \dots, y_n))$$

where X does not occur free in φ . In many-sorted logic, an analogous schema expresses the principle that for any appropriate formula φ , there is an entity of the second sort that stands in the predication relation exactly to the entities of the first sort that satisfy φ . That is,

$$\exists x^1 \forall y_1^0 \dots \forall y_n^0 (P^{<1, 0, \dots, 0>}(x^1, y_1^0, \dots, y_n^0) \leftrightarrow \varphi(y_1^0, \dots, y_n^0))$$

where x^1 does not occur free in φ . It is also straightforward to extend this translation procedure to standard logics of higher orders.

From the semantic point of view, it is common to interpret a many-sorted system by means of a set-theoretic model (or structure) consisting of many non-empty domains, one for each sort, and relations respecting the sortal restrictions of the predicates in the language. For instance, the relational symbol $R^{<1,0>}$ is interpreted as a set of pairs $\langle u, t \rangle$, where u is in the domain of sort 1 and t is in the domain of sort 0. The interpretation of the symbol for predication is not fixed but can vary across models. This marks an important difference with respect to the standard semantics for second-order logic, where predication has a fixed interpretation. Assuming a countable language, this semantic approach to many-sorted logic leads to a system that shares crucial metalogical properties with first-order logic, e.g. completeness, compactness, and the Löwenheim-Skolem property. In this sense, the system can be said to be *first-order* (see, e.g., Shapiro 1991, p. 14).

¹ For simplicity, we can leave out function symbols in the rest of the discussion. This will not affect my main points, though function symbols are relevant to the theme of section VI. See footnote 3.

Many-sorted logic can be a valuable tool when comparing different typed languages. It reduces some debates about the correct form of type theory to debates about which predicates to allow. This yields a simpler and arguably more tractable rendering of the original debates. For instance, the debate between critics and supporters of cumulative type theory (Degen and Johannsen 2000; Williamson 2013, 237-8; Krämer 2017; Linnebo and Rayo 2012; Florio and Jones 2021; Button and Trueman 2022) reduces to a debate about which predication relations may be admitted. Moreover, the treatment of predication in many-sorted logic seems better suited to represent some metaphysical ideas. For instance, it is more hospitable to views according to which predication is just one among many fundamental ontological relations, such as exemplification, participation, and realisation (e.g. Smith 2005). It also allows a clean separation between ontology (given by quantification) and ideology (given by the predicates admitted in the language). This distinction is less clear in the standard formulation of higher-order logic, which relies on quantification into predicate position, raising the question how to understand the commitments incurred by this form of quantification and their relation to those of quantifiers in nominal position. Additional considerations in favour of many-sorted logic can be found in Manzano 1996, which provides a systematic defense of this framework emphasising its unifying power.

Many-sorted logic can be a valuable tool also when comparing typed languages and untyped ones. Translating from the former to the latter is, as we will see, straightforward. Furthermore, a number of potentially useful results concerning theoretical equivalence become immediately available. In philosophy of science, there is a well-known problem of characterising various senses in which two theories can be said to be equivalent. The problem becomes more complex when we compare theories with very different type-theoretic structures. This is because the standard definitions of a number of relevant notions (e.g. mutual interpretability and definitional equivalence) apply directly only to theories formulated in languages with similar types. Many-sorted logic offers a fruitful setting to tackle this problem. For example, recent work by Thomas William Barrett and Hans Halvorson has clarified some of the ways in which a one-sorted system can be said to be definitionally equivalent to a system with many sorts (see Barrett and Halvorson 2017; Halvorson 2019, chs. 4-5).

So I propose to adopt many-sorted logic as a framework for our discussion. This means formulating typed theories as many-sorted theories. Our overarching question ('should we adopt a typed language?') becomes a question about sorts ('should we admit more than one sort?'). The contrast is thus between one-sorted systems and systems with more than one sort. It will be convenient henceforth to abuse terminology and call a system, language, or logic *many-sorted* if and only if it has more than one sort. (As we have seen, a one-sorted system is, strictly speaking, a limit case of a many-sorted system.)

I am now ready to develop and assess two arguments against untyped languages, which I set out to do in section I. Recall the gist of the arguments: typed languages have more expressive power than untyped languages; untyped language are too expressive in problematic ways.

Let us start with the first argument. Do many-sorted systems really have *more expressive power* than those with one sort only? There at least two ways of developing the argument, depending on how expressive power is understood. One understanding is semantic, the other is syntactic. On the semantic understanding, many-sorted logic has more expressive power than one-sorted logic because more can be represented in many-sorted logic than it can in

one-sorted logic. On the syntactic understanding, many-sorted logic has more expressive power because more can be said and proved in many-sorted logic than it can in one-sorted logic.

III

Can many-sorted logic represent more than one-sorted logic? The common model-theoretic semantics sketched in the previous section relies on an untyped metatheory, namely standard first-order set theory, to describe the range of models that can be represented by many-sorted languages. A model of the language consists of many non-empty sets representing domains of quantification, one for each sort, and set-theoretic relations denoted by predicates of the many-sorted object language. Denotations respect the sortal constraints of the object language. Let the domains be U_i with $i \in I$, the set of sorts. Then a constant a^i of sort i denotes an element of U_i , and a predicate $R^{<i_1, \dots, i_n>}$ denotes a subset of $U_{i_1} \times \dots \times U_{i_n}$. A many-sorted theory selects a class of models, those making all sentences of the theory true.

Models of many-sorted languages ('many-sorted models' for short) are seemingly more complex than models of one-sorted languages ('one-sorted models' for short). In effect, a many-sorted model combines multiple one-sorted models into one, with additional information about relations among entities of different sorts. This means that any one-sorted model is included in some many-sorted model. It may thus appear that the range of models that can be represented by many-sorted languages extends the range of models that can be represented by one-sorted languages.

However, this appearance is superficial. It is well known that every many-sorted model can be easily converted into an equivalent one-sorted model (see, e.g., Enderton 2002, pp. 296-9). The simple idea behind this *model conversion* is to capture sorts using special predicates. To this end, we add to the one-sorted language a new predicate S_i for each sort i . The rest of the one-sorted language is that of the many-sorted language, only treated as one-sorted. To obtain a one-sorted model from a given many-sorted model, one proceeds as follows. The domain is the union of domains U_i of the many-sorted model. Each predicate S_i denotes the set U_i . The rest of the non-logical vocabulary retains the denotation it has in the many-sorted model.

Let us explore how to capture formally the equivalence between a many-sorted model M and the one-sorted model M^* obtained from the model conversion just described (again, see Enderton 2002, pp. 297-8). Define a translation \bullet from a many-sorted language to a one-sorted language that maps quantification of sort i to quantification restricted to S_i , treating the rest of the language as if it were one-sorted. More precisely, the action of \bullet is given by the following recursive clauses, where $[\varphi]^\bullet$ is the result of applying \bullet to φ , t^i is a term (constant or variable) of sort i , and $R^{<i_1, \dots, i_n>}$ as well as t_{i_1}, \dots, t_{i_n} are one-sorted expressions in which the sortal decoration functions as a mere subscript.

$$\begin{aligned} [R^{<i_1, \dots, i_n>}(t^{i_1}, \dots, t^{i_n})]^\bullet &= R^{<i_1, \dots, i_n>}(t_{i_1}, \dots, t_{i_n}) \\ [t^{i_1} = t^{i_2}]^\bullet &= t_{i_1} = t_{i_2} \\ [\neg\varphi]^\bullet &= \neg [\varphi]^\bullet \end{aligned}$$

$$[\varphi \wedge \psi]^* = [\varphi]^* \wedge [\psi]^*$$

$$[\forall x^i \varphi]^* = \forall x_i (S_i(x_i) \rightarrow [\varphi]^*)$$

Terms and atomic formulas, including identities, are thus treated as if they were one-sorted. The translation commutes with negation and the logical connectives. Finally, quantification over entities of sort i is replaced by quantification restricted by the special predicate S_i .

Then we have that M satisfies a sentence σ if and only if M^* satisfies $[\sigma]^*$:

$$M \models \sigma \text{ if and only if } M^* \models [\sigma]^*$$

The differences between M and M^* seem unimportant to what we can express. The two models are about the same entities, and the denotations of the corresponding primitive expressions are the same. It can be shown that any sentence σ and its translation $[\sigma]^*$ have equivalent truth conditions when interpreted in M and M^* , respectively. This is obvious for atomic formulas: the denotations of terms and atomic predicates is the same. Now consider a quantified sentence, such as $\forall x^{i_1} \forall x^{i_2} R^{<i_1, i_2>}(x^{i_1}, x^{i_2})$, which can be read as ‘everything of sort i_1 is related by R to everything of sort i_2 ’. Both this sentence and its translation

$$\forall x_{i_1} (S_{i_1}(x_{i_1}) \rightarrow (\forall x_{i_2} (S_{i_2}(x_{i_2}) \rightarrow R_{<i_1, i_2>}(x_{i_1}, x_{i_2}))))$$

assert the same: every element of U_{i_1} stands in a certain relation to every element of U_{i_2} , the relation denoted by both $R^{<i_1, i_2>}$ and $R_{<i_1, i_2>}$. A simple induction establishes the equivalence in full generality.

Assuming that the extensions of the new predicates S_i are non-empty, the equivalence extends to logical consequence. Take any set Σ of sentences in the many-sorted language, and define Σ^* to be the set of their translations under \bullet . Let Δ be $\{\exists x_i S_i(x_i) : i \in I\}$, the set of sentences asserting that each S_i has a witness. Then $\Sigma \models \sigma$ in the many-sorted system if and only if $\Sigma^* \cup \Delta \models [\sigma]^*$ in the one-sorted system (see Theorem 43C of Enderton 2002, p. 298). This result will be used below (section V).

The preceding paragraphs establish that any model of a many-sorted language correspond to model of a one-sorted languages and any sentence in a many-sorted language has a one-sorted counterpart with equivalent truth conditions. In this sense, many-sorted languages are not more expressive than one-sorted languages. The discussion assumed an *untyped metatheory*: the semantics was formulated in standard first-order set theory, an untyped system. However, one may reject this assumption and insist on a *typed metatheory*. On this approach, a typed language is assumed as primitive in the metatheory, and the semantics is formulated using a typed system. Let us consider what happens if this option is chosen.

Adopting a typed metatheory opens up the possibility of selecting semantic values from the many types available. By contrast, an untyped metatheory can draw the possible semantic values only from the single type it comprises. The option of using semantic values from multiple types has become rather popular. Its main implementation is known as *higher-order semantics* (see, e.g., Boolos 1984, 1985; Rayo and Uzquiano 1999; Rayo 2002; Rayo and

Williamson 2003; Yi 2005, 2006; McKay 2006; Oliver and Smiley 2016). In the context of higher-order logic or plural logic, expressions of the object language are interpreted by means of corresponding expressions in the metalanguage. For example, the semantic value of a second-order variable is not a set or any other object. Rather, it is a second-level entity, which may be thought of as a *sui generis* property of objects described by primitive second-order vocabulary of the metalanguage. Similarly, a plural variable does not denote an object. Rather, it denotes some objects, which are specified by means of plural resources in the metalanguage.

From the perspective of a typed metatheory, the models of higher-order logic or plural logic appear richer than those of first-order logic, which only represents objects and their relations. The higher-order resources of the object language provide additional information: they represent higher-level entities (e.g. properties) and their relations. So differences between models of typed languages and those of untyped languages are significant, again from the perspective of a typed metatheory. In higher-order semantics, sentences of typed languages are generally not about the same entities as sentences of untyped languages. Some sentences of typed languages do not normally have one-sorted counterparts with the same truth conditions. Typed languages thus appear more expressive.

Recast in the present framework, higher-order semantics interprets expressions of the many-sorted object language by means of entities of the corresponding sort in the metatheory. For example, the denotation of a constant a^i is an entity of sort i , whereas quantification over a given sort is interpreted by the corresponding quantification in the metalanguage. The main point remains: from this semantic perspective, many-sorted languages appear more expressive than one-sorted languages. The truth conditions of some sentences of many-sorted languages go beyond those of sentences in one-sorted languages, involving entities of more than one sort.

Let us summarise the preceding discussion. We have reached a standoff between two consistent perspectives. If we assume an untyped metatheory like standard first-order set theory, many-sorted languages do not appear more expressive than one-sorted languages. From the perspective of a typed metatheory, however, the situation looks different: many-sorted languages appear more expressive than one-sorted languages. It follows that appealing to expressivity in support of many-sorted languages has little dialectical weight on its own. These languages appear expressively superior only if they are presupposed in the metatheory.

IV

Absolute generality. The semantic considerations of the previous section deliver a conditional conclusion: if we adopt a typed metatheory, we could come to see a many-sorted language as more expressive. Progress can be made by providing positive reasons to adopt a typed metatheory. A prominent argument due to Timothy Williamson (2003, pp. 425-7) aims to do precisely that. According to this argument, making reflective sense of absolute generality requires type distinctions. The argument can be given a general formulation, though we assume for definiteness that the relevant type is that of objects, as in Williamson's original discussion. For present purposes, the argument may be formulated as follows.

Let P be a predicate applying to objects. In model-theoretic semantics, we are interested in all possible interpretations of the language. So the interpretation of P will vary across models. In some model, the predicate applies to certain objects in the domain. In other models, it applies to different objects in the same domain or in another domain. Here is an intuitive principle about the range of possible interpretations: for any formula ϕ of the metalanguage, there is an interpretation according to which P applies exactly to objects in the domain that are ϕ . Call it the *liberal principle of interpretations*. To see how the principle works, suppose that the domain comprises the natural numbers and ϕ is the predicate ‘is prime’. Then the principle yields an interpretation according to which P applies exactly to the natural numbers that are prime. The principle is intuitive in part because it helps express the plenitude of interpretations of the object language, a model-theoretic ideal.

Let us state the principle a bit more precisely. For any domain D and for any formula ϕ , there is an interpretation J according to which P applies to an object in D if and only if the object satisfies ϕ . If absolute generality is possible, there is an unrestricted domain of quantification. This domain contains every object whatsoever so, when used in connection with the principle, it has the effect of rendering superfluous the restriction to domain D . This means that a specialised version of the principle holds: for any formula ϕ , there is an interpretation J according to which P applies to an object if and only if the object satisfies ϕ .

Assuming that interpretations are objects, a contradiction is obtained by taking ϕ to be ‘is not an interpretation x according to which P applies to x ’. Using an unrestricted domain, the liberal principle of interpretations entails that there is an interpretation J such that, for every object x , P applies to x according to J if and only if x is not an interpretation according to which P applies to x . It is possible to instantiate x with J , since J is assumed to be an object and is thus an appropriate instance of a universal quantifier ranging over objects. It is then straightforward to derive an inconsistency through classical reasoning: P applies to J according to J if and only if P does not apply to J according to J .

Williamson’s preferred response relies on a type distinction. He concludes that interpretations are not objects but primitive second-order entities, namely properties of objects (2003, pp. 452-4). As noted above, the argument can be given a general formulation, where the relevant predicate is $P^{<>}$, with i an arbitrary sort. The analogous response would be that interpretations are not of sort i but of different sort. So we have a positive reason to adopt a typed metatheory.

The argument presents us with a choice. We must reject one of three assumptions: the liberal principle of interpretations, the possibility of absolute generality, and the claim that interpretation are objects. Sacrificing the latter, as Williamson does, is certainly appealing. However, I believe that greater insight is gained if we analyse the liberal principle of interpretation.

The plenitude of model-theoretic interpretations is naturally expressed as the idea that, for any combination of *suitable semantic values* for the expressions of the object language, there is an interpretation that assigns that combination of semantic values to those expressions. In particular, a predicate can be interpreted by means of any suitable semantic value. Suppose the suitable semantic values for predicates are properties. Then plenitude requires:

(i) for any property, there is an interpretation assigning that property to P as semantic value.

On that interpretation, P applies to an object in the domain if and only if the object has the property assigned to P .

We can now factor the liberal principle of interpretations into two components. The first is (i). The second informs us about which properties there are:

(ii) for any formula φ , there is a property that an object has if and only if the object satisfies φ .

This more fined-grained analysis of the liberal principle of interpretations affords a different outlook on Williamson's argument. While (i) is very plausible, (ii) is a form of naïve comprehension and is, in light of the paradoxes, very controversial.

If one wishes to do justice to (ii), the introduction of type distinctions is a natural outcome. It is a safe and well-understood way to avoid inconsistency. As remarked earlier, Russell went as far as claiming that the introduction of type distinctions was *necessitated* by the paradoxes. Paradoxical notions (e.g. 'is predicated of itself', 'is not an interpretation x according to which P applies to x ') can no longer be expressed because they violate type distinctions. However, contrary to Russell's claim, there is nothing inevitable about this outcome.

There are non-classical attempts to capture naïve comprehension (e.g. Field 2004). But even if we adhere to classical logic, it is possible to develop a framework that, arguably, does justice to naïve comprehension (Schindler 2017). Moreover, one might also reject (ii) altogether, and this is what a number of type-free theories of properties do. Consider the analogy with set theory, where retaining naïve comprehension is not a decisive factor and where, in fact, the rejection of naïve comprehension prevails. That is because naïve comprehension is ruled out by a given conception of set, or because it is outweighed by broader theoretical considerations. The same holds for a theory of properties. Some conceptions of property (e.g. Linnebo 2006 and Roberts manuscript) do not sit well with naïve comprehension, and broader theoretical considerations might militate against it (e.g. Jubien 1989 and Menzel forthcoming).

In the context of absolute generality, there is an added complication. One needs a *universal* entity to represent an unrestricted domain of quantification. This is an entity that has every object as a member. If we operate within an untyped metatheory based on sets or properties, we would thus need either a universal set or a universal property. So there is an extra, non-trivial constraint in the development of the desired theory.

Let us take stock. Williamson's argument can be used to motivate a typed metatheory. However, I have argued that this move assumes a certain stance on the principle of naïve comprehension for the relevant entities. While viable, this stance is not inevitable. For example, if the relevant entities are properties, the assumption is that they satisfy a form of naïve comprehension and this is best captured by postulating a type distinction between objects and properties. Neither part of the assumption is forced on us. A one-sorted theory of properties might be able to do justice to the principle of naïve comprehension. Moreover, there might be a satisfactory one-sorted theory of properties that rejects the principle, as is the case for set theory. This is very much an open debate, with interesting new approaches being developed. No one can plausibly have the last word yet.

Can many-sorted logic say and prove more than one-sorted logic? In section III, we considered the thesis that many-sorted logic has more expressive power than one-sorted logic. The notion of expressive power received a semantic understanding. Here the thesis is assessed on a syntactic understanding of the notion, namely that many-sorted logic can say and prove more than one-sorted logic.

So understood the thesis seems trivially true. Any one-sorted system can be properly extended by a many-sorted system. We simply add new sorts and new predicates connecting them. This already yields more validities, owing to the logical axioms of the system.

However, the mere ability to say and prove more is not necessarily significant. An obvious case is the ordinary extension of a system by means of an explicitly defined predicate. Here the additional vocabulary and associated validities do not offer any significant expressive gain.

Recall the result, mentioned earlier, that the many-sorted relation of consequence has a one-sorted equivalent under the translation \bullet . That is, for any Σ and σ , $\Sigma \models \sigma$ in many-sorted logic if and only if $\Sigma^{\bullet} \cup \Delta \models [\sigma]^{\bullet}$ in one-sorted logic, where Δ conveys that each S_i has a witness. Both systems are sound and complete with respect to the model-theoretic semantics introduced in section III. This fact can be used to show that $\Sigma \vdash \sigma$ if and only if $\Sigma^{\bullet} \cup \Delta \vdash [\sigma]^{\bullet}$. This means that deducibility in a many-sorted system corresponds exactly to a certain kind of deducibility in a one-sorted system. This fact is significant.

A central motivation for the programme of higher-order metaphysics, which advocates the use of typed resources, has to do with the ability to regiment metaphysical reasoning (see Fritz and Jones forthcoming b). Examples include metaphysical claims and arguments involving quantification over propositions, properties, or modalities. Typed languages are shown to provide a fruitful framework for a systematic formalization of this kind of metaphysical reasoning (e.g. Dorr 2016).

The deductive correspondence between reasoning in many-sorted logic and in one-sorted logic suggests that untyped languages are expressively powerful in the sense relevant here. Any many-sorted system that faithfully represents reasoning in a metaphysical domain can be matched inside a one-sorted system. Thus the regimenting potential of typed languages does not exceed that of untyped languages. It should be emphasised that, by itself, this conclusion does not undermine the value of higher-order metaphysics. It might well be that typed languages provide a particularly good, perhaps even heuristically optimal, environment for metaphysical theorising. Still, any deductive insights obtained in that environment can be recovered in an untyped theory.

What applies to higher-order metaphysics also applies, *mutatis mutandis*, to any other area of inquiry that can be faithfully regimented in a many-sorted system. The deductive correspondence in question is a fully general metatheoretic result.

Nonsense and unnecessary theoretical choices. It was remarked in the introduction that the notion of expressivity plays a puzzling role in this debate. Typed languages seem more expressive, and yet a longstanding criticism against typed languages is that they have expressive limitations. According to critics, there are natural generalisations that cannot be expressed because they require quantification over types, which is not available in standard typed languages. In particular, there are ‘deep and interesting semantic insights that cannot properly be expressed’ (Linnebo 2006, p. 154), including the semantic principle of compositionality. (For discussion, see Linnebo 2006, section 6.4, and Krämer 2013.) It seems that these expressive limitations can be overcome in untyped languages, where types are objects and thus belong to the domain of quantification. In this respect, untyped languages appear more expressive than typed ones. In this section, I discuss some objections alleging that the expressivity of untyped language goes in fact too far, permitting the formalisation of nonsense or forcing us to make unnecessary theoretical choices.²

Switching from a many-sorted system to one-sorted system opens up new theoretical questions, as emphasised already by Quine (1956, p. 268-9). Questions that might not be grammatically formulated in the many-sorted system become expressible in the one-sorted system. The type-theoretic solution to the paradoxes advocated by Russell relies essentially on making problematic expressions (e.g. self-application) ungrammatical. In a many-sorted system rendering the theory of types, the relation of predication would not permit as arguments two variables of the same sort. Once we switch to a one-sorted system, self-predication becomes again expressible, and hence we can inquire into whether there is anything that can be predicated of itself. This phenomenon has been used to argue against untyped systems.

Consider a natural language like English. There is no significant morphological or syntactic distinction between, for instance, predicates that apply to concrete objects and those that apply to abstract objects. A speaker of the language can freely apply predicates of each kind to any class of entities. The results might be typical examples of category mistakes: ‘every prime number is green’ or ‘every chair has a prime factorisation’. A many-sorted language has the resources to avoid these results. For example, one sort could be reserved for the natural numbers, while another sort could be used for concrete entities such as chairs. Then appropriate sortal restrictions ensure that no variables can simultaneously occupy the argument positions of the predicates ‘is prime’ and ‘is green’, or of the predicates ‘is a chair’ and ‘has a prime factorisation’. These category mistakes can no longer be grammatically formulated. If we adopt a one-sorted language, this is obviously not an option. There is only one sort and thus no natural way to impose restrictions on how predicates can be linked by variables. Category mistakes can be expressed, just as in natural language.

How can these observations be turned into a philosophical argument against untyped languages? Let us consider two ways.

First, one might hold that, by allowing category mistakes, the one-sorted language permits the expression of nonsense, and that this is undesirable. There are several problems with this line of thought. The main problem is that there are good reasons to reject the view that category mistakes are nonsensical. This view is admittedly natural and has enjoyed a lot of support. However, a range of considerations strongly suggest that category mistakes are in fact meaningful (see Magidor 2013, ch. 3). To begin with, it is generally accepted that

² Related considerations can be found, for example, in Jones 2018 and Bacon forthcoming.

category mistakes are grammatical. The one-sorted framework respects this fact, unlike the use of sorts to ban category mistakes as ungrammatical. Moreover, the assumption that category mistakes are grammatical but meaningless clashes against linguistic evidence and theoretical considerations. Let us summarise the key points (again, see Magidor, 2103, ch. 3, for a detailed discussion): (a) category mistakes are context-sensitive in a way that meaningfulness is not; (b) category mistakes can be felicitously embedded in attitude reports, but nonsense cannot; (c) a compositional theory of meaning will likely imply that many paradigmatic cases of category mistakes have meaning; (d) uses of category mistakes in figurative speech are best explained if we assume that category mistakes are meaningful. While it may be debated how compelling each point is, they collectively make a strong case for the meaningfulness of category mistakes. It follows that formalisations permitting the expressions of category mistakes might be theoretically desirable rather than undesirable. Of course, category mistakes tend to be practically pointless: there is usually very little value in attempting to settle them. But protecting us from futility has never been the goal of a formal system.

A second way to argue against untyped systems focuses on the theoretical choices we face if we adopt a single style of variables, choices thought to be undesirable. Barrett and Halvorson write: ‘eliminating sort distinctions forces us to make unnecessary conventional choices about how to extend predicates beyond their original range of application’ (2017, p. 3578). Category mistakes give rise to such choices. If we can say ‘every prime number is green’ or ‘every chair has a prime factorisation’, we face the question whether we should take these sentences to be true or false. By contrast, the many-sorted framework can sidestep these conventional choices: ‘it does not force us to apply predicates in cases where we have no good reason to say that they do (or do not) hold of the items in question’ (Barrett and Halvorson 2017, p. 3578). In our examples, we can simply avoid the problem by using one sort for the natural numbers and another sort for concrete entities, with predicates restricted appropriately.

How significant is this problem? Let us consider two perspectives from which the question may be approached. From an axiomatic perspective, we are not always forced to make the relevant conventional choices. We may be happy for our theories to remain incomplete in these ways, given that the sentences left open are uninteresting. When conventional choices are required, either because we aim for completeness or because settling certain questions has theoretical value, there is no technical obstacle. Barrett and Halvorson observe that care is needed to ensure that the theory is extended consistently: on pain of contradiction, we cannot stipulate that all the sentences in question are false (2017, p. 2577). But we know that, if care is taken, we can succeed. By Lindenbaum’s lemma, any consistent first-order theory has a complete consistent extension.

Now consider the problem from a model-theoretic perspective. Any model of a consistent theory takes a stand with respect to the truth value of every sentence in the language. So, in building a model (e.g. to prove the consistency of a theory), conventional choices might have to be made. For instance, we need to fix the extension of every predicates. Still, these choices occur in the metatheory and are only relative to the *model*. They need not settle how the one-sorted *theory* is actually extended. In particular, these choices need not fix how the predicates’ ranges of application are extended in the theory.

Note that the issue under discussion is not confined to category mistakes but is more general. It concerns genuine theoretical possibilities that are not expressible in a many-sorted setting

but emerge in a one-sorted setting.³ While some of these cases seem to require unnecessary conventional choices, not all do. Some cases concern theoretical questions we wish to express and even settle in a certain way, because of a powerful intuitive pull towards one particular answer. This suggests that the one-sorted framework can afford greater expressivity of a kind that matters, as stressed in Florio and Linnebo 2021 (sec. 11.6 and 11.7). Quine (1956) offers some example. One concerns the possibility of identifying some classes with sets after translating a two-sorted class theory into a one-sorted theory (p. 269). Another example concerns the possibility of identifying empty classes from different types after translating the theory of types into a one-sorted set theory (p. 274).⁴ Let me provide a different example from plural logic.

The standard formulation of plural logic is two-sorted. There is one sort for objects and another sort for plural talk about objects. The distinction is often marked by the use of single variables (x, y, z, \dots) for the first sort and double variables for the second sort (xx, yy, zz, \dots). So we read ‘ $\exists xPx$ ’ as ‘something is P ’ and ‘ $\exists xxQxx$ ’ as ‘some things are (collectively) Q ’. The system includes a predicate for plural membership, corresponding to ‘is one of’ and relating variables of the first sort to variables of the second sort. Consider any object a . The axioms of plural logic entail that there are some things such that a and only a is one of them. Call these things aa . Roughly put, we have a ‘plurality’, aa , with one and only one member, namely a . An intuitive question is whether a and aa are identical. The answer we would intuitively like to give is that they are indeed identical. After all, there is nothing more to aa than a . The question cannot be addressed in the two-sorted setting, because there is no cross-type identity relation. By contrast, the desired identification can be expressed in a one-sorted system for plural logic (with special predicates ‘is an object’ and ‘is a plurality’), where it is perfectly consistent to assume that a is identical to the plurality whose only member is a . It is interesting to note here that the one-sorted approach to plural logic has gained popularity in recent literature (see Oliver and Smiley 2016, ch. 15; Florio and Linnebo 2021, sec. 11.7).

VII

Conclusion. Which form should our theories take? I have argued that expressivity does not favour typed languages over untyped ones. First, appealing to greater expressive power in support of typed languages is not effective. On a semantic understanding of expressive power, we reach at best a conditional conclusion: we can come to see a typed language as more expressive if such a language has already been adopted as a metatheoretic framework. On a syntactic understanding, the potential of typed languages to capture deductive reasoning does not exceed that of untyped languages. Second, untyped languages are not too expressive in problematic ways, despite permitting the formalisation of category mistakes and opening up new theoretical questions.

³ Functions provide a clear example. Consider a unary function $f^{<i_1>}$ taking an input of sort i_1 and returning an output of sort i_2 . In a many-sorted setting, the domain of the function is confined to sort i_1 . When we switch to a one-sorted setting, the domain naturally expands to the entire domain. One choice is whether to treat the function as partial or total. If the function is total, additional choices concern the function’s behaviour over new inputs.

⁴ In that case, Quine points out, it is also possible to identify an individual with its singleton set. Though natural and useful in some contexts (e.g. Schwarzschild 1996, p. 1; Fine 2010, pp. 574-5), this identification does not generally have a strong intuitive pull.

Of course, typed languages might still be preferable because of other theoretical or practical considerations (e.g. Manzano 1996, sec. 1.6; Halvorson 2019, sec. 5.4; Fritz and Jones forthcoming a). Here I have offered only a partial vindication of untyped languages. As far as expressivity is concerned, one may continue to follow Quine’s advice to ‘pool types and get on with homogeneous variables’ (Quine and Carnap 1990, p. 353).⁵

*Department of Philosophy
University of Birmingham
Edgbaston
Birmingham B15 2TT
United Kingdom*

*Department of Philosophy, Classics, History of Art and Ideas
University of Oslo
Postboks 1020
Blindern 0315
Oslo
Norway*

REFERENCES

- Bacon, Andrew forthcoming: ‘A Case for Higher-Order Metaphysics’. In Fritz and Jones forthcoming a.
- Barrett, Thomas William and Hans Halvorson 2016: ‘Morita Equivalence’. *The Review of Symbolic Logic*, 9(3), pp. 556-82.
- Barrett, Thomas William and Hans Halvorson 2017: ‘Quine’s Conjecture on Many-sorted Logic’. *Synthese*, 194(9), pp. 3563-82.
- Benacerraf, Paul and Hilary Putnam (eds.) 1983: *Philosophy of Mathematics: Selected Readings*, 2nd edn. Cambridge: Cambridge University Press.
- Boolos, George 1984: ‘To Be Is to Be a Value of a Variable (or to Be Some Values of Some Variables)’. *Journal of Philosophy*, 81(8), pp. 430-49.
- Boolos, George 1985: ‘Nominalist Platonism’. *The Philosophical Review*, 94(3), pp. 327–44.
- Button, Tim and Robert Trueman 2022: ‘Against Cumulative Type Theory’. *The Review of Symbolic Logic*, 15(4), pp. 907-949.

⁵ For helpful comments and discussion, I would like to thank Nick Jones, Øystein Linnebo, Carlo Nicolai, Stewart Shapiro, and Eric Snyder. This work was supported by the Arts and Humanities Research Council [grant number AH/V015516/1].

- Corfield, David 2020: *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. Oxford: Oxford University Press.
- Degen, Wolfgang and Jan Johannsen 2000: ‘Cumulative Higher-Order Logic as a Foundation for Set Theory’. *Mathematical Logic Quarterly*, 46(2), pp. 147-170.
- Dorr, Cian 2016: ‘To be F is to be G’. *Philosophical Perspectives*, 30, pp. 39-134.
- Enderton, Herbert B. 2002: *A Mathematical Introduction to Logic*, 2nd edn. San Diego: Academic Press.
- Field, Hartry 2004: ‘The Consistency of the Naive Theory of Properties’. *The Philosophical Quarterly*, 54(214), pp. 78-104.
- Fine, Kit 2010: ‘Towards a Theory of Part’. *The Journal of Philosophy*, 107(11), pp. 559-589.
- Florio, Salvatore and Nicholas K. Jones 2021: ‘Unrestricted Quantification and the Structure of Type Theory’. *Philosophy and Phenomenological Research*, 102(1), pp. 44-64.
- Florio, Salvatore and Øystein Linnebo 2021: *The Many and the One: A Philosophical Study of Plural Logic*. Oxford: Oxford University Press.
- Fritz, Peter and Nicholas K. Jones (eds.) forthcoming a: *Higher-Order Metaphysics*. Oxford: Oxford University Press.
- Fritz, Peter and Nicholas K. Jones forthcoming b: ‘Higher-Order Metaphysics: An Introduction’. In Fritz and Jones forthcoming a.
- Gödel, Kurt 1944: ‘Russell’s Mathematical Logic’. Reprinted in Benacerraf and Putnam 1983, pp. 447-469.
- Halvorson, Hans 2019: *The Logic in the Philosophy of Science*. Cambridge: Cambridge University Press.
- Jones, Nicholas K. 2018: ‘Nominalist Realism’. *Noûs*, 52(4), pp. 808-835.
- Jubien, Michael 1989: ‘On Properties and Property Theory’. In Gennario Chierchia (ed.), *Properties, Types and Meaning, Vol. I: Foundational Issues*, pp. 159-175. Dordrecht: Kluwer Academic Publishers.
- Krämer, Stephan 2013: ‘Semantic Values in Higher-Order Semantics’. *Philosophical Studies* 168(3), pp. 709-724.
- Krämer, Stephan 2017: ‘Everything, and Then Some’. *Mind* 126(502), pp. 499-528.
- Ladyman, James and Stuart Presnell 2018: ‘Does Homotopy Type Theory Provide a Foundation for Mathematics?’. *The British Journal for the Philosophy of Science*, 69(2), pp. 377-420.

- Linnebo, Øystein 2006: 'Sets, Properties, and Unrestricted Quantification.' In Agustín and Gabriel Uzquiano Rayo (ed.), *Absolute Generality*, pp. 149-78. Oxford: Oxford University Press.
- Linnebo, Øystein and Agustín Rayo 2012: 'Hierarchies Ontological and Ideological'. *Mind*, 121(482), pp. 269-308.
- Magidor, Ofra 2013: *Category Mistakes*. Oxford: Oxford University Press.
- Manzano, María 1996: *Extensions of First-Order Logic*. Cambridge: Cambridge University Press.
- McKay, Thomas J. 2006: *Plural Predication*. Oxford: Oxford University Press.
- Menzel, Christopher forthcoming: 'Pure Logic and Higher-Order Metaphysics'. In Fritz and Jones forthcoming a.
- Oliver, Alex and Timothy Smiley 2016: *Plural Logic: Second Edition, Revised and Enlarged*. Oxford: Oxford University Press.
- Quine, W. V. 1956: 'Unification of Universes in Set Theory'. *The Journal of Symbolic Logic*, 21(3), pp. 267-279.
- Quine, W. V. 1986: *Philosophy of Logic*, 2nd edn. Cambridge, MA: Harvard University Press.
- Quine, W. V. and Rudolf Carnap 1990: *Dear Carnap, Dear Van. The Quine-Carnap Correspondence and Related Work*, ed. and introd. Richard Creath. Berkeley and Los Angeles, CA: University of California Press.
- Russell, Bertrand 1908: 'Mathematical Logic as Based on the Theory of Types'. *American Journal of Mathematics*, 30(3), pp. 222-62.
- Shapiro, Stewart 1991: *Foundations without Foundationalism: A Case for Second-order Logic*, Oxford: Oxford University Press.
- Skiba, Lucas 2021: 'Higher-Order Metaphysics'. *Philosophy Compass*, 16(10), pp. 1-11.
- Rayo, Agustín and Gabriel Uzquiano 1999: 'Toward a Theory of Second-Order Consequence'. *Notre Dame Journal of Formal Logic*, 40(3), pp. 315-25.
- Rayo, Agustín and Timothy Williamson 2003: 'A Completeness Theorem for Unrestricted First-Order Languages'. In JC Beall (ed.), *Liars and Heaps*, pp. 331-56. Oxford: Oxford University Press.
- Roberts, Sam manuscript: 'The Iterative Conception of Properties'.
- Schindler, Thomas 2017: 'Classes, Why and How'. *Philosophical Studies*, 176, pp. 407-35.
- Schwarzschild, Roger 1996: *Pluralities*, Dordrecht: Kluwer Academic Publishers.

Barry Smith 2005: 'Against Fantology'. In Johann C. Marek and Maria E. Reicher (eds.), *Experience and Analysis*, pp. 153-170. Vienna: HPT&ÖBV.

Trueman, Robert 2021: *Properties and Propositions: The Metaphysics of Higher-Order Logic*. Cambridge: Cambridge University Press.

Yi, Byeong-Uk 2005: 'The Logic and Meaning of Plurals. Part I.' *Journal of Philosophical Logic*, 34, pp. 459-506.

Yi, Byeong-Uk 2006: 'The Logic and Meaning of Plurals. Part II.' *Journal of Philosophical Logic*, 35, pp. 239-88.

Williamson, Timothy 2003: 'Everything'. *Philosophical Perspectives*, 17, pp. 415-65.

Williamson, Timothy 2013: *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Wittgenstein, Ludwig 1979: *Notebooks 1914-1916*, 2nd edn. Oxford: Basil Blackwell.