1–20
https://doi.org/10.1093/philmat/nkad005
Philosophia Mathematica

OXFORD

Refocusing Frege's Other Puzzle: A Response to Snyder, Samuels, and Shapiro

Thomas Hofweber*

Department of Philosophy, University of North Carolina at Chapel Hill, Chapel Hill, N.C. 27599, U.S.A.

ABSTRACT

In their recent article 'Resolving Frege's other Puzzle' Eric Snyder, Richard Samuels, and Stewart Shapiro defend a semantic type-shifting solution to Frege's other Puzzle and criticize my own cognitive type-shifting solution. In this article I respond to their criticism and in turn point to several problems with their preferred solution. In particular, I argue that they conflate semantic function and semantic value, and that their proposal is neither based on general semantic type-shifting principles nor adequate to the data.

1. FREGE'S OTHER PUZZLE

Number words in natural language are of potentially great significance for the philosophy of arithmetic. If numerical symbols like '4' are merely symbolic notation for the natural number word 'four', then semantic features of number words can inform us of semantic features of number symbols in arithmetic. If so, and if number words in natural language aim to refer to objects, then arithmetic would aim to describe just these objects. But if number words in natural language are not referential at all, then this might also carry over to arithmetic, which then might do something completely different from aiming to describe objects. Not everyone will accept a close connection between arithmetical symbols and natural-language number words. But there surely is some connection or other, since it does not appear to be an accident that both are pronounced just the same way, systematically across all number word / number symbol pairs. But even assuming there is a close connection, how does the semantics

*Orcid.org/0000-0001-7095-6559. E-mail: hofweber@unc.edu

of number words in natural language inform us about the semantics of number symbols in arithmetic?

Back in [1884], Frege raised a puzzle about number words in natural language, which we will look at more closely just below and call 'Frege's other Puzzle', not to be confused with Frege's puzzle about identity statements.[1] Frege pointed out that number words can appear in natural language in apparently quite different ways: sometimes like adjectives, as in 'Jupiter has four moons' and sometimes like names, as in 'The number of moons of Jupiter is four'. The former suggest that number words in natural language do not aim to pick out objects, but rather modify nouns, while the latter suggest that they do aim to pick out objects, as names commonly do. How should we understand number words in natural language in light of this? Is their appearance as adjectives or their appearance as names more revealing about their semantic function?

In past work I have defended a view that takes number words to be primarily adjectives or determiners, but not names of objects. I proposed two different explanations of why number words nonetheless syntactically appear like names: one concerned syntactic displacement in a focus construction, see [Hofweber, 2007], the other concerned cognitive type coercion; see [2005]. How all this relates to the philosophy of arithmetic is spelled out in [2016].

In their recent article [2022], Eric Snyder, Richard Samuels, and Stewart Shapiro (*S3* from now on for brevity) critique my proposed solution to Frege's other Puzzle, and offer their own preferred solution instead. S3 defend a polymorphic solution to the puzzle: number words can take on different types in natural language, since general type-shifting principles support this. In this essay I hope to make clear why I find S3's solution unsatisfactory, and why I think that it misses the key target that needs to be aimed at to solve Frege's other Puzzle properly. Along the way I will also respond to some of their criticism of my own view.

Frege's other Puzzle is a puzzle about number words in natural language. There are several related, but different, aspects to this puzzle. Most centrally we can distinguish the following three aspects of what *prima facie* appears to be puzzling about number words in natural language:

1. **Syntax:** Number words seem to occur in different syntactic positions: that of an adjective and that of a singular term or noun phrase.
2. **Semantic values:** Number words seem to have different semantic values in different occurrences: that of type e and that of higher types.
3. **Semantic function:** Number words seem to have different semantic functions in different occurrences: that of reference and that of modifying a noun.

These three aspects of the puzzle are, of course, not unrelated. However, how they are related to one another is not as straightforward as it might seem, and a

[1] See [Hofweber, 2005] for this terminology.

good part of my discussion in this paper concerns in particular the relationship between the second and the third aspects: the relationship between the type of semantic value that an expression gets assigned and its semantic function. The simplest view of the relationship between the three parts of the puzzle closely connects all of them: syntactic position is closely associated with type of semantic value, which in turn is closely associated with semantic function. I hope to make clear below that although one of these close relations is indeed reasonable, the other is not.

S3 aim to establish that number words do not have a single semantic function and they hope to explain with semantic type-shifting principles why that is so. But much of their actual discussion does not concern semantic function directly, but semantic values: which semantic values number words can have, and what semantic type-shifting principles relate these different semantic values to one another. They move rather freely between talk of types of semantic values and of kinds of semantic functions, as I will discuss in more detail below. I think this is a key mistake in their argument. Semantic values need to be more sharply distinguished from semantic function, and one cannot directly draw conclusions from one about the other. This mistake is highlighted in particular in their discussion of the coordination objection to my own view. I would like thus to elaborate first on this difference and why no immediate inference from semantic values to semantic function is justified. In the process I also hope to make clear what the true difference between S3's semantic type-shifting account and my own cognitive type-shifting account comes down to.

2. WHAT IS AT ISSUE?

When one speaks of the semantic type of an expression, it is natural to understand this as what kind of an expression it is semantically: a predicate, noun phrase, determiner, *etc.* This in turn can easily be associated with the semantic function of the expression. But in much of our debate the notion of a semantic type has a more technical meaning, one related to type theory. On that way of understanding it, it is closely associated with what kind of semantic value the expression gets assigned in compositional semantic theories. It is important to make this contrast clear, and so I will briefly elaborate on it.

Types, in the technical sense, can be assigned to expressions in natural languages in two different ways: as syntactic types and as semantic types. Syntactic types simply correspond to how expressions can be combined to form a grammatical sentence. If we take the type of a sentence to be t, then combining an expression of type e with one of type $\langle e, t\rangle$ leads to type t and thus a grammatical sentence. Semantic types, on the other hand, are semantic values that get assigned to expressions, normally seen as functions from one domain to another one. In this sense, for an expression to have semantic type $\langle e, t\rangle$ means that the semantic value assigned to that expression in a particular compositional semantic theory is a function from the things of type e to the things of type t.

It is natural and reasonable to associate syntactic position closely with syntactic type. Being of a particular syntactic category in a sentence closely

corresponds to having a particular syntactic type. So understood, grammatical sentences require that the syntactic types of the expressions that occur in them combine properly to lead to type t overall. Furthermore, it is reasonable to associate syntactic and semantic types closely. An expression that grammatically belongs to type $\langle e, t\rangle$ most conveniently gets assigned a semantic value of type $\langle e, t\rangle$. Those semantic types will combine with each other in just the right way if the corresponding syntactic types combine in the right way. This association is so common that distinguishing syntactic and semantic types is sometimes neglected, but conceptually they need to be distinguished and they can in principle come apart.

The situation is trickier when it comes to semantic function and how it relates to semantic and syntactic types. By 'semantic function' I simply mean what an expression aims to do semantically. Paradigmatic possible answers are aiming to refer to an entity, or modifying a noun meaning, or to predicate of something, and so on. The standard view here, which I take to be correct, is that some phrases aim to refer, but most do not. Names, paradigmatically, aim to refer, but expressions like 'very', 'some', *etc.*, do not aim to refer. They do something else semantically, and thus have a different semantic function than reference.

A key part of Frege's other Puzzle is to say what the semantic function of number words is, and whether they have a single semantic function in their different occurrences. This is the aspect of Frege's other Puzzle which is the most important one for the philosophy of mathematics. Frege's other Puzzle is primarily a puzzle about natural language, but it is philosophically more significant than other puzzles that are merely about natural language. The significance of the semantic function aspect of Frege's other Puzzle is, naturally, that if number words have the semantic function of reference and number symbols in mathematics are semantically derivative on number words in natural language, then the truth of arithmetic seems to require the existence of the referents of number words. On the other hand, if number words do not have the function of reference and number symbols in arithmetic are derivative on number words in natural language, then the truth of arithmetic might not require the existence of anything.

This second option was the ultimate conclusion of my own attempt to solve Frege's other Puzzle. I argued that number words do not have the semantic function of reference, that number symbols are derivative on number words, and that the proper understanding of the content of arithmetical statements makes clear that arithmetic enjoys strong ontological independence: the truth of arithmetic is independent of what exists, including of how many things exist. The focus of this solution to Frege's other Puzzle was thus a claim about semantic function, and the significance for the philosophy of arithmetic came from a conclusion about the semantic function of number words. I tried to work this out in [Hofweber, 2016, chaps. 5, 6]. Along the way I defended a cognitive type-shifting approach for certain occurrences of number words, hoping to explain some of the syntactic aspects of Frege's other Puzzle. S3 disagree and defend a semantic type-shifting approach instead. But the disagreement between my

own approach and the one of S3 is not quite where it might seem and how it is presented in [Snyder *et al.*, 2022].

Everyone can accept that semantic type-shifting occurs, and so do I. Whether expressions in natural language can take on different semantic and syntactic types can be uncontroversial for our purposes. The only way to reject it reasonably is to reject the whole framework of theorizing in terms of types. But if we theorize that way at all, then it is clear that expressions like 'or' can occur with different syntactic and semantic types associated with them: they can combine NPs (John or Mary), determiners (two or three), sentences, and so on. Plausibly there will be general principles that determine in which types 'or' can occur, and how the semantic values of these different types relate to each other. These principles will not be uniquely tied to 'or', but be general principles that apply to expressions of this type: there will be similar ones for 'and' and so on. Thus general type-shifting principles will relate occurrences of expressions of a general type to one another, ones that will give us insights into both their syntactic flexibility and the variety of semantic values they can have on these occasions. This so far is all uncontroversial. As I discuss in [Hofweber, 2016, pp. 137–142], semantic type-shifting even for number words is compatible with my own preferred cognitive type-shifting view. There can be, and almost certainly are, general type-shifting principles which apply to number words and result in number words being of variable type. However, I argued that such general semantic type-shifting principles by themselves will not explain everything that needs to be explained about number words.

The issue is not whether semantic type-shifting occurs, but rather whether language-general semantic type-shifting principles by themselves solve Frege's other Puzzle. Do general type-shifting principles by themselves explain what needs to be explained, or is there something specific about number words that goes beyond what general type-shifting principles would license? My own positive proposal is a type-shifting proposal as well. But it holds that there is something specific about talk and thought about natural numbers, something that is not determined by general type-shifting principles, that we need to draw upon to solve Frege's other Puzzle. So, a purely semantic type-shifting proposal would hold that there are language-general type-shifting principles, ones that apply to categories of expressions like determiners, modifiers, etc., and that reliance on those principles alone explains the various features of natural-language number words. The cognitive type-shifting proposal that I favor rejects this. Not because it rejects semantic type-shifting — no one should reject that in general — but because it holds that general type-shifting principles by themselves do not explain what needs to be explained. There is something specific concerning number words, something tied to a cognitive difficulty we face when thinking about arithmetic, that needs to be relied upon in addition.

My own view has always been, and still is, that relying on purely semantic type-shifting principles is the second-most promising attempt to solve Frege's

other Puzzle.[2] But the even better, and I believe correct, solution goes beyond it and relies on cognitive type-shifting as well. My argument for this has simply been that the pure semantic type-shifting approach makes the wrong predictions about particular examples, an argument that S3 discuss and which I will revisit shortly, and that cognitive type-shifting can explain a particular feature of how arithmetical equations are read and how that reading changes over a speaker's mathematical education, something S3 do not discuss. The idea of this latter argument was that somewhat surprisingly, arithmetical equations like '$2 + 2 = 4$' have both singular and plural readings as 'two and two is four' and 'two and two are four', both of which are widely used. But what is more, there is a difference in whether the plural or the singular is used across the mathematical education of an individual. In early mathematics education the plural is more prominent than the singular, and later it shifts. This, I proposed, is to be explained by the singular leading to a cognitive advantage, one that corresponds to a type-lowering for cognitive reasons. This type-shifting was supposed to go beyond what general semantic type-shifting principles would license, and was distinctly tied to thinking and talking about natural numbers.

As I see it, there are two main differences that are at issue between S3 and myself, but they are not quite the differences S3 take us to have. Our disagreement is not about whether semantic type-shifting occurs at all, nor whether expressions in natural language can occur in various types, including number words. Instead they are, first, whether general semantic type-shifting principles by themselves explain the features of number words in natural language, and, second, what the relationship between semantic types and semantic function is. I will discuss the second one next, and then return to the first one below with more details.

3. SEMANTIC TYPES AND SEMANTIC FUNCTION

I suspect my main disagreement with S3 concerns the connection between semantic types and semantic function. My own view is that those two are reasonably independent: they are rather different levels of description of the semantic aspect of language and no easy inference can be drawn from one to the other. S3 hold the opposite view, and take there to be a clear and close connection. They frequently draw inferences from an expression's having a semantic value of a particular type to its having a particular semantic function, and they affirm this presumed connection explicitly at various places, as when they say "... 'four' takes on different semantic types, and thus functions ..." (p. 80) and " 'four' comes to serve both referential and non-referential semantic functions, thanks to independently available, and general, type-shifting operations" (p. 80). In this section I would like to work out why I think this is mistaken. This also directly relates to one of S3's arguments against my own view: the coordination argument.

---

[2] See, for example, [Hofweber, 2005, p. 205].

The semantic function of a phrase, *i.e.*, what it aims to do semantically, includes such things as reference, modification, predication, and others. The standard view here is that semantic function is diverse in the sense that not every phrase has the same semantic function: 'John' aims to refer, but 'some' does not. A less standard, but I suspect still popular, view is that every expression has one and only one function. This is mostly what is at issue now, in particular how this question relates to a possible variety of semantic values that an expression might get.

Talk of the semantic function of a phrase is one aspect or level of its semantic description. But it is not all that can be said about a phrase that is semantically relevant. Another concerns its semantic value in a compositional semantic theory. Every phrase has a semantic value in standard compositional semantic theories. In typed theories like the ones under considerations here, semantic values are things that belong to a certain type, standardly the type corresponding to the syntactic type of the expression. It is common in linguistic semantics to say that a phrase *denotes* its semantic value. In this sense 'very' denotes a higher-type object. But so understood denotation is not the same as reference. Every expression denotes its semantic value, in this linguistic sense, but only some expressions refer. Although I take it this distinction should be uncontroversial, it is not always made as clearly as it should. Denotation applies uniformly to all expressions: each one denotes whatever semantic value it has, in this sense of 'denotation' common in linguistic semantics. But semantic function is not uniform, and in particular not uniformly that of reference: some expressions aim to refer, while others do not. Talk of denotation in the linguistic sense thus belongs to the level of semantic description which is closely associated with semantic values, but not semantic function.

The question remains what the relationship is between the type of semantic value of a phrase and the kind of semantic function that phrase has. Here one could defend a close connection, as S3 do, or a loose connection, as I prefer. The close-connection view holds that there is a close association between semantic value and semantic function: From the fact that a phrase gets a certain semantic value one can infer that is has a certain semantic function, at least on that occurrence. In particular, a straightforward version of this view would hold that an expression of type e is referential, and it refers to the object which is its semantic value. An expression of a higher type is not referential, it has a different semantic function instead.

I believe there are very good reasons to reject the close-connection view, and those reasons are illustrated by an issue that goes back to the very beginning of this general approach to semantics: Montague's treatment of proper names. Montague proposed to treat proper names as generalized quantifiers, and assigned them semantic values of type $\langle\langle e, t\rangle, t\rangle$. His reason for this was closely related to the issue brought up in the 'coordination argument' discussed by S3: sometimes we combine proper names with quantifiers, and to have a uniform compositional semantic analysis of such complex noun phrases, it is helpful to have a uniform type of semantic value for noun phrases, be they

quantificational or not. Thus to treat 'John and every woman' properly in a sentence like

(1) John and every woman left

it is advisable to assign semantic values of the same type to 'John' as well as to 'every woman', and to take 'and' simply to combine two semantic values of the same type into one of that very type again. This works great for doing compositional semantics, and is considered a real breakthrough in semantic theorizing.[3]

But does this mean that Montague proposed that proper names do not refer to objects? Does it mean that 'John' does not refer to John, but to a set of properties instead? The answer must be 'no'. 'John' refers to John, but having its referent as its semantic value is inelegant and not as useful for semantic theorizing as using the set of John's properties instead. Both of them would work to generate the truth conditions correctly. Instead of holding that 'John is F' is true just in case John has the property of being F we can, equivalently, hold that it is true just in case being F is a member of the set of John's properties. Those are the same truth conditions, and thus assigning these semantic values allows us to generate the truth conditions correctly in both cases. But when it comes to more complex cases, ones involving 'John and every woman', the second, higher-type approach is more elegant and uniform. 'John' and 'every woman' get assigned the same type of semantic value, but this does not mean that they have the same semantic function. In particular, it does not mean that 'John' is not referential, nor that 'every woman' is referential. Their semantic functions are different, but their semantic values are of the same type in order to have the most uniform and elegant way to generate the truth conditions correctly. The goal in assigning semantic values is to generate the truth conditions correctly, and in the most elegant and uniform way possible. This goal is most easily met by assigning expressions with different semantic functions semantic values of the same type.

Overall, then, it would be a mistake to understand Montague's proposal to include that 'John' and 'every woman' have the same semantic function, since they combine with 'and' to form a complex noun phrase and thus they each get a semantic value of the same type. Anyone who would suggest that 'every woman' is referential, since it combines with 'John' via conjunction, or that 'John' is not referential, since it combines with 'every woman,' would be making just this mistake about conflating semantic values and semantic functions. In particular, they would be making the mistake of thinking that semantic function is determined by the type of semantic value. And just this mistake occurs in S3's 'coordination argument' that number words must be referential, since they can combine with names via 'and', as illustrated in their example (14) [2022, p. 68]:

> Fido and four are two of my favorite things.

[3] See [Montague, 1974] and [Barwise and Cooper, 1981].

This is obviously not exactly a natural English example, but that is not the main flaw of S3's use of it. Instead it is this: to draw any conclusions from an example like this about the semantic function of number words is simply a mistake based on conflating semantic function and type of semantic value, just as it would be a mistake to draw a conclusion about the referentiality of names or quantifiers from Montague's example.

This independence of semantic function and type of semantic value also applies to expressions that have a semantic value of type e. To have this semantic value in a semantic theory does not mean that the relevant phrase is referential and refers to an entity of type e. To be sure, this might often be the case, but nonetheless, these two things need to be distinguished. Type e semantic values are conveniently assigned to expressions that refer, the semantic value might often be the referent itself. But this is neither required nor guaranteed.

In the same spirit, it would be a mistake to criticize Montague's semantic type assignment to names by pointing out that it treats referential expressions as having a type other than e. That is not a legitimate criticism, since associating semantic function with a particular type of semantic value is not the goal of semantic type assignment. What would be legitimate is to criticize Montague's proposal since it treats NPs as too uniform, being unable to account for some of their differences, like which ones license anaphoric discourse reference. Partee argued [1987] just along these lines in favor of a type-shifting view, which was then taken to explain why some NPs, like 'John' can bind a discourse pronoun like 'he' later on, in a different sentence, while other NPs like 'every man' apparently cannot. Whether this is in the end a good argument can be debated, but it is the right kind of argument against a uniform type assignment, since variable type assignments in these cases might explain differences in truth conditions and interpretability. But to insist that Montague is wrong, since he gives referential expressions a type other than e just misses the point of such type assignments.

Reference clearly contributes to the truth conditions, but as Montague has shown, we can capture the contribution to the truth conditions of a referential expression by giving it an assignment of higher type like the set of all the properties of the referent. We can move up in type assignment, even for a referential expression. And the other way round, we can move down to type e independently of whether a particular phrase that gets assigned this type is referential. This is in essence what happens with Chierchia and Turner's neo-Fregean proposal [1988], which introduces special objects into the domain of objects to serve as the semantic values of property nominalization. These special objects are systematically assigned to properties with Chierchia's NOM type-shifting operation, which is relied upon by S3 in their own type-shifting account. That is all perfectly fine semantic machinery, but simply because property nominalization gets assigned special objects of type e on Chierchia and Turner's proposal does not mean that property nominalizations in English are referential. These are simply two separate issues and one cannot infer from one to the other.

But even if no immediate inference from type of semantic value to being referential or not is strictly licensed, could it nonetheless be a reasonable inference

to make anyway? And would a semantic theory that assigned all and only referential expressions type e semantic values not be better, even if it is not strictly necessary? If one wants to claim such a connection between these two different things, then it needs to be argued for. One would need to show that these and only these expressions with those types of semantic values have this semantic function. No such argument is given by S3, nor do I know one in the literature. These two notions are usually simply conflated and moving from one to the other is usually simply done with no further justification. This would not be good enough.

To be sure, it would be good to know which expressions are referential, and if we can group all those together somehow, that would be very informative. But this does not have to be done by assigning them a particular kind of semantic value, in particular that of type e. It would be good to know which phrases are referential, but that is not the same issue as which phrases should get what type of semantic value. Those are simply different questions, and the variability of which types can capture the contributions to the truth conditions of various expressions with different semantic functions makes this clear. We simply need to keep these issues apart.

Or to put the point a different way: semantics, the discipline, aims to find out which expressions are referential, what the truth conditions of sentences are, and so on and so forth. Constructing a compositional semantic theory is part of semantics, the discipline, but not all of it. A compositional semantics aims to generate the truth conditions correctly by assigning semantic values of various types, with the goal of getting the truth conditions right. Semantics, the discipline, goes further in that it also hopes to find out which expressions have what semantic function. What type of value works in a compositional semantics is largely independent of what the semantic function of an expression is. There is no direct route from one to the other, and any claim of a direct connection needs to be argued for and faces just the obstacles I pointed to above. Semantic function and semantic values are two quite different aspects of semantic theorizing that should not be closely associated with each other.

But such a close connection between semantic value and semantic function plays a crucial role in S3's paper. Their own positive view, which they label the *Polymorphic Strategy*, is spelled out as: "... the same expression serves different referential and non-referential semantic functions" [2022, p. 72]. But most of their paper aims to argue that number words have a variety of semantic values on various occasions. Without a close connection between semantic values and semantic functions, their arguments do not support the position they aim to defend. Different semantic values do not guarantee different semantic functions, or so I have argued here. Thus I do not see much support for the view that number words have different semantic functions in natural language from the considerations given in S3's paper.

I am just as guilty as the next person for not having carefully distinguished these different aspects of the semantic evaluation of number words, in particular in [Hofweber, 2005]. That unclarity was particularly tied to the ambiguous use of the phrase 'semantic type'. As noted above, it can be understood as what

kind of an expression something is semantically, which in turn can be closely associated with semantic function. But it can also be understood as what kind of a semantic value an expression gets in a particular semantic theory. When I said in [2005, p. 205] that cognitive type coercion holds that number words are not of variable semantic type and that they appear syntactically contrary to their semantic type, then this can be understood in two ways. First, it could be seen as the view was intended, namely that number words retain their semantic function as determiners when they syntactically appear as singular terms. In particular, they do not acquire the function of reference in those cases. Second, it could be seen as holding that number words have a semantic value of a type different from their syntactic type. S3 understand me, not unreasonably, to endorse the second, but the one I intended to endorse is the first. Although I did not properly distinguish these things in 2005, I believe I did a lot better in [2016], where I much more explicitly distinguish semantic function and semantic value, and explicitly endorse that cognitive type-shifting is compatible with semantic type-shifting, in particular in [2016, pp. 137–141]. We will discuss this issue in more detail just below, when we will look at the difference between semantic and cognitive type-shifting.

To wrap things up about semantic values, we need to distinguish more generally between *semantic facts* and *semantic values*: semantic facts are facts about the semantic features of our languages. They include such things as what readings a sentence has, what the truth conditions of these readings are, what the semantic function of various expressions is, and so on. These are facts that simply obtain; they do not merely hold within a theory. Semantic values, on the other hand, are theory-internal things. They are assigned within a particular compositional semantic theory. What constraints do the semantic facts put on what semantic values should be assigned in particular semantic theories? The *minimalist view* holds that the only constraints on the correctness of semantic values is getting the truth conditions right: what readings a sentence has and how these readings are compositionally determined by the semantic values. These constraints on correctness can be met more or less elegantly and simply, and thus besides correctness conditions we can also consider elegance conditions. But the key for the minimalist is that there are no constraints on the correctness of the assignment of semantic values that go beyond the generation of the truth conditions. In particular, it is perfectly fine to assign expressions with a particular semantic function a certain type of semantic value, as long as this leads to the right truth conditions. I have defended this minimalist view in more detail in [2016, Ch. 8]. Whether minimalism is correct is debatable, of course, and it might go too far in only considering truth conditions as the constraints on correctness. But even if minimalism goes too far, we should certainly not equate semantic values with semantic functions, and furthermore we should not conclude from the fact that an expression gets a higher-type semantic value that it is not referential, nor that an expression with semantic value of type e is referential. The real and most significant aspect of Frege's other Puzzle is about semantic function: are number words referential, never,

sometimes, or always? That issue is neither settled nor properly addressed by looking at semantic values, be they always the same or different.

4. COGNITIVE VS SEMANTIC TYPE-SHIFTING

Everyone who accepts a type-theoretic approach to semantics should be open to accepting that the types associated with particular expressions can shift, both syntactically and semantically. This is much more innocent than S3 take it to be, since it merely means that what types get assigned to these expressions can be different on different occasions. In the semantic case this simply means that they get assigned different types of semantic value on different occasions, which should not be taken to mean or imply that they have different semantic functions on these occasions. Although S3 sometimes can be read as suggesting that accepting semantic type-shifting alone supports their preferred solution to Frege's other Puzzle, the real issue with Frege's other Puzzle is not whether type-shifting itself occurs with number words, but whether general, language-wide type-shifting principles alone can explain what is going on with number words in natural language and how this might affect the semantic function of number words. And here I disagree with S3's proposal.

My own preferred view, based on cognitive type-shifting, is not at all in tension with the occurrence of semantic type-shifting. However, it holds that general, language-wide type-shifting principles by themselves do not explain what needs to be explained, and we need to draw on cognitive type-shifting as well. I would like to stress here that cognitive type-shifting is a kind of type-shifting. It simply is one that happens for cognitive reasons, not reasons tied to purely general, language-wide type-shifting principles. In particular, I argued in [2005] that the reason why arithmetic equations like '$2 + 2 = 4$' have both a singular reading as 'Two and two is four' as well as a plural reading as 'Two and two are four' is tied to overcoming a cognitive difficulty we face when we learn arithmetic.[4] In a nutshell, the idea is that: (1) our minds are best at reasoning about objects; (2) reasoning about objects can be associated with operating on representations of a syntactically low type, assuming a picture of reasoning as symbol manipulation; and (3) reasoning about arithmetic as semantically bare determiner statements requires manipulating expressions of syntactically unsuitable high types. To overcome this difficulty, we coerce the bare determiner plural statement into one where number words occupy type e by systematically lowering the types of the determiners down to e, and correspondingly the types of the other expressions in this sentence. This type lowering then brings with it a shift from the plural to the singular, and helps us overcome our cognitive difficulties. That this happens can be seen, I suggested, by looking at when arithmetical questions are first introduced in the process of learning arithmetic and how the plural readings come temporally first, while the singular ones are dominant later.

[4] That both the singular and the plural statements are widely used can easily be verified by consulting corpora like Google's Ngram viewer at http://books.google.com/ngrams/

This cognitive type coercion is distinct for number determiners, and not also to be found with other determiners. So, even though it is perfectly fine to use bare determiner statements with other determiners, like 'some are more than none', it is not the case that there are singular versions of these statements like 'some is larger than none' or the like. The reason for this, I proposed, is that we face no distinct cognitive difficulty with reasoning directly with non-number bare determiner statements, since no complex reasoning of this kind occurs. But with number determiners things are different. There we do have complex arithmetical problems to deal with, at least complex for the mind of a four-year old. Even though we can express them with bare determiner statements like 'Twelve and seven minus four are fifteen', there are also corresponding singular statements available like 'Twelve and seven minus four is fifteen', or more symbolically '$12+7-4=15$'. This is distinctly so for number determiners, and the reason is cognitive type coercion, or so my proposal.

Crucially, cognitive type coercion is a type-shifting principle, lowering the type of a determiner down to type e. This is a change in syntactic type first and foremost, since that is the one relevant for reasoning, on the conception of reasoning as symbol manipulation. In particular, the type-shifting does not affect the semantic function of a number word: it remains a determiner, understood as an expression with the semantic function of a determiner, and it remains non-referential. S3 take the fact that the type-coerced determiner retains its semantic function and remains a determiner to mean that it gets a semantic value of the higher type of a determiner. And they then proceed to argue that such semantic values lead to mismatches with other expressions, as in their example (7a) on page 69,

(2) Four is a number,

as well as other examples. But this is missing the issue to an extent. That an expression retains its semantic function in a cognitive type shift does not mean that it retains the type of semantic value it receives in a compositional semantics. Cognitive type-shifting is a form of type-shifting after all. Type of semantic value and semantic function need to be distinguished, and retaining a semantic function while being cognitively type-shifted is neutral with respect to what type of semantic value the resulting phrase receives in a semantic theory.

When it comes to (2), there are several different aspects of what needs to be addressed: first, is this sentence grammatical, and if so, why? Second, what are its truth conditions? Only after this has been settled does the job of the compositional semantics come in to assign semantic values that generate these truth conditions. To do that the semantic theory needs to find semantic values that give the right result. This, I argued above, can be done in any way that works. What needs to be found is a way to generate these truth conditions, not a way to capture or mirror the semantic function of the phrases that occur in the sentence over and above how the contribute to the truth conditions.

My own view on this matter is that (2) is grammatical and true, even though 'four' is not referential in it. This might seem counterintuitive; does there not

have to be this thing which is the number four, which 'four' refers to and to which we attribute the property of being a number? I think that is a good first reaction, but nonetheless, it is mistaken. I tried to show in more detail in [2016, Chap. 8] that this reaction assumes what I called 'the referential picture of language', which maintains that the function of singular terms in true subject-predicate sentences is reference. But we have good reason outside of the issue about number words to hold that this picture is mistaken, or so I argued. Be all that as it may, the point remains that it is a mistake to equate semantic function with type of semantic value. Thus what type of semantic value 'four' gets in (2) does not settle what semantic function it has. My claim was and is that 'four' does not have the function of reference in (2), not that 'four' needs to get a semantic value of a particular type. If it works for the semantics, a value of type e is fine with me, but one should not confuse having value of type e with being referential and with referring to its semantic value. To be clear once more, I certainly deserve blame for not making these issues clearer in my [2005], but I did try to make them explicit in [2016], in particular on pages 137–142, which discuss the difference between cognitive and semantic type-shifting, and Chapter 8 which discusses the status of semantic values and their relation to semantic facts and semantic function.

5. IS SEMANTIC TYPE-SHIFTING ALONE ADEQUATE?

To solve Frege's other Puzzle with semantic type-shifting principles requires one to uncover general type-shifting principles that explain the various occurrences of number words in natural language. S3 aim to do just that. They specify a number of type-shifting principles that are supposed to generate all the relevant occurrences of number words. To carry out this project it is important to have two things in mind. First, the type-shifting principles have to be general, not just specific to number words. That is to say, there have to be general principles that relate words from one type to another one. It is not enough simply to formulate some principles along the lines of: number words can occur as NPs and as determiners, and when they are NPs they do this, and when they are determiners they do that. Instead, one needs general principles that relate, say, NPs to determiners, and say how their semantic contributions are related on these different occurrences. Second, it is not enough simply to say how number words may occur in different ways; one also needs to specify how they may not occur. So, it is not enough simply to have permissive principles, ones that say what occurrences are permitted, but also to have restrictive principles that say what occurrences are forbidden. At least one needs to do this unless one accepts that all permitted occurrences are indeed fine and meaningful. This latter restriction applies to both straightforward ungrammaticalities as well as to what readings of a sentence are available and what readings are preferred. A perfectly grammatical sentence can have a number of readings, and sometimes these readings can be attributed to the different types that an expression can have. In the end, all of this needs to come together.

As S3 note, I have in the past suggested that semantic type-shifting is the second most promising attempt for solving Frege's other Puzzle, but it faces

obstacles that I find difficult to overcome. One of these obstacles is that the relied-upon type-shifting principles are not general, and the second is that it makes the wrong predictions about cases. Because of these difficulties I proposed cognitive type-shifting, which is specific to number words, and which helps to make the right predictions. S3 discuss my argument about predictions in their paper, but I am not persuaded by their response to it, and I would like to discuss this next. After that we will consider S3's own type-shifting proposal and its problems with generality.

My original argument against semantic type-shifting was this: if the different occurrences of number words in natural language were simply the result of general type-shifting principles, then they should obey the general constraints of such type-shifting principles. In particular, when numerous readings are available which result from several different possible type assignments, then the one that assigns the lowest type should be the preferred reading. But this is not so, as examples like

(3) I want two or three beers.

make clear. That sentence has a reading where 'or' combines determiners (which are of higher type) or NPs (which are of lower type), but the preferred, and close to only, reading is the one with higher types, contrary to general principles that the type-shifting approach relies upon. That speaks against a semantic type-shifting account, or so I argued in [2016, pp. 139 ff.], in particular against there being a general type-shifting principle that moves a determiner to a noun phrase.

S3 respond in two different ways [2022, pp. 81–83]: first, they argue that taking lowest-type readings to be preferable is an optional addition to the theory, and second, that the argument confuses two kind of priority: priority in processing and priority in interpretation. I remain unpersuaded by both of these responses. First, a preference for lower type readings is not just an optional part of the theory, but a crucial aspect of it. This is quite explicitly endorsed by the proponents of this theoretical approach, cited just below, and it does play an important role in it. The problem is simply that general type-shifting principles give rise to numerous readings, some quite outlandish. But there is no such plurality to be found in the data of our semantic theorizing, namely what readings sentences are taken to have by speakers. We must account for this in our theories, including the fact that some readings are vastly preferred over others. The standard way of dealing with this problem is to give preference to the lowest-type reading, and to cut off interpretation if the lowest-type reading is reasonable in context. The lowest-type reading does not have to be the only one, but it is the preferred one: the one you would expect, unless there are overriding reasons against this reading in the context. To illustrate the problem, suppose I am thinking about which animals to adopt at the shelter, and I say

(4) I'm considering Fido or two cats.

This sentence has numerous readings according to various type-shifting principles. For example, according to S3, proper names can be type-raised to be modifiers of nouns, as in 'Einstein idea' in 'I had one of those Einstein ideas the other day'. On such a proposal a sentence like (4) semantically has a reading where 'Fido' either is a bare modifier, picking up what it modifies from an earlier utterance or context (*i.e.*, a Fido animal), or it combines with a properly type-adjusted 'two' to form a complex modifier of 'cats' (*i.e.*, Fido cats or two cats). But none of this is to be found as a reasonable interpretation of (4), and it needs to be explained why not. One straightforward explanation is due to a principle that says that in interpretation the lowest types are tried first, and if that results in a reasonable reading in the context, then the issue is done and dusted. In that case no other in-principle possible type assignment is taken as a reasonable reading on an ordinary occasion of utterance. That is an explanatory proposal, one endorsed by the proponents and originators of this general approach to semantics.[5] Barbara Partee, one of the founders of this general approach, states this explicitly, for example in [1987, p. 359], where she says "... there is a general processing strategy of trying lowest types first, using higher types only when they are required in order to combine meanings by available composition rules." This is the third of three general principles of the polymorphic-type approach, according to Partee [1987] at least. S3 take it to be optional, but I do not see what alternative proposal would explain what needs to be explained here.

Secondly, S3 take relying on lowest types to generate preferred readings to confuse two senses of priority: priority in processing and priority in interpretation. They suggest that these senses of priority should not be conflated. Here they are correct that these are two different senses of priority conceptually, and so they should conceptually be distinguished. But even though they are conceptually different, they are closely related. Priority in processing can help explain priority in interpretation, although only defeasibly. It can be that in a particular context the reading that is processed first is unsuitable as the proper one for contextual reasons and then we move on to look for another one. So, the two senses of priority can certainly come apart. But nonetheless, there often is a preferred reading and the question is what explains that this is the preferred among several readings. One answer, favored by Partee and also relied upon for my argument, is that the priority in processing defeasibly explains the priority in interpretation. The priority in processing is tied to lowest feasible type assignment. And within this framework (3) indeed is a problem for the pure type-shifting approach to Frege's other Puzzle like the one favored by S3.

I should add that this problem is distinct to a pure type-shifting approach. On a cognitive type-shifting approach number words can be of type e, but there is no claim made that type e has interpretive priority. The reason that number words appear in type e, on the cognitive type-shifting approach, is to simplify reasoning in the case involving complex bare determiners. So, we

[5] See [Partee and Rooth, 1983; Partee, 1987].

lower the type of the plural 'Two and two are four', which involves three bare determiners, to that of lower type, and with it the singular 'two and two is four' to ease cognition. But this type-lowering is specifically for cognitive reasons, not for purely semantic reasons. It is not the result of general semantic type-shifting principles, and with it it does not require a priority of the lowest-type reading in interpretation. The lowest-type reading will still have priority among all the readings licensed by purely semantic type-shifting principles, but that does not include the reading of (3) where 'two' is of type e. The source of type e readings of number words is to help with cognition and reasoning in arithmetic, not the interpretation of ordinary assertions. The cognitive type-shifting proposal might have other problems, but the standard reading of (3) is not one of them, contrary to the purely semantic type-shifting proposal, which predicts the wrong preferred reading of this sentence.

But that is not the only problem with the pure type-shifting approach. The second problem is tied to whether the approach is based on universal type-shifting principles, ones that connect expressions of general semantic categories to each other, or whether it is tied to specific type-shifting principles, ones that only work for number words. Relying on principles that apply distinctly and exclusively to number words is not enough. It is clear how the different occurrences of number words are broadly related to one another when it comes to truth conditions: everyone in this debate can agree that the number of moons is four iff there are four moons, and so on. And everyone can agree that there are different kinds of semantic values that could be assigned to these different occurrences of number words to generate these truth conditions. The problem is to explain why number words have these different features, and in particular how they do this in light of what their semantic function is, whatever it may be in the end. This could be explained if there were general type-shifting principles that gave the right results. That is to say: there are general, language-wide type-shifting principles which generate those sentences with these truth conditions. For these principles to be explanatory, they need to be general in the sense of not merely applying to number words, but to expressions of the same semantic kind more broadly. S3's approach has the ambition to do just that. They motivate their approach by considering other examples of expressions, like adjectives, and their variable types. But despite this ambition, some of the type-shifting principles they propose are specific for number words only. The most obvious example here is their principle NUM, which maps type e to type $\langle e, t\rangle$, but not in general. NUM maps a number n to the property of being a plurality with n elements. But this is a type-shifting principle specific to numbers only. It does not in general map type e to type $\langle e, t\rangle$, and, as S3 also note (fn. 23), it is not a generally accepted type-shifting principle in the semantics literature. This is particularly relevant for polymorphic substantivalism, in S3's terminology, which takes number words lexically to be referring expressions of type e, which derivatively achieve their predicative meaning via NUM. This would distinguish number words from other lexically referring expressions like presumably proper names, who do not come with a special type-raising principle that gives them a unique predicative meaning. Here it is important to note that S3's NUM is

different from a universal type-raising principle, which does apply to all type e expressions and raises j to $\lambda x.x = j$, *i.e.*, John to being John. That applies to all type e expressions. But NUM is specific to numbers. It is like a type-raising principle that is specific to people and raises John to the property of being one of John's siblings. That type-raising principle is not one that comes from general semantic type-shifting principles, and neither does NUM.

But the issue is not with NUM alone. Chierchia's NOM type-shifting operation, which connects a property to a postulated object of type e which is the semantic value of the corresponding property nominalization, is a general one, it is assumed to apply to all properties, with some possible exceptions like properties leading to paradoxes. But it is not restricted to properties associated with numbers, like the one that results from applying NUM to a particular number of type e. NOM applies not just to properties like $\lambda x.\mu(x) = 4$, but also to properties like $\lambda x.\mu(x) \geq 1$, *i.e.*, not just to properties like being a plurality of size 4, but also to being a plurality with size larger than or equal to 1, which is to say, being some things. Applying NOM to the former leads to the nominalization of that property, *i.e.*, the property of being a plurality of size 4, which, according to S3, is identified as the number 4 via Rothstein's identification of numbers with the denotations of these property nominalizations [2017]: $n = {}^{\cap}[\lambda x.\mu(x) = 4]$, which S3 label (RSE). But since being some things can also be nominalized, it has a denotation of type e as well, which could be identified in a similar way with a type e reading of 'some': $some = {}^{\cap}[\lambda x.\mu(x) \geq 1]$. But that is absurd. There is no type e reading of 'some', and there is no singular NP use of 'some'. But why not? S3's account does not explain this, since their proposal uses principles specific to number words, not general principles that explain these differences. NUM is specific to number words, and RSE is used only to identify the denotation of number words with the denotations of certain nominalizations of property terms, but their natural generalization to other expressions does not seem to hold.[6]

None of this is supposed to be taken to show that the type-shifting account proposed by S3 is mistaken, only that it is not general enough and not explanatory. Some of their type-shifting principles are specific to numbers, and their list of type-shifting principles is surely only a start. For example, their type-raising principles should be augmented with corresponding type-lowering principles, as is standard in the type-shifting literature. And they seem to need further principles, since their principles so far do not reach all the required higher types; in particular that of a determiner is not reached with them as formulated. To be sure, that their proposal is incomplete does not speak much against it in the present state of the debate, since all alternative proposals, including, of course, also my own, are equally incomplete. No one in this debate can account for all the data, as far as I can tell, and that should be no surprise. The issue is complex and difficult, not because of problems tied to semantics in general, but also because of the large variety of uses that number words have in natural

[6] Similar remarks also apply to the more detailed proposal made in [Snyder, 2021].

language. But there are clear divisions in what strategies and frameworks different parties in this debate prefer to rely upon in their attempts to come up with an adequate account. And here there are real differences in overall vision and general outlook.

The lack of generality of purely semantic type-shifting accounts is in part what motivates the cognitive type-shifting approach. On that approach we acknowledge that there is something distinctly different about number words in natural language, something that is not to be found in general with expressions of the same category as number words, be they adjectives, determiners, or modifiers. The reason for this distinctness, and with it the explanation for why number words have this distinct feature, is something cognitive. There is something different about our thought about numbers and how it is represented in language which does not carry over to other types of thoughts that involve expressions of the same kind as number words. It is the complexity of reasoning with number determiners that is distinct for number determiners, and not to be found in other determiners. That is why number words have features in natural language that are not to be found in general, and that is, in part, why no general type-shifting principles explain everything that needs to be explained about number words. For number words we need something else in addition: specific cognitive type-shifting principles.

6. CONCLUSION

Obviously, I could not address all the arguments and points raised in S3's rich paper, but I hope to have made the main points of disagreement explicit and clarified my own position. We disagree, in particular, about a close association of semantic value with semantic function, and whether semantic type-shifting principles by themselves can explain what needs to be explained about number words. We do not disagree about whether semantic type-shifting occurs at all, including with number words, only about whether semantic type-shifting principles alone explain the features of number words that we find in natural language.

Although reasonable people can disagree about what the most promising line of explaining what needs to be explained is, I hope we can all agree on the significance of this general issue for the philosophy of arithmetic. Understanding number words in natural language is a promising route towards understanding arithmetical statements in mathematics, in particular with respect to what the basic semantic function of the number symbols in these arithmetical statements is. Although a good part of the debate in the philosophy of mathematics is concerned with much more advanced mathematics, it is notable, and maybe even slightly scandalous, that we still do not have a proper agreed-upon understanding of even simple arithmetical equations and the function of number symbols in them.

ACKNOWLEDGEMENTS

My thanks to two anonymous referees for their thoughtful comments on an earlier version, which led to numerous improvements.

REFERENCES

Barwise, Jon, and Robin Cooper [1981]: 'Generalized quantifiers and natural language', *Linguistics and Philosophy* **4**, 159–219. DOI: https://doi.org/10.1007/bf00350139.

Chierchia, Gennaro, and Raymond Turner [1988]: 'Semantics and property theory', *Linguistics and Philosophy* **11**, 261–302. DOI: https://doi.org/10.1007/bf00632905.

Frege, Gottlob [1884]: *Die Grundlagen der Arithmetik: eine logisch-mathematische Untersuchung über den Begriff der Zahl*. Breslau: W. Koebner.

Hofweber, Thomas [2005]: 'Number determiners, numbers, and arithmetic', *The Philosophical Review* **114**, 179–225. DOI: https://doi.org/10.1215/00318108-114-2-179.

_____ [2007]: 'Innocent statements and their metaphysically loaded counterparts', *Philosophers' Imprint* **7**, 1–33. Available at https://quod.lib.umich.edu/p/phimp/3521354.0007.001.

_____ [2016]: *Ontology and the Ambitions of Metaphysics*. Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780198769835.001.0001.

Montague, Richard [1974]: 'The proper treatment of quantification in ordinary English', in P. Portner and B.H. Partee, eds, *Formal Semantics: The Essential Readings*, pp. 17–34. Yale University Press, 2002. DOI: https://doi.org/10.1002/9780470758335.ch1.

Partee, Barbara [1987]: 'Noun phrase interpretation and type-shifting principles', in J. Groenendijk, D. de Jongh, and M. Stokhof, eds, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, pp. 115–143. Berlin, New York: Foris Publications. DOI: https://doi.org/10.1515/9783112420027-006.

Partee, Barbara, and Mats Rooth [1983]: 'Generalized conjunction and type ambiguity', in Rainer Bäuerle, Christoph Schwarze, and Armin von Stechow, eds, *Meaning, Use and the Interpretation of Language*, pp. 361–393. Berlin: Walter de Gruyter. DOI: https://doi.org/10.1515/9783110852820.361.

Rothstein, Susan [2017]: *Semantics for Counting and Measuring*. Cambridge University Press. DOI: https://doi.org/10.1017/9780511734830.

Snyder, Eric [2021]: *Semantics and the Ontology of Number*. Cambridge University Press. DOI: https://doi.org/10.1017/9781108602259.

Snyder, Eric, Richard Samuels, and Stewart Shapiro [2022]: 'Resolving Frege's other Puzzle', *Philosophia Mathematica* (3) **30**, 59–87. DOI: https://doi.org/10.1093/philmat/nkab032.