What can we mean?
on practices, norms and pluralisms

Greg Restall
Arché, Philosophy Department,
University of St Andrews[1]

Draft paper for the *Aristotelian Society*, October 2024

* * *

In his presentation to this Society 65 years ago, Michael Dummett (1959) took his initial steps on what became a 50-year exploration of the connections between truth, logic, the foundations of semantics and fundamental issues in metaphysics. Dummett's work on realism and anti-realism was at the centre of fervid discussion and active debate from the 1970s to the 1990s. At the centre of Dummett's programme was revisionary view of fundamental principles in logic (Dummett 1977, 1991).

According to Dummett, many of the principles of logic can be shown to be self-justifying on neutral semantic grounds, without prejudging the debate between the realist and the anti-realist about any domain.[2] However, not all traditional logical principles pass muster. In

---

[1] Heartfelt thanks to my *Arché* colleagues, including Franz Berto, Aaron Cotnoir, Viviane Fairbank, Sandy Goldberg, Sophie Nagler, Sabina Domínguez Parrado, Stephen Read and Crispin Wright, for helpful comments, questions and discussion on earlier presentations of this material, and many of the ideas discussed here.

[2] We do not need to go into the details here, but the general idea goes as follows: For a given logical concept such as conjunction, the conditional or the quantifiers, we can treat the rule *introducing* a judgement as its definition. (For the conditional, the introduction rule says that we can infer $A \to B$ when we can make a deduction from the supposition of A to the conclusion B). Its corresponding *elimination* rule should be in harmony with the introduction rule, allowing us to infer *from* the judgement only what we could use to deduce it in the first place (so, with a conditional $A \to B$ you may utilise this in deduction by deducing B when you have already deduced A). See Dummett's *Logical Basis of Metaphysics*, Chapter 11 (1991), for more details on the approach. There has been much discussion on whether Dummett's criterion can be defended and *exactly* how it is to be understood (Steinberger 2011). Dummett's argument to the conclusion that the distinctively *classical* laws are not so justified depends on an account of the general rules governing the proofs in which the rules are given. Different rules governing the *assumption contexts* in which a proof may be constructed give rise to different logical systems as self-justifying, including classical logic, or even a relevant logic, or a number of other non-classical logics (Restall 2023), but the details of *that* argument are beside the point here.

particular, the inference of double negation elimination (the step from $\neg\neg p$ to p) is not as harmless as other principles. To adopt it involves a metaphysical commitment that skews debate in favour of the realist—or so the story goes. The philosophically neutral logical perspective, acceptable to all sides, would be to accept only intuitionistic logic, whose constructive reasoning principles can be justified on grounds independent of the debate over realism or anti-realism concerning any subject matter.

The debate over whether Dummett is right, and whether intuitionistic logic is the neutral foundation for semantic and metaphysical theory reached its peak in the 1980s and early 1990s (Taylor 1987, Wright 1992. See Green 2002, for a retrospective view). Fashions change in philosophical logic as much as anywhere else, and this disagreement has, by-and-large, receded from view in contemporary philosophical logic. Since the heyday of the debate between realism and anti-realism, philosophical logic turned its attention to many different issues, including vagueness, contextualism and assessment sensitivity, paradox and paraconsistency, pluralism and monism, whether logic is in any way *exceptional* among the sciences, and much more besides. Even when we narrow our focus to philosophical logic in *Oxford*, which was ground zero for the debate over Dummettian anti-realism, the current scene could not be more different. Oxford philosophical logic is now dominated by discussions of higher-order modal logic, and the whole host of metaphysical concerns that arise (Williamson 2013, Bacon 2023, Fritz and Jones 2024). Almost everywhere, Dummettian concerns are sidestepped, rather than addressed head on.[3]

A rather curious fact is that the situation is almost exactly reversed in the world of *mathematics*. In Dummett's own day, the clarion call to adopt *intuitionistic logic*—with its revisionary treatment of the law of the excluded middle, double negation elimination, and other classical inference principles—fell largely on deaf ears. The tradition of *constructive mathematics*, in which mathematics is refounded on properly intuitionistic principles, was, in the second half of the 20th Century very much in the minority, despite Dummett's attempts to spread the intuitionistic good news (Bishop and Bridges 1985, Bridges et al. 2023).

Things look different in mathematics today, as intuitionistic logic and the distinctive forms of reasoning involved in doing mathematics constructively has gained a significant new foothold in mathematical practice, not directly through the arguments or the example of philosophical logicians like Dummett, but with the rise of *proof assistants* (Avigad 2024). As mathematicians rush to formalise mathematical results in the language of proof assistants like *Agda* (Bove et

[3] See Williamson's "Must Do Better" (2006) for a vivid and opinionated account of *why* this debate was sidestepped. I say *almost* everywhere, because there are exceptions. Dummettian concerns are still in view in recent work on inferentialism and proof-theoretic semantics (Incurvati and Schlöder 2023, Peregrin 2014, Tennant 2017, Rumfitt 2015), though even there, the debate between realism and anti-realism and the issue of the neutrality of intuitionistic logic, has receded from centre stage.

al. 2009) and *Lean* (Avigad et al. 2023), the most fundamental and basic logical principles in formalised mathematics turn out to be those principles Dummett advocated as metaphysically neutral. Popular proof assistants such as Agda and Lean are, at their core, Dummettian constructivists, reasoning intuitionisticially, while being prepared to countenance additional classical principles as an additional commitment you might make when needed (see, e.g., Avigad et al. 2023, Section 3.5). So, the advent of wide-scale use of constructive reasoning methods in formalised mathematics is a prompt to revisit these decades-old discussions in philosophical logic, with the new insight we can gain from this unexpected development.

The questions raised by this development are of concern to more than just specialists in philosophy of mathematics and logic—they have broader philosophical interest. The use of proof assistants in mathematics is one example among many of the use of machines as a part of our cognitive activities. *Semantics* is not just a matter of ascribing meanings to our own thoughts and words—semantics matters at the human/machine interface too. How should we understand the role that computational systems play in our own practices of explanation, inference, and justification? There is scope to revisit Dummett's concerns about meaning, truth and reality, broadening our field of view to take in the role that computational devices play in our thought and talk and action.

The upshot of all this will not be an argument for anti-realism over realism (or *vice versa*), or for one logic over another. As a logical pluralist (Beall and Restall 2006), I am inclined to try to better understand the kinds of practices we might engage in, and what we can *do* by taking part in those practices, without attempting to isolate one practice somehow *correct* while the others are incorrect. How do the different things that we *do* when we think and talk and reason and argue—and that our *devices* do when we co-opt them in that process—manage to *mean* things, and make claims on the world and on each other?

* * *

Computers have changed the face of mathematical research in many different ways, from the revolution in publishing provided by the internet and preprint servers, the use of calculators and computers to eliminate routine and mundane arithmetic calculation, to the possibilities afforded by high-performance computation in running simulations of large-scale dynamical systems — the practice of mathematics has changed significantly from the days of Euclid, or of Euler, or even since the post-war growth of mathematical research in the middle of the 20th Century. The rise of *proof assistants* is another significant change in mathematical practice.

A proof assistant is not primarily an artificial theorem *prover*—it does not construct mathematical proofs for you—it acts as a patient research assistant, checking your work, making sure that your statements are consistent with your definitions, and checking that the proof that *you* write out is correct (Avigad et al. 2023). There is nothing in the abstract notion of a proof assistant that would favour one way of understanding valid proof over another. A formalist about proof could specify some arbitrary system of "proof rules", and a proof assistant programmed to check those rules could check that any given string of symbols supplied to it satisfies whatever the constraints the "proof" checker was designed to check. Contemporary proof assistants are *not* formalists in this sense. Proof assistants such as *Agda* (Bove et al. 2009) and *Lean* (Avigad et. al. 2023) are more opinionated about what kind of thing a proof *is*, and this is why it turns out that the logic of proof assistants is at heart, constructive.

Proof assistants like *Agda* and *Lean* treat proofs as *functions*, where a valid argument from some premises to a conclusion is represented as a *function* that transforms grounds for the premises into grounds for the conclusion. The ground for a conjunction of two propositions consists of a *pair*, the first element of which is a ground for the first conjunct, the second, a ground for the second conjunct. The ground for a conditional statement ($A \to B$) is a function that transforms grounds from the antecedent A to grounds for the consequent B. The ground for a universally quantified judgement of the form "all Fs are Gs" is also a function, which, given an object (of the relevant type) and a ground for the claim that this object has type F, is able to provide a ground for the claim that the object has type G. We should define negation, since it is important when distinguishing constructive reasoning from classical logic: The negation of A is understood as $A \to \bot$ where $\bot$ is a given contradictory proposition that *never* has a ground. So, a ground for the negation of A is a function which can provide, for any ground given for A, *per impossibile*, a ground for $\bot$. So, we can ground the negation of A only when A can have no ground.

So much sounds relatively straightforward, at least for anyone familiar enough with logical vocabulary. Straightforward enough, that is, except for the unspecified notion of *ground*. What *is* a ground, in general? The philosopher is interested in this question (Prawitz 2012), but the designer of proof assistants need not worry about the metaphysics or the epistemology of grounds. As far as the *logic* goes, proofs are functions that combine grounds and supply new grounds from old in regular ways. When it comes to *mathematics*, the definitions of the basic concepts will tell us what we need to know, structurally, about the grounds of atomic judgements. In fact, in *type* theory, propositions are just a special instance of the more general class of *types*, and proofs are a special instance of *terms* inhabiting those types. Constructive type theory is a general account of types and terms, inside which proofs and propositions. A proof π from A to B and a function f from $\mathbb{R}$ to $\mathbb{N}$ are *exactly* the same kinds of thing (Martin-Löf 1985).

The reasoning principles arising naturally in type theory are familiar from intuitionistic logic (Dummett 1977, Heyting 1956, Rathjen 2023), and proof assistants like *Agda* and *Lean* are implementations of Martin-Löf's dependent type theory (1984). Since this framework takes the construction of proofs to be a specific case of constructing *functions*, functional programming languages provide a useful framework for developing of proof assistants. So, as mathematicians labour to encode their definitions and proofs in the vocabulary of proof assistants like Agda and Lean, they are learning to express their results in the language of dependent type theory (Escardó and collaborators 2024, Lean community 2024).

Mathematics encoded in this way is, at its core, *constructive*. A proof of a disjunction $A \vee B$ (from an appropriately determinate background context) may be transformed into a proof of one of the disjuncts, A, or B. A proof of an existentially quantified statement $\exists x\ \phi(x)$ may be transformed into an algorithm supplying a witness term t where we can prove $\phi(t)$. Such results are *impossible* in classical logic, since $p \vee \neg p$ is a classical tautology, but we cannot expect to prove an arbitrary p or $\neg p$.[4] Classical mathematical theories can tell us that f is a continuous function where $f(0) < 0$ and $f(1) > 0$, and so, that there is some number r between 0 and 1 where $f(r) = 0$ (this is the *intermediate value theorem*), but we may be in no position to *find* such a number r.

Mathematicians nonetheless regularly make use of classically valid principles, and proof assistants allow for this, by allowing for the development of proofs where classicality is an added *assumption* (Avigad et al. 2023, Section 3.5). The situation is strikingly similar to Dummettian semantic anti-realism where constructive reasoning principles are the neutral agreed-upon core, and distinctively classical principles are an optional extra, to be adopted when the metaphysics—or the mathematical theory—asks for it.

However, the addition of classical principles is *optional*, and distinctively constructive mathematics is possible, in which classical assumptions are avoided.[5] The resulting mathematical results do not only have the distinctive *computational* properties mentioned above, they also apply to a wide range of mathematical structures, which are of independent interest whether you start out as committed to intuitionistic logic or not.[6]

This well established, if still minority, practice of constructive mathematical theorising raises a significant question. How are we to understand the relation between constructive mathematics

---

[4] $\neg\neg(p \vee \neg p)$, on the other hand, *is* provable. It is straightforward to refute $\neg(p \vee \neg p)$ (since this entails both $\neg p$ and $\neg\neg p$), an obvious contradiction. So, in an important sense (discussed further below), $p \vee \neg p$ is constructively *undeniable*.

[5] The online repository *TypeTypology* of Martín Escardó and collaborators (2024) is a good example of the kind of depth and breadth of distinctively constructive mathematical results.

[6] The "internal logic" of cartesian closed categories is intuitionistic (Lambek and Scott 1986). (This is another way to understand the functional interpretation proofs and types mentioned above: a conjunction $A \wedge B$ is understood as the cartesian product $A \times B$ and the conditional $A \to B$ is the function space from A to B.) Different kinds of cartesian closed categories provide natural examples of "spaces" which are governed by an intuitionistic logic (e.g. Hyland 1982). So, constructively proved mathematical results apply in a range of different "mathematical universes." Andrej Bauer (2016) gives an account of what it is like to learn to do mathematics constructively by way of attending to the different spaces in which these results can apply.

and classical mathematics? A relatively standard account of the difference is illustrated in Bishop and Bridges' 1987 monograph on constructive analysis:

> …take the assertion that every bounded non-void set *A* of real numbers has a least upper bound. (The real number b is the *least upper bound* of *A* if $a \leq b$ for all a in *A*, and if there exist elements of *A* that are arbitrarily close to b.) … If this assertion were constructively valid, we could compute b, in the sense of computing a rational number approximating b to within any desired accuracy… (Bishop and Bridges 1987, p. 7)

Here, the thought is that we can prove *less* when we reason constructively than when we reason classically. Something that we might be able to classically prove we might *not* be able to constructively prove. Constructive mathematics is a *restriction* on classical mathematics.

However, we need not think of constructive mathematics as a *restriction*. Consider this, for contrast:

> …constructive logic is stronger (more expressive) that classical logic, because it can express more distinctions (namely, between affirmation and irrefutability), and because it is consistent with classical logic. Proofs in constructive logic have computational content: they can be executed as programs, and their behaviour is described by their type. Proofs in classical logic also have computational content, but in a weaker sense than in constructive logic. Rather than positively affirm a proposition, a proof in classical logic is a computation that cannot be refuted. (Harper 2016, p. 104)

From this contrary standpoint, constructive mathematics is an *expansion* of classical mathematics, because more distinctions can be drawn, and the constructive mathematician has more expressive power. For the classical reasoner, $\neg\neg(p \vee \neg p)$ and $p \vee \neg p$ say the same thing, while the constructivist takes them to have different content.

What *should* we say? Is constructive practice a *restriction*, or an *expansion* of classical reasoning? In the remainder of this paper, I will attempt to clarify what is at stake in either of these perspectives, by paying attention to proof assistants, and what we *do* with them.

* * *

I will start with an analogy. Consider the humble calculator—a device that plays an essential role not only in giving answers to arithmetical questions, but in giving us *knowledge* that we would not otherwise have: when a calculator says that $345 \times 678 = 233{,}910$, we thereby *learn* that 345 times 678 is 233,910. How does the calculator perform this task? A calculator is more than an abacus or a pencil and paper, which serve as an extension of our memory when we are doing our sums. Calculators are *devices* that we use to *do* calculation for us.

We acquire our knowledge of basic facts of arithmetic by way of an education involving *counting* things, adding collections together, and so, learning about addition of numbers, generalising to multiplication, while at the same time, learning some kind of notational system for numbers, and methods for doing elementary arithmetic exploiting features of those representational systems. The details of different people's training will differ, but at root, it is difficult to see how someone could learn *arithmetic* without knowing how to count things.

Calculators do not count things. They manipulate patterns—states of the computational system—in ways that *we* recognise as representing numbers, in regular ways. It is enough for our purposes for it to serve as a reliable intermediary and a tool in our counting and calculating practice. It does not need to be able to count five things, any more than an abacus does. But it does need to be doing things such that the regularities observed in the action of the calculator can be *read* by a user as a part of a mathematical explanation. That much seems necessary if we want the calculator to play a role in a demonstration that $345 \times 678 = 233{,}910$.

Exactly *what* regularities are required for the actions of a calculator to count as reliably *doing* arithmetic? The simple answer is that it needs to get arithmetic *right*. That is fair enough, but since there are infinitely many different arithmetic statements, put this way, it is an infinite set of requirements, and one that is, if we leave it in this form, not feasible either to *impose* in the first place as we build a machine, or to *check* for compliance, once a machine is built.[7] We can check this infinitely large (or even, just a stupendously large finite) collection of constraints by verifying that the calculator's output matches the content of some *formal theory*, which can be finitely specified.

What theory might we use? Here, there is more than one candidate, because arithmetic (and our counting practices) can be made rigorous in more than one way. A familiar formalisation of arithmetic is in the axioms of *Peano Arithmetic*. Here, there are three axioms governing the notion of *zero* and the *successor* function, which supplies, for each number x its *successor* sx.[8]

- $sx \neq 0$
- $sx = sy \rightarrow x = y$
- $x \neq 0 \rightarrow \exists y\, x = sy$

Then, familiar arithmetic functions on the natural numbers, like addition and multiplication, can be defined recursively in familiar ways

[7] Traditional pocket calculators can represent numbers only up to some finite bound. However, more sophisticated calculating devices can work with natural numbers of arbitrary size, limited only by available time and memory, where the available memory of the device can be expanded as needed.

[8] Here, as always, any unbound variables are implicitly universally quantified. $sx \neq 0$ can be understood as $\forall x\, sx \neq 0$; $sx = sy \rightarrow x = y$ as $\forall x \forall y(sx = sy \rightarrow x = y)$, and so on.

- $x + 0 = x$
- $x + sy = s(x + y)$
- $x \times 0 = 0$
- $x \times sy = (x \times y) + x$

We represent the idea that the natural numbers are *only* those numbers found by starting with zero and taking successive applications of the successor function by adding the principle of induction, according to which if zero has a feature ϕ and whenever a number has ϕ so does its successor, then *all* numbers have that feature.

- $[\phi(0) \wedge \forall x(\phi(x) \rightarrow \phi(sx))] \rightarrow \forall x\phi(x)$

If the output of our calculator *agrees with* the judgements of Peano Arithmetic, it is reliably doing finite arithmetic, no matter *how* it does this. Such a calculator could be understood as *calculating*. The theory of Peano Arithmetic is a well-understood regularisation and formalisation of our arithmetical practice, even though it does not involve any notion of *counting* or *enumerating*.

Other formalisations of arithmetic *do* make some kind of use of a notion of counting. Neo-Fregean formalisations of arithmetic introduce arithmetical concepts by way of a notion of *predicate abstraction*. For any one-place predicate F, we have a singular term $\sharp F$, to be read as "the number of F's", and the key principal governing this term-forming operator is *Hume's Principle* (Wright 1983),

- $\sharp F = \sharp G \leftrightarrow \exists f(f{:}\, F \leftrightarrow G)$

which, using the resources of second-order logic, states that the number of Fs is the number of Gs if and only if there is a bijection between the Fs and the Gs. Neo-Fregean arithmetic more explicitly corresponds with the general conception of numbers as involving *counting* things. With the help of lambda abstraction,[9] we can introduce the finite numbers using the notion of identity:

- $0 =_{df} \sharp\lambda x\, x \neq x$
- $1 =_{df} \sharp\lambda x\, x = 0$
- $2 =_{df} \sharp\lambda x(x = 0 \vee x = 1)$
- $3 =_{df} \sharp\lambda x(x = 0 \vee x = 1 \vee x = 2)$, etc.

[9] If $\phi(x)$ is a formula in which the variable x may occur free, then $\lambda x\, \phi(x)$ is a one-place predicate, where for any singular term t (that is free for x in $\phi(x)$), $\lambda x\, \phi(x)$ holds of t if and only if $\phi(t)$. So, $\lambda x\, x \neq x$ is a 'non-identity predicate' which holds of t if and only if $t \neq t$, i.e., it holds, *never*.

We can define addition by first settling $\sharp F + \sharp G$ to be $\sharp\lambda x(Fx \vee Gx)$ when nothing is both F and G (corresponding to the naïve idea of addition as counting up two disjoint collections), and continuing from there. If our calculator's output agreed with a neo-Fregean theory, it would *also* count as recognisably doing arithmetic.

There are more formal theories for arithmetic than just these: any formal theory is an account of possible *patterns* in which our many and varied practices can run, and these can be more or less complex.[10]

A calculator might, under-the-hood, implement a neo-Fregean arithmetic, or a Peano Arithmetic, or be doing something else besides. What is required for it to be intelligible as *doing arithmetic* is that there is some translation between what *it* is doing with some recognisable arithmetic practice. And the same holds for *you* and for *me* and for anyone else who uses arithmetic vocabulary.

Now, these counting practices agree on a great deal, but disagree at the margins. Ask yourself (or ask someone with competence in elementary school arithmetic) this: Is there a number n where $n = n + 1$? The answer is *no* for someone whose concept of arithmetic complies with the conditions of Peano Arithmetic,[11] where, and the answer is *yes* in a neo-Fregean arithmetic, since the number $\sharp\mathbb{N}$ of finite natural numbers satisfies $\sharp\mathbb{N} = \sharp\mathbb{N} + 1$, since we can put the natural numbers in bijection with the natural numbers plus one extra thing (recall *Hilbert's Hotel*).

It would not be a surprise for a competent user of arithmetic vocabulary to find that their own concept of number simply does not settle the issue as to whether a number can be its own successor. On some precisifications of the number-concept (finite ordinal numbers, as modelled in Peano Arithmetic), no number is its own successor. On others (cardinal numbers, as modelled in a neo-Fregean arithmetic), there are numbers, like $\sharp\mathbb{N}$, the number of finite naturals, that *equal* their own successor. Everyday mathematical practice need not settle on one way of understanding the concept "number", and we get away with not distinguishing these concepts in our everyday arithmetical life. We tend not to concern ourselves with abstract generalisations about numbers. Our practices are settled-enough to get by, the general rules are nailed down only when our aims require it—such as when we build an exact arithmetic calculator, start doing abstract mathematics, or get into an argument in the playground about

[10] The formal theory of arithmetic in Whitehead and Russell's *Principia Mathematica* (1925, 1927) did a lot of heavy lifting before deriving the simple arithmetic truth that $1 + 1 = 2$ at ∗110·04.

[11] It is immediate that $0 \neq 0 + 1$ by the first axiom of Peano Arithmetic (since $s(x) = x + 1$), and it is also immediate that $n \neq n + 1 \rightarrow s(n) \neq s(n + 1)$ by the second axiom. So, by induction, *no* number is its own successor.

whether there is a biggest number. When we stray into those areas, we need to become more rigorous and define our terms well-enough for the task at hand.

So, is it *correct* to say that there is some number n where $n = n + 1$? To get a *useful* answer to that question, we must be more specific about how we will interpret the word "*number*" in the question. Our everyday practice is unsettled enough to allow for different ways of settling this issue.

So, suppose I come across a neo-Fregean calculator, which is able to solve elementary equations. I ask it to solve the equation $x = x + 1$, and it returns an answer, $x = \sharp\mathbb{N}$, rather than saying there is no solution. Should I conclude, then, that there is some amount of money I could have in my bank account such that adding one pound makes literally *no* difference to my balance? Not unless there is some bank that allows for a literally *infinite* balance. To interpret the results of such a calculator (which, by hypothesis, proves whatever is derivable in second order logic supplemented with Hume's Principle), we must attend to what such results *mean*. In neo-Fregean arithmetic, two predicates have the same number if and only if they are equinumerous and the only numbers so-defined that are equal to their own successors are *infinite*. To *interpret* the findings of such a calculator, we appeal to the the patterns that it instantiates, and use *those patterns* to understand the significance of the results the calculator produces.

* * *

What goes for understanding the counting and calculating functions of devices can also serve for interpreting the *assertoric* and *inferential* processes instantiated in proof assistants. Just as a calculator does not literally *count* anything, neither does a proof assistant *assert* anything. As we have seen, in a proof assistant implementing a dependent type theory, a *proof* represents a function that can take grounds for the premises as input, and delivers a ground for the conclusion as its output. I will argue that the structure of such a type theory—instantiated by a proof assistant system—stands to our everyday assertoric and inferential practice, as a particular formal theory of arithmetic—instantiated by some particular calculating system—stands to our everyday practice of counting and calculating.

To start, we should consider what *we* do when we assert and infer. It is no surprise that there are many proposals for how to understand assertion (Brown and Cappelen 2011), but for our purposes it is enough to briefly consider two different kind of approaches—those describing the function of assertion at the point of *production* (speaker norms), those at the point of *reception* (hearer norms).[12] *Speaker norms*: e.g. assert only what you *know* (the knowledge norm);

[12] There are also norms governing the *shared space* between the speaker and hearers (conversational norms), reflecting the role assertion has on the *common ground* in conversation. These are also important, but there is no space to discuss them here.

or assert only what is *true* (the truth norm), etc. *Reception norms*: to assert p entitles the hearer to (a) ask for a justification of the assertion and (b) to reassert p, handing back the request for justification to the original speaker.

Here it is not hard to see[13] how the *inferential* structure instantiated in proofs as represented in a proof assistant can play a role in allowing the deliverances of proof assistants to playing a role in assertion. The proof function shows how grounds of the premises of an argument may be used to produce grounds for the conclusion (Prawitz 2012). For the human who wants to *assert* the conclusion, given a context in which the premises have been granted, the proof is available to show *how* the conclusion follows from the premises (Restall *to appear*). So, something *proved* by a proof assistant becomes apt for assertion, provided that having such a ground is sufficient for knowledge, and therefore, truth.

So, *using* the proof as a means to produce grounds, production norms for assertion may be satisfied. Further, the proof of a proposition can be used to fulfil a *justification request* for the assertion, and thereby, so there is something to answer the hearer who asks for a justification request, or who refers back to the proof assistant to justify *their* re-assertion of the claim, should it be questioned. To represent a theorem formally in a proof assistant is taken as an epistemic achievement, reassuring the audience that indeed the proof is complete, and so, it can play the justificatory role when called upon.[14]

In the context of our attempt to understand the difference between constructive and classical logic, it is not sufficient to stop here. Our point of contention is not primarily about what *can* be proved with the aid of a proof assistant, but what *cannot* be so proved. When we learn that some result—such as the intermediate value theorem, mentioned above—cannot be given a proof in a proof assistant without making explicit classicality assumptions, does this have any significance? To answer this question, we should return to what, precisely, the proof assistant is *doing*, in the same way that if a calculator tells us that there is some number x where $x = x + 1$, we should attend to *what theory* the calculator is encoding. If this is a calculator solving equations in a neo-Fregean account of cardinalities (including *infinite* cardinalities), then all is well and good. If this is a calculator programmed to solve equations in the finite ordinals, then something has gone wrong somewhere. What is the corresponding account of the constructive

[13] In this brief discussion, I do not have space to consider the issues arising concerning how the representations in the code of the proof assistant can be *read* as sentences of a natural language that a user might understand, in just the same way that the inputs/outputs of a calculator can be read as denoting numbers. This is a non-trivial requirement, but not one I have time to discuss here.

[14] See Section 2 of Jeremy Avigad's explanation of the role of proof assistants (2024) for an account of this epistemic safeguarding role. The rest of that paper recounts *other* advantages of using proof assistants.

invalidity of the intermediate value theorem?[15] It is that there is no function that supplies, for each continuous $f\colon [0,1] \to \mathbb{R}$ where $f(0) < 0$ and $f(1) > 0$ to provide a *ground* for the claim that there is some $x \in (0,1)$ where $f(x) = 0$. Further, there are different constructive mathematical "universes" inside which this formulation of the intermediate value theorem can be refuted, while there are other models (including *classical* spaces) inside which the theorem nonetheless holds.

This result has *epistemic* significance, if the standards of evidence in operation in the discussion are appropriately high. When we are doing constructive mathematics, the standard of evidence asks for constructive grounds, and proof assistants using dependent type theory are designed to model such grounds.[16] So, if a claim fails to have those grounds, it may be rejected as out of bounds, for not having met the standard of assertion appropriate when doing mathematics constructively. A bald assertion of an instance of the law of the excluded middle $p \vee \neg p$ in the context of a constructive proof may ruled out, since no grounds can in general be provided, since any such ground must bring with it the means to ground p, or to ground $\neg p$, and there is no way to do this, in general.

Constructive mathematics is recognisably a kind of *assertoric* and *inferential* practice, in which claims are made, and constructive proof is the coin by which assertions are justified. With the advent of proof assistants based on dependent type theory, many mathematicians are becoming fluent in constructive proof, and the practice is emerging into the mainstream of mathematics.

Note that nothing in *this* explanation of the rise of constructive mathematics leads inexorably to favouring mathematical anti-realism over realism. The importance and usefulness of the constructive practice is motivated on internal mathematical grounds, and not by any particular view of the metaphysics of mathematics (Bauer 2018).

* * *

All that being said, constructive mathematics is not the *only* way that the norms of mathematical proof are made precise. The majority tradition in mathematical reasoning remains *classical*. A great deal of everyday mathematical reasoning appeals to the law of the excluded middle, to double negation elimination, and to other nonconstructive reasoning

---

[15] Note, though, that there is a reformulation of the intermediate value theorem that *is* constructively provable: if $f\colon [0,1] \to \mathbb{R}$ is continuous and for every $x \in [0,1]$ either $f(x) < 0$ or $f(x) > 0$, then for either for every $x \in [0,1]$, $f(x) < 0$ or for every $x \in [0,1]$, $f(x) > 0$ (Bauer 2018, Theorem 5.3).

[16] The higher standard of evidence in criminal legal proceedings compared to civil court is a useful analogy to keep in mind.

principles. These nonconstructive reasoning principles are found *everywhere*, both in mathematics, and in philosophy. Consider this relatively recent philosophical monograph:

> It is unclear whether there is here a genuine disagreement between Gadamer and Davidson. It is *undeniable* that someone may lack a concept that others have, and that we now have many concepts that no one had three hundred years ago. New concepts are continually introduced. They cannot always be defined in the existing language, but they can be explained by means of it; a study of how we acquire concepts, such as the concept of infinity, that could not even be expressed before their introduction would be highly illuminating. It is also *undeniable* that we can now recognize, of certain concepts that were used in some previous age, that they were incoherent or confused. (Emphasis mine.)

Here, the author is treating *it is undeniable that* as an intensifier, twice in quick succession. It would be a very strange thing, in the context of this discussion, were one to agree with the author and continue "yes, I cannot *deny* that someone may lack a concept that others have… but I do not see why it follows that I should *grant* it." Yet, the claim that *it is undeniable that p* is a form of double negation, since the denial, $\neg p$, is ruled out. The natural reading of this passage to take the author to be committed to the inference, here, from $\neg\neg p$ to p.[17]

We need not treat this as a *mistake*—it is appropriate, in the given context. Here is one way to understand that context. It is very natural to think that there is a certain kind of discourse in which we seek to *settle issues*. We would like to know whether p holds or not, and to rule out *one* of these options out is to leave the other. In this course of reasoning, the author asks us whether someone may lack a concept that others have … or *not*. The *no* case is ruled out, and so, only *yes* remains. It wins by being the last option standing, not necessarily because it has been given any positive (constructive) ground. This move—according to which, showing that something is *undeniable* is enough to show that it is true—is at the heart of a certain kind of deductive reasoning, and it is one that we make, time and again in our thought and our talk, even if we also practice constructive mathematics, in which we refrain from applying double negation elimination.[18]

Rather than asking whether the classical practice of inference is *correct* or not, let's consider what it practice is good for, and what it isn't—by analogy with the idea that different and

[17] This is a cheeky example, since it is an extract from *The Nature and the Future of Philosophy*, by Michael Dummett (2010, p. 94). Such reasoning is ubiquitous in philosophy, and elsewhere.

[18] If you start off as a committed constructivist, you can understand the family of *settleable issues* as given by the *negations* of propositions. The inference from $\neg\neg\neg p$ to $\neg p$ *is* constructively valid, and so, if we restrict attention to the constructive universe of *negative* propositions, we see that *it* behaves classically.

diverging *counting* practices might in an important sense be equally correct, and can be seen to be so when we understand that we are doing different things when we use *cardinal* numbers and when we use *ordinals*.

"Issue settling" discourse is fundamentally *bilateral* (taking *yes* and *no*, or assertion and denial, on a par).[19] Since $p \lor \neg p$ is *undeniable* (as we saw above) it follows that we have grounds for $p \lor \neg p$. We have not suddenly been able to ground one disjunct of $p \lor \neg p$ or the other—we have settled it only because it is undeniable, and not necessarily because we have any positive ground for p or for $\neg p$. Restricting ourselves to classical inference (and imposing the bilateral inference norms) means that we might be in a position to assert a disjunction without possessing a ground for either disjunct. Similarly, we may be able to categorically classically prove $\exists x \phi(x)$ without thereby constructing some term t where we can prove $\phi(t)$. We can prove $\exists x \phi(x)$ on the constructively unacceptable basis of a refutation of $\neg \exists x \phi(x)$.

What we *lose* in terms of the constructive power of assertion, when adopting classical reasoning principles, we *gain* with regard to the ability to express denial. Consider some domain of constructive mathematics, and some proposition A where we have no ground for $A \lor \neg A$, and furthermore, we *know* that we have no ground. Then someone asks us the question: is it the case that A? What can we say? We cannot answer *yes* (since A has no ground) and I cannot answer *no* (since $\neg A$ has no ground). Our constructive theory will have some *models* where A holds (since $\neg A$ fails, A is at least *consistent* with our theory), and some models where $\neg A$ holds (since A fails, $\neg A$ is consistent with our theory), but the fact that my theory is incomplete, and has two extensions, one where A holds and another where $\neg A$ holds does not mean that our indecision about $A \lor \neg A$ a matter of *ignorance* that might be settled with more information. Such ignorance is consistent with a classical theory, in which $A \lor \neg A$ *is* true, but our theory does not decide on which disjunct holds. The constructive reasoner wants to be able to *rule A out*, without going so far as to say that $\neg A$ is true. But to do this, constructively speaking, requires some kind of semantic ascent—we can say A is not *proved*, or A is not *known*, or some such thing.[20]

[19] There is more to say about the form of bilateralist inference, and the literature has a number of different proposals (Incurvati and Schlöder 2023, Restall 2005, Rumfitt 2000). The most direct way to understand the shift from constructive to classical proof is to expand our language to include a primitive speech act of *denial* alongside assertion (write the denial of p as '$\not{p}$'), with two structural rules connecting them: (1) from A and $\not{A}$ the contradiction $\bot$ follows, and (2) if we can derive a contradiction from the assumption $\not{A}$ (that is, if A is *undeniable*) then we can derive the conclusion A, discharging that assumption (Restall 2023). Given this background context, the harmonious proof rules Dummett takes to be semantically neutral behave classically: since $p \lor \neg p$ is undeniable, we can now *prove* it, without having to revise the inference rules for any of the connectives. We have expanded what counts as a proof (since we are more generous toward counts as ground for an assertion) and so, without changing the rules of any *connective*, more can be proved.

[20] Or we can say that the statement A is a constructive *taboo*: a principle which is not *false*, but which violates the *spirit* of constructive mathematics (see, e.g. Rathjen 2023, Section 1.2.1). Typically, taboo statements are

But this, it seems, is to change the subject from whatever it was we were talking about when asking whether A holds. We have not answered the question *about whether A or not*, we have only said something about our state of knowledge, or of our theory. If I restrict myself to constructive reasoning about some domain (whether that be mathematics or something else), I can only go so far in describing what is going on with the phenomena at hand.

* * *

We return, then, to the divide between *realism* and *anti-realism*, which was Dummett's original concern. Some classical mathematicians express their preference for classical mathematics over constructive reasoning in realist terms: their theory tells them that $A \vee \neg A$ and they would like to discover which disjunct is true, because the phenomena they study is *really* one way or the other. They are studying *the numbers* (the sets, the topological spaces, or whatever else…) and the success criterion is whether or not those descriptions are correct, not whether we are able to *construct* grounds for our claims.

There is something to this intuition: if we picture the phenomena in this way, we implicitly treat each issue as in fact *settled* (by Reality, with a capital "R", I suppose) and so, treating all of our claims as we theorise as issues that may be settled one way the another is appropriate. The realist picture fits naturally with classical reasoning. To restrict the grounds for our reasoning to what can be explicitly and positively constructed when the phenomena at hand might exceed our grasp, seems to be an artificial restriction if the aim is correct description. This does not mean that the restriction has no *point*. You can be as realist as you like about the mathematical universe, and still see the value of constructively theorising about that universe. Here, we return to the first of the two perspectives on constructive mathematics mentioned above. On this view, we may not be able to constructively prove all the classical *facts* about mathematics: constructive mathematics is a restricted subset of classical mathematics.

The reverse connection between realism and classical reasoning, is harder to establish. There is no reason to think that *classically* reasoning about a phenomenon means that there is some implicit realist commitment to it, over and above what is incurred in the use of constructive reasoning. It is well known that we can take a constructive theory (say, of arithmetic, thought of as a construction of the thinking subject, and not the description of some independently existing "realm"), and we find *inside* it, a perfectly classical theory, if we focus on the *settlable issues* in our language (the sentences of the form $\neg A$).[21] When might be tempted to say, in our

---

true in *classical* models of a constructive theory, but fail in other interesting models of the theory which have useful or interesting constructive features.

[21] This is one way to understand the Gödel–Gentzen double negation translation, which embeds classical Peano Arithmetic inside the constructive *Heyting* Arithmetic (Gödel 1933, Gentzen 1933). If we can justify a

native constructive tongue $A \vee B$, we instead say the classical substitute, $\neg(\neg A \wedge \neg B)$. When we might say $\exists x\ \phi(x)$, we say $\neg\forall x\ \neg\phi(x)$, and so on. As far as a *classical* semantics of disjunction and the existential quanifier goes, this makes no difference, but the result is a constructive vindication of classical reasoning about this domain, at the cost of making claims that are weaker than their constructively stronger counterparts. If there was no controversial metaphysical commitment before, we have incurred no new commitments, because we make made no new claims. The constructivist is able to translate the classical theoretical commitments into their own tongue, at no theoretical or metaphysical cost. We have here a vindication of the *second* perspective on constructive mathematics mentioned above: we can constructively recover classical theorems, when we isolate the classically-behaving propositions inside our constructive theory.

* * *

So, if all this is correct, when we say $p \vee \neg p$, is what we have said *true*? Here this depends on how we are *taken*. To take something to be true is to *evaluate* it. Speech is a communicative act, and requiring both speaker and audience. If the audience treats our claim constructively, it *may* have no proof, and thus, fail to meet its mark. (It may not meet the standard of evidence required for admission in *this* court). Note, though, that to say that it is *not* the case that $p \vee \neg p$ would be to exceed those very same constructive bounds. It is not as though we have grounds for $\neg(p \vee \neg p)$, either.

If we treat the claim $p \vee \neg p$ as expressing an *issue* to be settled, with all the classical norms of reasoning applying, then the answer is *yes*. It is true, since it is undeniable.

Notice, though, that to ask the question of whether $p \vee \neg p$ is true or not is simply, again, to ask about $(p \vee \neg p)$. The question has been asked, and we are in the business of evaluating it. To evaluate it *well*, it seems best to pay close attention to the norms we are applying, and to reflect on whether we *want* those to apply, instead of taking one and only one set of evaluative norms as given.

References

Avigad, Jeremy, Leonardo de Moura, and Soonho Kong 2023: *Theorem Proving in Lean*, Online manual (Release 3.23.0) https://leanprover.github.io/theorem_proving_in_lean, Accessed September 2024.

Avigad, Jeremy 2024: "Mathematics and the Formal Turn," *Bulletin of the American Mathematical Society*, **61**, pp. 225–240. https://doi.org/10.1090/bull/1832

---

constructive arithmetic on anti-realist grounds, then classical arithmetic, understood in this way, proves no more problematic.

Bacon, Andrew, 2023: *A Philosophical Introduction to Higher-Order Logics*, Routlege.

Bauer, Andrej 2016: "Five stages of accepting constructive mathematics," *Bulletin of the American Mathematical Society*, **54**, pp. 481–498. https://doi.org/10.1090/bull/1556

Beall, Jc and Greg Restall 2006: *Logical Pluralism*, Oxford: Oxford University Press.

Bishop, Errett and Douglas Bridges 1985: *Constructive Analysis*, Springer-Verlag.

Bridges, Douglas, Hajime Ishihara, Michael Rathjen and Helmut Schwichtenberg (editors) 2023: *Handbook of Constructive Mathematics*: Cambridge, Cambridge University Press.

Bove, Ana, Peter Dybjer and Ulf Norell: 2009. "A Brief Overview of Agda – A Functional Language with Dependent Types" Stefan Berghofer, Tobias Nipkow, Christian Urban, Makarius Wenzel (editors) *Theorem Proving in Higher Order Logics. TPHOLs* 2009. Lecture Notes in Computer Science, vol 5674. Berlin, Heidelberg: Springer.

Brown, Jessica and Herman Cappelen (editors) 2011: *Assertion: New Philosophical Essays*, Oxford: Oxford University Press.

Davis, Martin (editor) 1965: *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvable Problems and Computable Functions*. Hewlett, New York: Raven Press.

Dummett, Michael 1959: "Truth". *Proceedings of the Aristotelian Society*, Supplementary Volume, **59**:1, pp. 141–162.

Dummett, Michael 1977: *Elements of Intuitionism*, Oxford: Oxford University Press.

Dummett, Michael 1991: *The Logical Basis of Metaphysics*, Cambridge, Mass: Harvard University Press.

Escardó, Martín Hötzel, and collaborators, 2024: *TypeTypology*, an online repository of univalent mathematics written in *Agda*. https://www.cs.bham.ac.uk/~mhe/TypeTopology (Accessed September 27, 2024).

Fritz, Peter, and Nicholas K. Jones (editors) 2024: *Higher-Order Metaphysics*, Oxford: Oxford University Press.

Gentzen, Gerhard 1933: "On the relation between intuitionist and classical arithmetic." In M. E. Szabo (1969) *The Collected Papers ofGerhard Gentzen*, Amsterdam: North-Holland, pp. 53–67.

Gödel, Kurt: 1933: "Zur intuitionistischen Arithmetik und Zahlentheorie." *Ergebnisse eines mathematischen Kolloquiums*. Heft 4, 34–38. English translation: "On intuitionistic arithmetic and number theory", in Davis (1965), pp. 75–81.

Green, Karen 2002: *Dummett: Philosophy of Language* (Key Contemporary Thinkers), Cambridge, UK: Polity.

Harper, Robert 2016: *Practical Foundations of Programming Languages*, (second edition), Cambridge, UK: Cambridge University Press.

Heyting, Arend 1956: *Intuitionism: An Introduction*, Amsterdam: North Holland.

Hyland, Martin 1982: "The Effective Topos", in *L. E. J. Brouwer Centenary Symposium*, edited by A. S. Troelstra and Dirk van Dalen, Amsterdam: North-Holland, pp. 165-216.

Incurvati, Luca and Julian J. Schlöder 2023: *Reasoning with Attitude: Foundations and Applications of Inferential Expressivism*. New York: Oxford University Press.

Lambek, Joachim and Philip J. Scott 1986: *Introduction to Higher-Order Categorical Logic*: Cambridge: Cambridge University Press.

Lean Community, The 2024: *A Mathlib Overview*, https://leanprover-community.github.io/mathlib-overview.html (Accessed September 27, 2024).

Martin-Löf, Per 1984: *Intuitionistic Type Theory*, Naples: Bibliopolis.

Peregrin, Jaroslav 2014: *Inferentialism: Why Rules Matter*, Basingstoke: Palgrave Macmillan.

Prawitz, Dag 2012: “The Epistemic Significance of Valid Inference”, *Synthese*, **187**:3, 887–898.

Rathjen, Michael 2023: “An Introduction to Intuitionistic Logic”, in Bridges, Douglas, Hajime Ishihara, Michael Rathjen and Helmut Schwichtenberg (editors) *Handbook of Constructive Mathematics*: Cambridge, Cambridge University Press, pp. 3–19.

Restall, Greg 2005: “Multiple Conclusions” in P. Hájek, L. Valdés-Villanueva, and D. Westerståhl, editors, *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*, London: KCL Publications, pp. 189–205.

Restall, Greg 2023: “Structural Rules in Natural Deduction with Alternatives”, *Bulletin of the Section of Logic*, **52**:2, pp. 109–143. https://doi.org/10.18778/0138-0680.2023.6

Restall, Greg *to appear*: “Questions, Justification Requests, Inference, and Definition”, to appear in *Synthese*. Preprint available at https://consequently.org/writing/qjrid/

Rumfitt, Ian 2000: “Yes” and “No”, *Mind*, **109**:436, 781–823.

Rumfitt, Ian 2015: *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*, Oxford: Oxford University Press.

Steinberger, Florian 2011: “What harmony could and could not be,” *Australasian Journal of Philosophy*, **89**:4, 617–639.

Taylor, Barry M. (editor) 1987: *Michael Dummett: Contributions to Philosophy*, Dordrecht: Martinus Nijhoff Publishers.

Tennant, Neil 2017: *Core Logic*: Oxford: Oxford University Press.

Whitehead, Alfred North and Bertrand Russell (1925, 1927), *Principia Mathematica* (Second Edition), Cambridge: Cambridge University Press.

Williamson, Timothy 2006: “Must Do Better.” *Truth and Realism*, edited by Patrick Greenough and Michael P. Lynch, Oxford: Oxford University Press, pp. 177-187, reprinted as an appendix in *The Philosophy of Philosophy* (Willamson 2007).

Williamson, Timothy 2007: *The Philosophy of Philosophy*, Malden, Mass: Blackwell Publishing.

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

Wright, Crispin 1983: *Frege’s Conception of Numbers as Objects*, Aberdeen: Aberdeen University Press.

Wright, Crispin 1992: *Truth and Objectivity*, Cambridge, Mass: Harvard University Press.