What can we mean?—on practices, norms and pluralisms

Greg Restall, *Arché*, Philosophy Department, University of St Andrews
Handout for the *Aristotelian Society*, October 2024

Michael Dummett spoke to this Society 65 years ago (Dummett 1959), inaugurating a long-running debate over semantic realism and anti-realism, and the role of logic as a necessary prolegomenon to fruitful discussion in metaphysics.

The *Issue*

Dummett argued that not all traditional logical principles are metaphysically neutral.

Some logical principles (those of *intuitionistic* or *constructive* logic) are self-justifying on neutral semantic grounds.[1]

Double negation elimination (DNE), the inference from $\neg\neg p$ to p, is *not* self-justifying. To adopt DNE involves metaphysical commitment that skews debate in favour of the realist.

The philosophically neutral logical perspective, acceptable to all sides, would be to accept only *intuitionistic logic*.

The debate over whether Dummett is *correct* reached its peak in the 1980s and early 1990s.

That peak has receded: e.g. Oxford philosophical logic is now dominated by discussions of higher-order modal logic (see Williamson 2013).

Almost everywhere, Dummett's concerns are sidestepped, rather than addressed head on.[2]

The situation is reversed in *mathematics*.

Constructive mathematics was in the minority in the second half of the 20th Century (Bishop and Bridges 1985, Bridges, *et al* 2023).

With the rise of *proof assistants* (Avigad 2024) like *Agda* (Bove, *et al* 2009) and *Lean* (Avigad, *et al* 2023). An increasing number of mathematicians are doing *constructive mathematics*, in line with Dummett's scruples.

What is happening here?

The use of proof assistants raises issues of broader philosophical interest. *Semantics* matters at the human/machine interface. How should we understand the role of computational systems in our own practices of explanation, inference, and justification?

I approach this as a *pluralist* (Beall and Restall 2006), though if you are inclined to find one framework as *correct* and another *incorrect*, I'll indicate where the different choice points lie along thew way.

What Proof Assistants *Do*

Proof assistants are just one way computers have changed the face of mathematics in the late 20th and early 21st Century.

A *proof assistant* 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).

Proof assistants, in general, could be totally agnostic about the choice of logical principles to use, and could be practicing "formalists".

Contemporary proof assistants are *not* formalists. *Agda* (Bove et al. 2009) and *Lean* (Avigad et. al. 2023) are more opinionated about what kind of thing a proof *is*, and the *logic* they encode is constructive.

Agda and *Lean* treat proofs as *functions*, where a valid deduction is represented as a *function* that transforms grounds the premises into grounds for the conclusion.[3] This is straightforward, except for the unspecified notion of *ground*. What *is* a ground, in general (Prawitz 2012)?

As far as the *logic* goes, proofs are functions that combine grounds and supply new grounds from old in regular ways.

---

[1] For a given logical concept we can treat the rule *introducing* a judgement of that form as its definition. 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. See Dummett's *Logical Basis of Metaphysics*, Chapter 11 (1991). 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 for specific logical concepts 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).

[2] See Williamson's "Must Do Better" (2006) for a vivid and opinionated account of *why* this debate was sidestepped.

[3] The ground for a conjunction is *pair*, the first element of which is a ground for the first conjunct, the second, for the second conjunct. The ground for a conditional ($A \to B$) is a function transforming grounds from the antecedent A to grounds for the consequent B, etc.

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.[4]

The reasoning principles that arise naturally here 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*, proof assistants are special kinds of *functional programming languages*.

So, mathematicians learn to express their results in the language of dependent type theory (Escardó and collaborators 2024, Lean community 2024).

Mathematics encoded in this way is *constructive*.

A proof of a disjunction $A \vee B$ 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 $\neg p$.[5]

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 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).

This is strikingly similar to Dummettian semantic anti-realism where distinctively classical principles are an optional extra, to be adopted when the metaphysics asks for it.

This well established, if still minority, practice of constructive mathematical theorising raises a question.

How are we to understand the relation between constructive mathematics and classical mathematics?

OPTION 1: Constructive mathematics is a *restriction* on classical mathematics.

> …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)

OPTION 2: Constructive mathematics is an *expansion* of classical mathematics.

> …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)

What *should* we say? Is constructive practice a *restriction*, or an *expansion* of classical reasoning?

An *Analogy*

Consider the 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 that work?

We acquire our knowledge of basic facts of arithmetic by way of an education involving *counting* things.

Calculators do not count things, but the system involves the reliable manipulation of patterns.

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* but that is an *infinite* task, since there are infinitely many arithmetical equations. We need a finitary way to specify these infinitely many facts.

---

[4] 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).

[5] $\neg\neg(p \vee \neg p)$, 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, $p \vee \neg p$ is constructively *undeniable*.

PEANO ARITHMETIC:

Here, there are three axioms governing the notion of *zero* and the *successor* function s.[6]

- $sx \neq 0$
- $sx = sy \to x = y$
- $x \neq 0 \to \exists y\, x = sy$
- $x + 0 = x$
- $x + sy = s(x + y)$
- $x \times 0 = 0$
- $x \times sy = (x \times y) + x$
- $\left[\phi(0) \wedge \forall x\big(\phi(x) \to \phi(sx)\big)\right] \to \forall x \phi(x)$

If the output of our calculator *agrees with* the judgements of Peano Arithmetic, it is reliably doing finite arithmetic. But it need not be *counting* in any sense.

NEO-FREGEAN ARITHMETIC:

Other formalisations of arithmetic *do* make some kind of use of a notion of counting.

For any one-place predicate F, we have a singular term $\sharp F$, to be read as "the number of Fs", 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.

With the help of lambda abstraction,[7] we introduce the finite numbers using 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.

Define addition by setting $\sharp F + \sharp G$ to be $\sharp \lambda x(Fx \vee Gx)$ when nothing is both F and G, and continuing from there.

If our calculator's output agreed with a *neo-Fregean* theory, it would *also* count as recognisably doing arithmetic.

A calculator might 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. (The same holds for *you* and for *me*.)

These counting practices agree on a great deal, but disagree at the margins: Is there a number n where $n = n + 1$?

The answer is *no* for Peano Arithmetic, and the answer is *yes* in a neo-Fregean arithmetic.[8]

A competent user of arithmetic vocabulary could well find that their own concept of number simply *does not settle the issue* as to whether a number can be its own successor.

So, is it *correct* to say that there is some number n where $n = n + 1$? To get a *useful* answer to this question, we must be more specific about how we will interpret the word *"number."*

If I have a calculator and I ask it to solve the equation $x = x + 1$, to interpret the significance of the *answer* of that calculator, I must have at least *some* sense of what the calculator is *doing*.

The *Claim*

> *What goes for understanding the* counting *and* calculating *functions of devices also goes for interpreting the* assertoric *and* inferential processes *instantiated in proof assistants.*

There are many proposals for how to understand *assertion* (Brown and Cappelen 2011).

SPEAKER NORMS: e.g. assert only what you *know* (the knowledge norm); or assert only what is *true* (the truth norm), etc.

HEARER 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.

The *proof function* in a proof assistant 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*).

---

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

[7] 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*.

[8] In Peano arithmetic, this is an easy proof by induction. Zero is not its on successor by the first axiom, and by the second, if the *successor* of x is its own successor, so is x, so, using induction, *no* number is its own successor. In Neo-Fregean arithmetic, 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.

Something proved by a proof assistant becomes apt for assertion, provided that having such a ground is sufficient for knowledge, and therefore, truth.

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 in a proof assistant is an epistemic achievement.[9]

However, 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 *cannot* be given a proof in a proof assistant without making explicit classicality assumptions, does this have any significance?

(Recall the issue of understanding what it means when our calculator tells us that there is—or isn't—a solution to the equation $x = x + 1$.)

What is the corresponding account of the constructive *invalidity* of the intermediate value theorem?[10] It is that there is no function that supplies, for each continuous $f : [0, 1] \to \mathbb{R}$ where $f(0) < 0$ and $f(1) > 0$ a *ground* for the claim that there is some $x \in (0, 1)$ where $f(x) = 0$.

This result has *epistemic* significance, if the standards of evidence in the discussion are appropriately high.[11]

If a claim fails to have those grounds, it may be rejected. An assertion of $p \vee \neg p$ in the context of a constructive proof may be ruled out, since in general, we have no means to ground an arbitrary p or an arbitrary $\neg p$.

Constructive mathematics is *recognisably assertoric* and *inferential.* Claims are made, and constructive proof is the coin by which they are justified.

Note: nothing here favours mathematical anti-realism over realism. The motivation is on internal mathematical grounds (Bauer 2018).

This said, the majority tradition in mathematical reasoning is *classical*. Nonconstructive reasoning is *everywhere*, in mathematics, and philosophy. Consider this:

> 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*.)

The author treats *it is undeniable that* as an intensifier.

(It would be strange to *agree* with the author, but to 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.")

The claim that *it is undeniable that p* is a form of double negation. The natural reading is to take the author to be committed to the inference from $\neg\neg p$ to p.[12]

There is a kind of discourse in which we seek to *settle issues*. We want to know whether p holds or not. To rule out *one* option is to leave the other. It wins by being the last option standing, not necessarily because it has been given any positive (constructive) ground.[13]

This is a norm applying to issues that diverges from the norms applying in constructive reasoning.

Let's *not* ask whether the *constructive* norms or whether *issue-settling* norms are objectively *correct*, by analogy with asking whether *cardinal* or *ordinal* numbers are *the correct numbers*.

---

[9] 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.

[10] Note that a *reformulation* of the intermediate value theorem *is* constructively provable: if $f : [0, 1] \to \mathbb{R}$ 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).

[11] Consider the higher standard of evidence in criminal legal proceedings compared to civil court.

[12] This cheeky example is an extract from *The Nature and the Future of Philosophy*, by Michael Dummett (2010, p. 94).

[13] 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.

"Issue settling" discourse is fundamentally *bilateral* (taking *yes* and *no*, or assertion and denial, on a par).[14] Since $p \vee \neg p$ is *undeniable* we thereby have grounds for $p \vee \neg p$.

(We have settle it only because it is *undeniable*, and not 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 can ground 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 $\phi(t)$.

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 *rejection* by way of assertion.

Consider some domain of constructive mathematics, and some proposition A where we have no ground for $A \vee \neg A$, and we *know* that we have no ground.

We are asked: is it the case that A? What can we say?

We cannot answer *yes* (since A has no ground) and we cannot answer *no* (since $\neg A$ has no ground).

Our indecision about $A \vee \neg A$ is not merely a matter of *ignorance* that might be settled with more information.

(Such ignorance is consistent with a classical theory, in which $A \vee \neg A$ *is* true, but our theory does not decide on which disjunct holds.)

The constructive reasoner would like to *rule A out*, without going so far as to say that $\neg A$ is true.

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,[15] which involves changing the subject: we have not answered the question *about whether A or not*.

If I restrict myself to constructive reasoning about a domain, I can go only *so* far, describing the phenomena at hand.

Some *Consequences*

Return, to the divide between *realism* and *anti-realism*.

Some classical mathematicians express their preference for classical mathematics 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.

There is something *to* this: they implicitly treat each issue as in fact *settled* (by Reality) and so, treating all of our claims as we theorise as issues that *may* be settled one way the another is appropriate.

To restrict the grounds for our reasoning to what can be constructed when the phenomena exceed our grasp, seems artificial if the aim is *correct description*.

This does not mean that the constructive restriction has no *point*. You can still value of constructively theorising for its other virtues. (This way lies OPTION 1 above: constructive mathematics is a subset of classical mathematics.)

However, there is no reason to think that *classically* reasoning about a phenomenon means that there any more realist commitment implicit over and above a constructive theory.

Take a constructive theory: we find *inside* it a perfectly classical theory, if we focus on the *settlable issues* in our language (the sentences of the form $\neg A$).[16] When might be tempted to say, in our 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 $\neg \forall x \neg \phi(x)$. As far as a *classical* semantics goes, this makes no difference, but the result is a constructive vindication of classical reasoning, at the cost of making claims that are (constructively) weaker than their constructive counterparts.

If there was no controversial metaphysical commitment before, we incur no new commitments, because we make no new claims. The constructivist is able to translate classical theoretical commitments into their own tongue, at no change in *ontology*.

---

[14] The literature has a number of different proposals considering bilateralism (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 'p'), with two structural rules connecting them: (1) from A and A the contradiction $\bot$ follows, and (2) if we can derive a contradiction from the assumption A (that is, if A is *undeniable*) then we can derive the conclusion A, discharging that assumption (Restall 2023). Given this context, the harmonious proof rules Dummett takes to be semantically neutral behave *classically*: since $p \vee \neg p$ is undeniable, we can now *prove* it, using Dummett's own definitions for the connectives.

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

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

This perspective vindicates OPTION 2 mentioned above: we can constructively recover classical theorems when we isolate the classically-behaving propositions inside our constructive theory. Classical commitment is found *inside* a constructive theory.

A *pluralist* does not have to endorse one option and reject the other, any more than a mathematician has to endorse one kind of number and reject the others. However, pluralism is not mandatory: if you have a preferred unitary theory of assertion and of propositional content and you are convinced that propositions understood in *that* way are all that should count as *propositions*, properly so-called, then nothing I have said here need count as a decisive argument against that view.

But let's say you are tempted by this thoroughgoing pluralism, about logic and about propositional content. 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*. Speech is a communicative act, requiring a speaker and audience. If the audience treats our claim constructively, it *may* have no proof, and thus, fail to meet that mark. (It may not meet the standard of evidence required for admission in *this* court). 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 that to ask the question of whether $p \vee \neg p$ *is true* or *not* is simply 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*, we must pay close attention to the norms we apply, and to reflect on whether we *want* to apply them, 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* (v 3.23.0) https://leanprover.github.io/theorem_proving_in_lean

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

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

Beall, Jc and Greg Restall 2006: *Logical Pluralism*, OUP.

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

Bridges, Douglas, Hajime Ishihara, Michael Rathjen and Helmut Schwichtenberg (ed.) 2023: *Handbook of Constructive Mathematics*, CUP.

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 (ed) *Theorem Proving in Higher Order Logics. TPHOLS 2009*. LNCS, vol 5674. Springer.

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

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

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

Dummett, Michael 1977: *Elements of Intuitionism*, OUP.

Dummett, Michael 1991: *The Logical Basis of Metaphysics*, HUP.

Escardó, Martín Hötzel, and collaborators, 2024: *TypeTypology*. https://www.cs.bham.ac.uk/~mhe/TypeTopology.

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

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

Harper, Robert 2016: *Practical Foundations of Programming Languages*, (second ed.), CUP.

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

Incurvati, Luca and Julian J. Schlöder 2023: *Reasoning with Attitude*. OUP.

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

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

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, *et al* (ed.) *Handbook of Constructive Mathematics*: CUP, 3–19.

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

Restall, Greg 2023: "Structural Rules in Natural Deduction with Alternatives", *Bulletin of the Section of Logic*, **52**:2, 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*. https://consequently.org/writing/qjrid/

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

Williamson, Timothy 2006: "Must Do Better." *Truth and Realism*, ed, by Greenough and Lynch, OUP, 177-187, reprinted in *The Philosophy of Philosophy* (Willamson 2007).

Williamson, Timothy 2007: *The Philosophy of Philosophy*, Blackwell.

Williamson, Timothy 2013: *Modal Logic as Metaphysics*, OUP.

Wright, Crispin 1983: *Frege's Conception of Numbers as Objects*, AUP.