Truth Through Proof, by Alan Weir. Oxford: Clarendon Press, 2010. Pp. xiv + 281. H/b £35.

Speaking in the metaphysics room, there are no numbers. There are no four-groups, no models built out of sets, nothing like that. But despite this, when speaking in the mathematics room, we can still correctly assert the likes of 'for every prime there is a greater one', 'the Klein four-group is the smallest non-cyclic group', 'there are non-standard models of PA'.

What makes it correct to assert such claims if there are no abstracta 'out there'? Answer: facts about what we can *prove*. Since we don't want to let abstracta back in through the rear door, that has to mean: facts about the concrete proofs we can produce in real-world practice.

Such, in the baldest terms, is the claim of *Truth Through Proof*. But isn't this going to be radically revisionary about the reach of mathematical truth? Concrete proofs are thin on the ground: and if such proofs are all that can make for correctness in arithmetic – and for brevity's sake, let's concentrate on arithmetic in this review – doesn't a very limited strict finitism threaten? Won't the story be implausibly revisionary about what mathematicians mean by what they say and write? Surely when we talk about the infinitude of primes we aren't *reall*y talking about some concrete proof-inscription. And if mathematics is some formal game, how can it be usefully applied?

The 'neo' aspects of Alan Weir's advertised neo-formalism consist in extended answers to such questions. Weir thinks his brand of formalist can after all, with a good conscience, go along with standard infinitary mathematics (though perhaps bailing out before the esoteric upper reaches of set theory). The formalist isn't committed to attributing some quite implausible content to what mathematicians say and moreover can have a workable theory of applied mathematics. So this will indeed be formalism without tears.

Wonderful, if Weir can pull it off. But the devil, as always, is in the details. And there are a *lot* of details and some seeming digressions, piled high and often obscuring the shape of the emerging position. The book is tough going. Here I can only touch on a few themes.

A metaphysical story about correctness conditions, Weir insists in Ch. 2, can't always be parlayed into a theory of literal content (of what is grasped by the competent user). Recall, for example, the projectivist about morals who says that what makes a moral utterance correct (here, among us) is in the end a question of our attitudes, or of what our attitudes would be when thought through, or some such. Still, it is no part of the projectivist story that ‘abortion is often permissible’ is *about* attitudes. That utterance isn’t in the business of representing the metaphysics of the situation: indeed it isn’t in the business of genuine *representation* at all. It plays a different role in a non-representational mode of discourse. Similarly for claims that go-along-with-the-fiction like ‘Sherlock lived quite near Westminster’. Such claims have correctness conditions as moves in the intended game, something quite complicated to do with reasonable inferences from texts. But their content isn’t *about* texts.

So far, so good, perhaps. But for Weir these are at best half-analogies, for he doesn’t want to be a projectivist or fictionalist about mathematics. In making arithmetical claims like ‘74 + 46 = 120’ we aren’t projecting an attitude or playing along with a fiction. We are engaged in another kind of non-representational discourse, he says, making assertions in a ‘formal’ mode. But what does that mean?

Here’s a toy example, which isn’t Weir’s own but I think fairly introduces both the bones of his semantic picture and his story about formalist correctness conditions. It involves an abacus game developed in three stages.

Stage one. We teach children to play shuffling beads on an abacus. In fact the rules of the game are such that permitted moves track stages in the addition of two numbers under 100 (represented in the initial state of the abacus by units and then tens of one number on the top two rows, then units and tens of the other number on the next two rows) with the result represented by the state of the final three rows at the end of the game (there’s an extra row for a ‘carry’ digit). But the children don’t know this: they are taught allowed plays by pattern-recognition and ‘you can do *this*’ and ‘you can’t do *that*’ until they latch on.

Stage two. We now teach the children an additional routine. They learn to write in a ledger e.g. the symbols ‘74 + 46 = 120’ when they have just got from a certain initial configuration (the configuration that *we* could describe as representing 74 and 46 on the top four rows of the abacus) to a final configuration (that we say represents 120 on the bottom three rows). Note, however, it isn’t enough just to have ended up with the appropriate configuration of beads: for a legitimate entry in the ledger, the configuration must have been achieved by correct play.

Stage three. The rules are loosened. The children learn that correct 'addition' tokens can be allowed in the ledger whether or not they are cued to actual bead-shufflings. They learn that if the entry of an 'addition' is challenged, the challenge can met by going through the abacus game and getting the result the token is cued to, and the challenge is lost if they get a different result.

Now, in the first-stage game (at the G-level for short), the tokens the children are moving around – the beads on wires – have no significance. But, at the level of the written tokens, where the children are giving a kind-of-commentary (call this the C-level), the tokens in the new language game can be thought of as having a certain significance. There are now correctness conditions for the issuing of a token '74 + 46 = 120'.

However, although the correctness condition for issuing '74 + 46 = 120' is – as we would metatheoretically put it – that you can get from a certain initial state of the abacus to a certain final state via moves according to the rules of the abacus game, it would seem to be over-interpreting to suppose that this is what the 'equation' *means* for the child. After all, she need have no reflective grasp of e.g. the concept of a rule-of-the-game, and indeed no articulated descriptive concepts for initial and final states either. The tokens in her language game don't explicitly represent their correctness conditions as such.

This, then, gives us a toy language-game for which it's natural to say that (i) there are correctness conditions for issuing 'equation'-like tokens, (ii) these correctness conditions are given in terms of the availability of legitimate moves in an abacus game, but (iii) it would be over-interpreting to suppose that the players, in issuing such an 'equation'-like token *mean* that the correctness conditions obtain.

Now move from the toy model to an example Weir actually considers in Ch. 3. In this case, the stage-one game is DA, a de-interpreted version of decimal arithmetic. At this level there is a contentless symbol-shuffling game, this time involving not beads but tokens like '74 + 46 = 120'. Again, what passes for allowed plays is taught by example (not by giving explicit rules). Then, at the next stage, players are again taught a kind-of-commentary game. A player is taught to enter one of those equation-like tokens in a ledger if they are produced at the end of a DA game. Then, thirdly, the practice is expanded to allow a player to write down such a token even if they haven't done the DA 'derivation' routine, so long as they could respond to a challenge by 'doing the proof'.

This time, tokens of the very same equation-like type '74 + 46 = 120' get into the story twice over. Firstly, tokens can appear in the G-level DA symbol-shuffling game, where they are as empty of content as bead-arrangements in the abacus game. But second, a token can appear again at the C-level in a kind-of-commentary on the DA game, making a move in a language game which has the correctness condition that you can derive that sort of token inside DA.

As before, at the C-level, (i) there are correctness conditions for the issuing of equation-like tokens, (ii) these correctness conditions are given in terms of the availability of moves in a game, but (iii) it would again be over-interpreting to suppose that the players, in issuing such an equation-like token *mean* that the correctness conditions obtain.

Let's grant that. But does it imply that these C-level tokens aren't representational at all? They certainly express something quite different from an affective attitude or belief-suspending pretence (as in the case of moral discourse, or discourse that goes-along-with-a-fiction). In fact, the natural suggestion is that they serve to express inchoately representational beliefs about the G-level game: think of the sort of way that 'he's in pain' might be argued to represent – albeit inexplicitly – the presence of certain neural state. But Weir doesn't really consider such a possibility that C-level tokens are inchoately representational: so his arguments for outright non-representationalism about them seem far too quick.

Be that as it may, here's the pay-off. The neo-formalist claims that *our* basic arithmetical claims are like the tokens at the C-level cued to the DA game: they too have correctness conditions to be given in term of what can be proved. But there's semantic sugar on this formalist pill. For unlike the (strawman?) formalist who implausibly treats arithmetical sentences as contentless, or another kind of formalist who equally implausibly treats them as being explicitly about moves in a formal game, the neo-formalist has a story about content that treats our arithmetical claims as like the C-level claims in having content, but not so rich a content as to explicitly represent their correctness conditions (Weir would say that the content is not representational at all, but we don't have to buy that). Will the semantic sweetener be enough, though, to get us to swallow the formalism?

One obvious question is: how can we identify arithmetical truth with provability, given Gödel's incompleteness theorem? And proofs in exactly *which* derivation game give the correctness conditions for our actual claims? Something like the abacus game is too limited; likewise, surely, for quantifier-free DA. You might have thought it was incumbent on the formalist about arithmetic to really nail down the proof-providing practice *S* that we are using (at least up to equivalence-of-theorems). Oddly, Weir doesn't try to do this, but takes a rather open-ended view about what might count as a proof: he's an arithmetical formalist without a canonical formalism.

We'll return to Gödel; but first consider another theme. In his reconstruction of arithmetic, Weir's story starts with moves in some uninterpreted game, and then assigns content to '74 + 46 = 120' in its role as in effect commenting on moves in the game. And only *afterwards*, with the content of equations fixed by the liaisons to the uninterpreted formal game, do we then get to talk about applying the 'arithmetic' to the world (via 'bridge principles' which allow us to pass between nominalistically harmless worldly claims of the form '$\exists_n xFx$' and the corresponding 'The number of Fs = n', and then we can use equations in the obvious ways). A Frege or Wittgenstein would say that this is just topsy turvy, and that it's the embedding in the applied practices of counting and adding and multiplying that gives the content to pure arithmetical equations, properly so called. Weir dissents.

In fact, as Weir acknowledges in Ch. 5, his general position on how mathematics gets applied is very close to Field's. True, there are significant disagreements: the semantic gloss on what is happening is different (on the maths side, it is fiction for Field, and a putative 'formal' mode of assertion for Weir), and their metaphysical stories are different (Weir is happy to populate the world with considerably more physical properties than Field). But readers who thought that Field's project was not just wrong in detail but basically misguided aren't likely to find Weir's revised version much more to their taste. And they may also be rather suspicious of how Weir thinks he can establish the sort of conservativeness results that his Fieldian position surely needs. He appeals to the claim that we can 'model' e.g. an empirical theory plus arithmetical superstructure in ZFCU and prove a model-theoretic conservativeness theorem. But given that, officially, what makes *that* theorem true (speaking in the mathematics room) is according to Weir a concrete proof of it in an uninterpreted ZFCU game, we might well wonder why facts about the concrete symbol-shuffling in that formal game should make us trust a certain device for extending the machinery of an *empirical* theory.

The basic worry here – how can facts about plays in one game can tell us about something else outside that game – crops up again when we look at how Weir deals with what would seem to be damning problem with his position about what makes for correctness in arithmetic. Suppose *P* is a short true Δ_0 proposition of arithmetic whose proof in a standard system of arithmetic is astronomically long (there are many such!) – i.e. it has no practically possible concrete proof. Then, according to Weir's official semantic story, neither *P* nor not-*P* is correct. So, given that he gives a correctness-value semantics for the propositional connectives, it seems at first blush that he is not entitled to assert the classically (and intuitionistically!) true disjunction, either *P* or not-*P.* So either Weir is going to have to be radically revisionary about arithmetical truth, or he's going to have to give a special explanation of why he is still entitled to the application of excluded middle here.

He takes the second line. The idea seems to be this. Here is our practice of producing concrete formal proofs (in our system *S* for arithmetic, whatever it is). Like other happenings in the world, we can theorize about all this, doing some applied mathematical theory *M* which – like other bits of applied mathematics – no doubt idealizes the real-world practice of proving things in *S*. And now, the claim goes, we *do* have a concrete *M*-proof of the model's representative for the claim that either *P* or not-*P* (where *P* is our 'concrete undecidable'). So, in sum,

> the existence of concrete indeterminables [like *P*] should not inhibit reasoners from applying excluded middle to them so long as there is a concrete proof that, in a legitimate idealization, the image of the indeterminable is decidable in the formal sense. (p. 205)

But that's very puzzling. Why should facts about practical symbol shuffling in one formal game *M* (the facts which, for Weir, are the truth-makers of *M*-claims) give us any justification for a rule – excluded middle – added to a different game *S*? And anyway, while a practicable applied theory may of course idealize, if it is truth that we are after, a more accurate theory with 'friction and air resistance' trumps a smoother but over-idealizing theory. So won't a proof theory *M* about our actual concrete practices in producing token *S*-proofs model aim to reflect something of our actual limitations, and not idealize to allow wildly long proofs (for example)? But then *M* won't always establish that the 'image' of *P* is provable or is refutable in the (now somewhat less idealized) image of *S*.

Rather than trying to keep idealization to a workable minimum, however, Weir seems happy to go in the other direction. To summarize in his own words,

> So long as there exist concrete tokens of the theorems [and proofs?] … of the applied mathematics which one needs in the idealization, then it is permissible for a neo-formalist to use the theorems of the idealizing theory. … [This] reveals that the emphasis on the finite nature of proof which took hold in logic after Gödel's work in the 1930s is a harmful prejudice which should be abandoned. There should be no size limit on the abstract structures which can be invoked in idealizations of our finite corpus of mathematical utterances. (p. 10)

Which – to return to an earlier issue we raised – gives us an idea of how the Gödelian gap between truth and provability (as idealized) is supposed to be closed: just appeal to an infinitary idealization which isn't recursively enumerable so the incompleteness theorem doesn't apply.

Yet *this* will strike many readers as a rather astonishing turn for the formalist to take. We started with what seemed to be a robust emphasis on the honest toil of concrete proofs in an arithmetic system *S* (whatever it is) as making for arithmetical correctness. And now we are invited to see another source of arithmetical correctness, one still arising from concrete proofs, but now proofs in a (wildly!) idealizing applied mathematics modelling the producing of *S*-proofs. Which might look like a very considerable revision in the notion of arithmetical truth. Or perhaps – some will think – this just reflects what is in the end the rather slippery open-endedness and unclarity in the notion of proof that is supposed to make for arithmetical truth. Either way, I suspect many readers initially tempted by other aspects of Weir's neo-formalism will in the end find his pivotal Ch. 7 on idealization a real sticking point.

Peter Smith
Faculty of Philosophy
University of Cambridge
Sidgwick Avenue
Cambridge, CB3 9DA
UK