Ancestral Arithmetic and Isaacson's Thesis

Peter Smith

April 21, 2007

1 Isaacson's Thesis stated

First-order Peano Arithmetic (PA) is incomplete. The question naturally arises: what kind of sentences of PA's language L_A (that's 'the language of basic arithmetic', with the standard interpretation) can we establish to be true even though they are unprovable in PA?

There are two familiar classes of cases. First, there are sentences like the canonical Gödel sentence or the canonical consistency sentence for PA. Second, there are sentences like the Π_2 arithmetization of Goodstein's Theorem.

In the first sort of case, we can come to appreciate the truth of the Gödelian sentences which are undecidable in PA in part by reflecting on PA's consistency or by coming to accept instances of the Π_1 reflection schema for PA. And those routes surely involve deploying ideas beyond those involved just in grasping the basic arithmetical notions expressed in L_A and hence in understanding PA: to reason to the truth of the Gödel sentence, we need not just to be able to do basic arithmetic, but to be able to reflect on our practice.

In the second sort of case, we come to appreciate the truth of the sentences which are undecidable in PA by deploying transfinite induction or other infinitary ideas (we may even have to use assumptions about large cardinals, for the sort of cases explored by Harvey Friedman). So the reasoning again involves ideas which go beyond what's involved in grasping basic arithmetic.

Thinking about these sorts of cases, then, suggests a general conjecture, which we'll call

> *Isaacson's thesis* If we are to give a proof for any true sentence of L_A which is independent of PA, then we will need to appeal to ideas that go beyond those that are required in understanding PA (and more generally, but more vaguely, beyond those that are required to understand elementary pure arithmetic).

Is this true? Well, I'll suggest that the conjecture's fate turns at least in part on the provability or otherwise of a certain technical result. More precisely, I'll define a certain theory of arithmetic which I'll call *ancestral arithmetic*, and argue that the conceptual resources need to grasp PA in fact gives us a grasp of ancestral arithmetic (but also, arguably, no more than a grasp of ancestral arithmetic). So the question arises: is formalized ancestral arithmetic conservative over PA for sentences in the language L_A? If it is, then – as I'll argue – that is strong support for Isaacson's thesis. If it isn't, then Isaacson's thesis fails: there will be truths which are independent of PA but which still can be proved without going beyond what is given us in our our understanding of basic arithmetic .

To give this talk a spurious sense of excitement, I won't tell you in advance how things are going to pan out.

2 Hidden higher-order concepts?

Before proceeding, however, let me pause to distinguish the claim that I am calling Isaacson's thesis from the rather stronger version that is more or less explicitly to be found in Isaacson's 1987 paper.

Recall that, in that paper, Isaacson talks of 'hidden higher-order concepts'. Sometimes, in fact, he claims that the truths of arithmetic which are independent of PA in some sense themselves '*contain* essentially hidden higher-order, or infinitary, concepts'. I think we should resist such talk. For example, I'd want to firmly insist that the canonical Π_1 Gödel sentence is, as far as questions of content go, exactly on a par with other Π_1 sentences, and in no good sense contains any infinitary concept, hidden or otherwise. But in fact, some of his other remarks suggest that Isaacson's talk of truths which '*contain* hidden higher-order concepts' is rhetorical overkill. For example, consider the concluding words of his paper: 'truths in the language of arithmetic which lie beyond what is provable in Peano arithmetic must be *perceived as true* in terms of hidden higher-order concepts'. And if 'perceived as true' means something in the same ball-park as 'proved', then what we have here is not a tendentious and obscure claim about the *content* of what is proved, but a much more transparent claim about the *means* of proof. In other words, we have here a version of what I called Isaacson's thesis, but now a stronger one which says not merely that we need to appeal to ideas beyond those constitutive of our understanding of arithmetic in proving truths independent of PA, but says – more specifically – we'll need to appeal to higher-order/infinitary ideas to do the proofs.

However, I doubt whether we should buy this stronger claim about higher-order proofs (even though it is more plausible than the claim about hidden higher-order contents). For why should we suppose that, to get beyond PA, we have to make play with what are compellingly to be described as 'infinitary' ideas?

To take the obvious example, Isaacson himself notes that if we accept Π_1 reflection for PA, we are in a position to derive the Gödel sentence for PA. So let's ask: what does it take for a cognitively responsible agent to get herself into the position to rationally endorse Π_1 reflection?

Well, this issue is at the heart of a recent spat in *Mind* between Jeff Ketland and Neil Tennant. And when the wraps are off, their dispute in part seems to come down to this. Ketland supposes that if one is to rationally endorse Π_1 reflection for PA, one must already be in receipt of a good argument that PA is consistent. After all, he would say, you wouldn't continue to endorse $Prov(\ulcorner\bot\urcorner) \to \bot$ if it turned out that PA were inconsistent and hence $Prov(\ulcorner\bot\urcorner)$ were true. But, he contends, to get an explicit argument that PA *is* consistent *does* involve making play with infinitary ideas – for example, we might argue by adding a full Tarskian truth-theory (which is in effect equivalent to working with a certain second-order arithmetic with an infinitary comprehension principle), or by going via Gentzen's route of using induction up to ε_0 added to Primitive Recursive Arithmetic. Hence, on Ketland's line, arguing for the truth of the Gödel sentence for PA via considerations of Π_1 reflection does depend ultimately on the use of higher-order ideas to ground an essential presumption of consistency.

To which Tennant would, I think, riposte that being rationally confident in the arithmetic reliability of PA does *not* require having a prior guarantee that you can't be in for a nasty surprise. And this is a general epistemological point. For compare: being rationally confident in the general reliability of my eyes doesn't require – and indeed, surely can't require, on pain of rampant scepticism – a prior guarantee that they won't deceive me. I don't need a prior guarantee that I don't live in a demonic world where eyes deceive for it to be rational for me to rely on my eyes in beliefformation: similarly, you don't need a prior guarantee that it won't lead you astray for you to be rational in continuing in your default reliance on PA (nor do you need such a guarantee for it to be rationally in order explictly to endorse your continuing reliance, and be disposed to accept instances of the Π_1 reflection schema). So for Tennant, the cognitive responsible agent who endorses her own practices and accepts Π_1 reflection for PA does have a route to a derivation of the Gödel sentence for PA which needn't go via an argument for consistency which depends on the use of infinitary concepts (so, on Tennant's line, the stronger version of the Isaacson conjecture fails). Though of course, our agent does need to be able to stand back and reflect on her practices, and so indeed does need to deploy ideas which go beyond those which suffice to constitute her understanding of PA 'from the inside', as the core Isaacson conjecture insists.

Though I won't argue the point further here, I'm inclined to go with Tennant, and so would myself say that in some cases deriving arithmetical truths independent of PA *needn't* involve making play with infinitary ideas. There needn't be hidden higher-order concepts at work (in any natural substantive sense of 'higher-order'). But be that as it may. The point for present purposes is that the general epistemological issues at the root of the Ketland/Tennant dispute seem only obliquely related to the question of the truth of the core Isaacson conjecture. The way you jump on *that* dispute doesn't affect the truth of core conjecture as I've stated it.

3 Second-order arithmetics

If you were presented with Isaacson's thesis cold, taken out of the context of his paper, then your first thought might well be that it faces an immediate and very obvious challenge. For isn't there a very evident way of going beyond first-order PA while still keeping within the confines of what is given to us in our understanding of elementary arithmetic – for, you might say, can't we exploit the second-order ideas which are surely already involved in, for example, our informal understanding of induction? Recall the often quoted remark due to Georg Kreisel:

> A moment's reflection shows that the evidence of the first-order [induction] schema derives from the second-order [axiom].

Now, Isaacson himself says something by way of response to this sort of thought. But let me very quickly put a few points in my own way, which will enable me to mention a theorem – one which is familiar now but wasn't in 1987 – which nicely supports his thesis.

I suppose that this much may indeed be plausible in Kreisel's claim: we are prepared to give a blanket endorsement to all the instances of the first-order Schema because we accept the intuitive thought that if zero has some property, and if that property is passed down from one number to the next, then all numbers have that property, whatever property that is. But now let's ask: what properties are we implicitly quantifying over here?

Well, there is no obvious reason to suppose that our intuitive thought in fact aims to generalize over *more* than some 'natural' class of arithmetical properties (for example, those expressible by open wffs of L_A). In particular, there is no reason at all to suppose that our intuitive thought aims to use the sophisticated conception of properties whose extensions can be quite arbitrary subsets of the numbers.

Careful mathematicians are alert to exactly this point. For example, in his classic *Mathematics: Form and Function*, Saunders Mac Lane quite explicitly first presents the second-order principle of induction as an induction over the natural arithmetical properties that can be identified by first-order formulae. And he then distinguishes this intuitive 'induction over properties' (as he calls it) from the stronger 'induction over [arbitrary] sets'.

So grant Kreisel his thought that we stand prepared to accept first-order arithmetic with all the instances of its induction schema because we already find in effect committed to a second-order theory with a second-order induction axiom. That doesn't settle whether the intuitively compelling second-order theory should involve, in Mac Lane's terminology, induction over properties or induction over sets.

Now, the second option, with its infinitary notion of quantification over arbitrary sets of numbers, must *surely* go beyond what we are essentially committed to in our grasp of elementary arithmetic. But a second-order theory which quantifies only over those arithmetical properties which we can already express in L_A arguably *is* available to someone who grasps arithmetic plus some logical ideas. So here we have intimations of one natural sort of theory of arithmetic which in a sense goes beyond first-order PA while not calling on more sophisticated mathematical concepts. But however exactly we spell that out, such a theory cannot be stronger than the standard second-order arithmetic ACA_0 which has a comprehension principle which in effect restricts the domain of the second-order quantifiers to those properties expressible in L_A. And now here comes the theorem I said I'd mention: it is now well known that ACA_0 is *conservative over PA*. The second-order theory does have interesting properties of course: for a start, it exhibits massive speed-up over PA (there are PA theorems which have vastly shorter proofs in ACA_0). But it doesn't prove *more* L_A truths than PA.

In sum, if we go second-order but in a way that deploys ideas that arguably don't go beyond those available in elementary arithmetic plus logic, then we don't get any new truths of L_A over and above those which are already available in PA. Which is an observation thoroughly in harmony with Isaacson's thesis.

4 Induction and the ancestral

So far, so good. But now let's just dig a bit deeper. Even if we restrict the second-order induction scheme to run over kosher arithmetical properties that can be expressed in L_A, why should we accept it to be true? Why should we stand prepared to accept all the instances of the first-order induction schema?

We all know the answer: Suppose, for some property P, we are given (i) $P0$ and (ii) $\forall x(Px \to Px')$. By (ii), if 0 has P so does $0'$. Hence $0'$ has P. By (ii) again, if $0'$ has P so does $0''$. Hence $0''$ has P. Likewise, $0'''$ has P. And so on and so forth, through all the successors of 0. But the successors of 0 are the only natural numbers. So *all* natural numbers have property P, i.e. $\forall x Px$. QED.

Note, however, the crucial assumption here: *the successors of* 0 *are the only natural numbers.* That evidently underlies our acceptance of the induction axiom. And, I'd argue, is an idea which surely has to be available to anyone who grasps the full content of basic arithmetic.

A quick aside. Of course, I don't want to deny that there could be a level of arithmetical competence – a grasp of 'baby arithmetic' – which falls short of grasping e.g. the idea that every number has a successor. The baby arithmetician, presented with any number is happy to add one and arrive at the next, but – let's suppose – she never advances from the grasp of a general rule to a grasp of explicit quantifications over numbers, so can't frame the thought that every number has a successor. But the possibility of a level of conceptual competence with arithmetic which falls short of what is required for grasping PA is irrelevant for our purposes. For overall, we are considering whether someone who *has* got what it takes to understand PA can in fact get further with arithmetic using the same conceptual resources, or whether proving

L_A truths that run beyond PA requires bringing to bear new concepts.

To return to the argument. A full grasp of PA involves being able to handle quantifications over the numbers (to be sure, we can treat its purely Π_1 recursion axioms as schematic, but instances of the induction axiom can be as complex as you like and so can be essentially quantified). And understanding quantifications over the numbers involves understanding that the numbers are 0, the next number, the one after that, *and so on, without limit* – and understanding too that those are the only numbers. Which in effect is to grasp the thought is that every number stands in the ancestral of the successor relation to 0.

And indeed, grasp of the idea of an ancestral – or what comes to much the same, a grasp of the idea of transitive closure under iterable operations – is equally needed at another place in understanding PA. To appreciate the import of the induction schema, we also need to understand what counts as an open wff of L_A that can be plugged into the schema. And to understand what counts as a wff of L_A we need to understand the idea of starting from a class of atoms and then applying connectives or prefixing quantifiers, and then doing it again, and doing it again, and so on without limit.

Now, what does it take to grasp the operation of forming the ancestral? Well, we all know that the operation is *not* first-orderizable. Hence, in particular the property of being 0-or-one-of-its-successors, something grasped by someone who understands interpreted PA, is not itself definable in first-order PA. Equally, we all know that the ancestral *is* definable in a second-order framework (using full quantification over arbitrary subsets of the domain). Thus something has the property of being 0-or-one-of-its-successors if it belongs to every set of numbers which contains 0, and if it contains n contains n'.

Was I wrong, then, so rapidly to dismiss the idea that a grasp of full second-order quantification over arbitrary sets of numbers is built into our understanding of elementary arithmetic? Well, the fact that the notion of the ancestral *can* be defined in second-order terms doesn't mean it *is* an essentially higher-order notion. Here's a comparison. We can define identity in second-order terms by putting $x = y =_{\text{def}} \forall X(Xx \equiv Xy)$. But that certainly doesn't show that understanding identity presupposes an understanding quantification over arbitrary subsets of the domain. It *can't* show that, because understanding the idea of arbitrary sets of objects of some kind already requires an understanding of what makes for an object of that kind, and that in turn involves an understanding of what makes candidates the same or different objects, i.e. it already involves an understanding of identity.

Similarly for the notion of the ancestral. The child learns that her parents have parents, and that they have parents too. And in sepia tones, her great-grandparents have parents in their turn. And she learns that other children have parents and grandparents too. The penny drops: she realizes that in each case you can keep on going. The child gets the idea of an ancestor, i.e. the idea of being someone who turns up eventually as you trace people back through their parents. And *is* the idea of the ancestral essentially higher order? Well, it would seem flatly implausible to say that – because we *can* define the idea of an ancestor from the idea of a parent in second-order terms, and we *can't* do it first-order terms – the child who so easily grasps the concept *ancestor* must already be in cognitive command of the idea of arbitrary sets of people.

Similarly, the child learns to count and do arithmetic. The hundreds are followed by the thousands, and the tens of thousands, and the hundreds of thousands and the millions. Again the penny drops: you can keep on going. And she gets the idea of being a natural number, i.e. of turning up eventually as you keep on moving to the next number. Again, it seems badly over-shooting to say that to get this far, she has to be in cognitive command of the idea of arbitrary sets of numbers. (Here's another straw in the wind, deploying a thought that goes back to Bernays. We might well hesitate over the very idea of arbitrary infinite sets of numbers – sets which are supposedly perfectly determinate but are in the general case beyond any possibility of our specifying their members. We can readily make sense of the membership of a set of numbers being determined by possession of some characterizing property which gives a recipe for picking out the numbers; and we can readily make sense of the membership being merely stipulated (more or less arbitrarily). However, the first idea gives us infinite sets but not arbitrary ones; and the second idea may give us arbitrary sets (whose members share nothing but the gerrymandered property of having being selected for membership) but not infinite ones – unless we are prepared to conceive of a completed infinite series of arbitrary choices. Neither initial way of thinking of sets uncontentiously makes sense of the classical idea of arbitrary infinite sets of numbers. Now, in pointing this out, I'm not suggesting that we should be sceptical about the infinitary classical idea. But I do suggest that such worries about the infinitary classical idea, whatever their significance, surely do not carry over to become worries about the relatively homely idea of the ancestral. Which again indicates that the ancestral is not in the same infinitary boat.)

In sum, I've suggested that the child who moves from a grasp of a relation to a grasp of the ancestral of that relation doesn't thereby manifest an understanding of second-order quantification interpreted as quantification over arbitrary sets. It seems, in fact, that she has attained a distinct conceptual level here, something whose grasp requires going beyond a grasp of the fundamental logical constructions regimented in first-order logic, but which doesn't takes as far as an understanding of full second-order quantification.

We don't want to be multiplying conceptual primitives unnecessarily. But, if what I've been saying is right, the concept of the ancestral of a relation doesn't seem at all a bad candidate for being a fundamental logical idea. It is something (to repeat) that can't be explained in simpler terms, but which you can grasp without getting your head round more sophisticated ideas like the classical conception of quantification over arbitrary subsets of a domain. So let's run with this thought for a bit: what happens if we try to build the ancestral operator into a formal theory of arithmetic which is otherwise basically first order?

And you can immediately see that raising this question suggests a potential worry for Isaacson's thesis. For if fully understanding PA involves getting hold of the idea of the ancestral of the successor relation, then an 'ancestral arithmetic' which does have a built-in ancestral operator won't take us beyond what is already within the grasp of someone who understands PA. So it will matter crucially whether such a theory is or isn't conservative over plain PA. If it *isn't* then the thesis is in trouble, for we'll then have a clear counterexample to the claim that proving sentences of L_A which are undecided by PA must involve ideas that go beyond those which are constitutive of our understanding of basic arithmetic. So how do things pan out?

5 Ancestral arithmetic

Let's start thinking, then, how we could extend a first-order language with an ancestral-forming operator.

To begin, we need to expand our first-order logical vocabulary with an operator – we'll symbolize it with a star – which attaches to two-place expressions $\varphi(x, y)$: $\varphi^*(x, y)$ is to be interpreted as expressing the ancestral of the original relation expressed by $\varphi(x, y)$. We'll be forgivably careless about the syntactic details here.

Suppose that we augment the language L_A with such an operator. Abbreviate

$y = x'$ by Sxy. And first consider then what happens when you take the familiar axioms of PA but add in the new axiom

S. $\forall x(x = 0 \vee S^*0x)$

which says – as we want – that every number is zero or one of its successors.

Any interpretation of this expanded set of axioms which respects the now-to-be-kept-fixed logical meaning of the ancestral-forming operator evidently must be *slim*, that is to say, have just the 'zero' and its 'successors' in the domain. But, by a standard argument, the slim models of the remaining axioms of PA are all isomorphic. So our augmented theory is categorical.

Let's define, then, the appropriate semantic entailment relation $\Gamma \vDash^* \varphi$, which obtains if every interpretation which makes all of Γ true makes φ true – where we are now generalizing just over interpretations which give the star operator its intended meaning (and otherwise treats the logical vocabulary standardly). Then, because of categoricity, our expanded axioms semantically entail any true sentence of the expanded language. So in particular, if φ is a true sentence of the unexpanded language L_A, then PA + S $\vDash^* \varphi$.

There can't, however, be a complete axiomatization of this 'ancestral arithmetic'. That follows by a simple compactness argument (for $\vDash^*$ is evidently non-compact). Or we can argue as follows. Take any formally axiomatized theory which is sound for $\vDash^*$ and which includes PA. The incompleteness theorem then applies (assuming consistency), so there will be an unprovable-yet-true L_A sentence G^* for this theory. But since the L_A truths are all semantically entailed by ancestral arithmetic, this means that there is an unprovable-yet-semantically-entailed sentence, so the deductive system can't be complete.

So far, this should all be more or less familiar from the last chapter of Shapiro's book on second-order logic, where he mentions what he calls ancestral logic. But now note that the unaxiomatizability of $\vDash^*$ of course doesn't rule out *partial* axiomatizations of ancestral arithmetic. Shapiro doesn't explore the options here: but R. M. Martin and John Myhill did, in a series of articles in *JSL* going back to 1943.

Suppose we write $H(\psi, \varphi)$ as short for $\forall x \forall y((\psi(x) \wedge \varphi(x, y)) \rightarrow \psi(y))$ – which says that the property expressed by ψ is *hereditary* with respect to the relation φ (i.e. is passed down a chain linked by φ). Then Myhill's proposed axioms for the star operator are tantamount to the following schematic rules:

From $\varphi(a, b)$ infer $\varphi^*(a, b)$.
From $\varphi^*(a, b)$, $\varphi(b, c)$ infer $\varphi^*(a, c)$.
From $H(\psi, \varphi)$, infer $H(\psi, \varphi^*)$.

These first two rules serve like 'introduction' rules for the star operator; and the third rule is easily seen to be equivalent to the following 'elimination' rule:

From $\varphi^*(a, b)$ infer $\{\psi(a) \wedge H(\psi, \varphi)\} \rightarrow \psi(b)$.

What are in effect the same rules seem to have been rediscovered by Avron Arnon in a 2002 paper, and taken over from there by Richard Heck in a forthcoming piece. (Of course, when I call those 'introduction' and 'elimination' rules, I don't mean to imply that they are rules that suffice to fix the content of the star operator – they can't, given that the rules aren't complete.)

Our 'elimination' rule is evidently a kind of generalized inductive principle which says that given that b is a φ descendant of a then if a has some property which is passed down from one thing to another if they are φ related, then b has that property too. And it is easy to see that – in the presence of these rules and the axiom S – our generalized induction principle entails all instances of the ordinary first-order induction schema.

So let's now consider the formal system PA* which extends Robinson Arithmetic Q by adding our new axiom (S) plus Myhill's rules for handling the ancestral operator (and so also extends PA). The obvious next question is: what is the deductive power of this system with respect to sentences of L_A? Is this another case like the canonical partial axiomatization of full second-order arithmetic, PA_2 (a.k.a. Z_2), where we also only have a partial axiomatization of the relevant semantic consequence relation, though axiomatized PA_2 can prove more L_A sentences than PA (e.g. it can prove Con_{PA}). Or is PA* another extension of PA like ACA_0, which is conservative over PA?

If the first case holds, then – as I said – we'd surely have a refutation of Isaacson's thesis. For PA* is arguably well within the conceptual reach of someone who has fully understood the inductive schema of PA; and so, if we could use it to prove new sentences of basic arithmetic not provable in PA, and Isaacson's thesis would fall.

6 PA* is conservative over PA

But in fact, as you'll probably have guessed, things fall out Isaacson's way:

PA* is conservative over PA.

In other words, for any L_A sentence ψ such that $\mathrm{PA}^* \vdash \psi$, it is already the case that $PA \vdash \psi$.

As far as I know, this result is not in the literature: it may be folklore. But anyway, here's a ...

Proof Recall that we can express facts about sequences of numbers in PA by using a β-function where for any finite sequence $k_0, k_1, \ldots, k_x$ there is some c such that for all $u \leq x$, $\beta(c, u) = k_u$. So suppose R is some relation. Then its ancestral R^* is such that

A. R^*ab is true just so long as, for some x, there is a sequence of numbers $k_0, k_1, \ldots, k_x$ such that: $k_0 = a$, and if $u < x$ then $Rk_u k_{Su}$, and $k_x = b$.

Using a two-place β-function, that means

B. R^*ab is true iff for some x, there is a c such that: $\beta(c, 0) = a$, and if $u < x$ then $R(\beta(c, u), \beta(c, Su))$, and $\beta(c, x) = b$.

So consider the following definition:

C. $\varphi^{**}(a, b) =_{\mathrm{def}} \exists x \exists c\{B(c, 0, a) \wedge$
$(\forall u < x)\exists v \exists w\{(B(c, u, v) \wedge B(c, Su, w)) \wedge \varphi(v, w)\} \wedge$
$B(c, x, b)\}$

where φ expresses some relation R, and B is a formal three-place wff of L_A which captures a β-function. Then φ^{**} expresses the ancestral R^* of the relation R expressed by φ.

It is now easy to check that the Myhill inference rules for starred wffs apply equally to our defined double-starred construct φ^{**} in PA (that's because the moves are then evidently valid semantic entailments within PA, and the theory's deductive system is complete), and the double-starred analogue of axiom (S) is also a theorem of PA. So corresponding to any proof involving starred wffs in PA* there is an exactly parallel proof in plain PA involving double-starred wffs. Hence any proof of a pure (star-less) L_A wff in PA* has a parallel proof in plain PA. Which is what we needed to show.

7 Other arithmetics?

But is PA* (or some extension to it, which can similarly be proved to be conservative over PA) as far as we can get solely by recourse to ideas already involved in the understanding of PA? Well, the literature on ancestral arithmetics is surprisingly thin. To be sure, there is an extensive treatment of logics with a transitive closure operator in the literature on finite model theory; but that is all model-theoretic, and doesn't seem to discuss deductive systems. There's the paper I mentioned by Avron Arnon which puts it down as an open question what other rules for the star operator should be added 'in order to make the system "complete" in some reasonable sense'. And there is a passing discussion in Heck's forthcoming paper, but Heck's interest is in stronger systems which draw on essentially richer conceptual resources. So I don't know of other essentially stronger competitors to PA* for being formal regimentations of what is available to us through our grasp of basic arithmetic and the ancestral in particular. So things are now looking rather good for Isaacson's thesis!

8 Conclusion

Let's summarize. Our everyday understanding of basic arithmetic pins down a unique structure for the natural numbers, at least up to isomorphism. Hence our grasp of basic arithmetic involves more than is captured by first-order PA. But what more? It seems enough that we have the idea that the natural numbers are zero and its successors and nothing else. And getting our head round this idea, we suggested, involves the general idea of the ancestral: the numbers are what stand in the ancestral of the successor relation to zero.

Now, the ancestral of a relation can be defined in second-order terms, but it seems overkill to suppose that our understanding of ordinary school arithmetic is essentially second-order. Why not treat the operation that takes a relation to its ancestral to be a (relatively unmysterious, even though not purely first-order) logical primitive? If we do, we can construct a formalized theory PA* which naturally extends PA in a way that arguably still reflects our everyday understanding of arithmetic.

Since PA* is still properly axiomatized, however, we know that it will be incomplete (assuming it is consistent). But we might have suspected that it would at least have proved more arithmetic truths than PA: but not so. PA* is deductively conservative over PA for L_A sentences: so we can't in fact use this expanded theory to deduce new truths of basic arithmetic that are left unsettled by PA.

Hence, to put it the other way around, it seems that if we *are* to come up with proofs of truths unsettled by PA, then we'll have to deploy premisses and/or logical apparatus that go beyond PA*. Which arguably implies that we'll need to invoke ideas which go beyond those essential to our ordinary understanding of basic arithmetic. Why? Because – apart from a grasp of the recursive definitions of addition and multiplication – the idea that all the numbers can be reached from zero by repeatedly adding one, i.e. the idea that all numbers are related to zero by the ancestral of the successor relation, does seem to be at the limit of what is *necessary* to the grasp of the purely arithmetic concepts which underlie PA and – to speak more vaguely – underlie our grasp of basic pure arithmetic. And if that's right, Isaacson's thesis (at least in the version I stated it) seems secure.