Mathematics as Language

From *Benacceraf and his critics*, Adam Morton and Stephen Stich, eds., Blackwell 1996.

There are surprisingly many words in the average mathematical book or article. Much of most mathematical works consists of plain English (or French, or Japanese): setting up, explaining, and drawing consequences from characteristically mathematical-looking symbols. And there do not have to be any symbols at all; interesting mathematics can be done in purely natural language. When mathematical symbols do enter they occur in linguistic contexts defined to a large extent by the more ordinary language around them.

But what occurs in a linguistic context is language. And undoubtedly there is a sense in which forms of expression that exist only in pure and applied mathematics are linguistic, just as linguistic as the language of gossip and comedy. There are also ways in which the characteristically mathematical parts of mathematics are very different from the rest of language. The aim of this paper is to discuss the likeness and unlikeness of mathematics to the rest of language, purely for its own sake. There are connections with standard issues in the philosophy of language, and I shall begin to draw them. The connections are rather Benacerrafian.

Old arguments between Platonists and formalists turn on the 'is it language' axis? 'In mathematics you conjoin and disjoin, and what you say is true or false' say archetypal Platonists, 'and if it is true or false it is true or false of something, and that something sure isn't the world around us'. 'The syntax is not made for talking about anything', say archetypal formalists 'it is made for drawing conclusions, and that's what it's all about: whether things follow and not whether thecae arc true.' There is obviously a *prima facie* case on both sides. Mathematics is indeed language; it is asserted and denied, used to make questions and commands. And mathematics is indeed not typical language; the syntax is different, and something *feels* very different about the pragmatics. But does this make a semantical difference?

Modes of Language

Begin in order to set up a contrast, with a simplistic picture of one non-mathematical linguistic mode. Narrative mode, the way we speak when we tell stories. I take this to be one of the very simplest modes of languagc use. That is, it represents a particularly natural way in which speakers intend thoughts to be organized in the minds of hearers. So it is a matter of the presentation of the facts rather than the facts presented. Narrative presentation works by focusing on a (small) number of protagonists at a given moment in time, then it unwinds time forwards successively attributing properties to these protagonists. So the core facts attributed and thus communicated are of the form A is true of p at time t. (Or during interval i. Or relationally: R holds between p_1 and p_2 at time t etc.) The result is a kind of a database. The facts are presented as structured in a particular way. The normal presentation of a story invites such a structuring, since the events are narrated in the order in which they are represented as happening, and the syntax of the language - proper names, anaphora, tenses, modals - encourages the hearer to organise The facts presented in this way.

An experiment I have tried illustrates this: present e.g. *Little Red Riding Hood* to a class, written out in inverse chronological order. Or in terms of what is true of the wolf, what is true of the grandmother, what is true of the basket, etc., with the identity of these characters disguised. It takes the class a long time to answer simple narrative questions expressed in terms of time and protagonist. And it is only after being
prompted with such questions that they realise that they have been given
a familiar story.

The idea of a mode of language can be put a bit more formally. Consider a language consisting of the sentences which are actually produced (or, less naively their interpreted deep structures) augmented with a range of sentences of richer and more explicit logical form. Thus we might have the natural language sentence 'She filled the basket with cakes, and then she left her mother's house.' The richer language might include sentences such as 'At time t Little Red Riding Hood filled the basket with cakes and at some time t + ☐ where ☐ is small Littlc Red Riding Hood left the house which belonged to the mother of Little Red Riding Hood.' Sentences in the augmented language would follow from sets of produced sentences plus background information. (If the natural language is expressive enough it may be its own augmented language.) Now a mode of language distinguishes a subset of sentences of the augmented language as targets of interpretation. In the case of narrative language the subset might be that of sentences of the form 'P does A at t'. Then speakers intend that hearers will deduce particular sentences of this distinguished set from what they say, and hearers will take their task of comprehension to have been accomplished when they deduce a sentence of the distinguished set which is relevant to the narrative in progress.

Thus there is a program for hearers of narrative to follow. It is (to a
first approximation):

Hear s

```
Assign s an interpreted logical form, call it F.
Set F = a
FOR v = immediate consequence of F
IF v is of the form 'P does A at t' and is relevant to the story STOP
ELSE set F = v
```

The important thing to notice about this program is that it tells hearers when to stop deducing consequences. It tells them when they can take themselves to have understood.[1] (There will be several such programs in this paper. Their terminology should be self explanatory.)

There are other modes of language. A descriptive mode is also common. One of many descriptive modes is taxonomic. The discourse sets up and then presupposes a data-organizing system in which objects are filed under a number of predicates or relations. Each utterance either adds a new object under an existing heading or expands the list of headings. ('Now I'm going to tell you the capitals of European countries. The capital of Scotland is Edinburgh, and the capital of Wales is Cardiff.
France's capital is Paris, and for Germany it is Bonn although sometimes they pretend it is Berlin. A country's cultural capital may not be the same as its political one. Glasgow, for example....') The distinguished subset then consists of sentences of the form '*o* has P' or '*o* bears R to *p*', where P or R are from a list of predicates specific to the discourse in question.

There are obviously many standard descriptive modes, just as there are many standard narrative modes, and speakers obviously make many variations on them. But in each of them there is a rough uniformity of the kind of organization into which speakers expect hearers to fit the facts. And in each of them the program hearers follow when understanding sentences tells them when they can stop.

When someone says something, in either narrative or descriptive mode, each hearer makes sense of it as something said in that context in that mode. The mutual knowledge that that is the relevant mode is vital to the hearer's task. If one of us thinks we're telling a story and the other thinks we're doing taxonomy then misunderstanding is inevitable. Indeed some individual words are ambiguous in ways that are resolved, at least partially, by knowledge of the mode. For example 'then' indicates passage of time in narrative but the drawing of a consequence in argument. The most important thing for my purposes is that the knowledge of the mode tells the hearer when to stop interpreting the utterance. Thinking of interpretation simply as drawing consequences (in a possibly richer language), it tells the hearer when to stop deducing. Given an utterancc in a context in a mode the hearer deduces consequences from (a disambiguation of) it until she arrives at something in the target form for that mode. Then she may have to answer questions or carry out some other task ('who got to grandmother's house first, Little RRH or the wolf', 'put the scissors in the top left hand drawer'.) The accumulation of consequences in the target mode should be the right sort of information for carrying out the task. If not, then that was not the right linguistic mode for a conversation directed at that kind of task.

I shall assume, then, that speakers and hearers usually know what linguistic mode they are in, and that they find communication much harder when they do not have a manageable mode. If there is no mode there is no STOP instruction to their interpretation, and the comprehension task is in danger of being ungrounded. Given a linguistic mode, hearers can extract information from what is said in terms of a specific data structure. (You can think of this as a language of thought, mentalese, if you like; but you do not have to.) The assumption is clearly not that there is a single canonical structure imposed on all facts understood from all speech. It is that for every conversation there is such a level. And it is fixed for that conversation. I take this to be a characteristic of narrative and descriptive conversation. I don't expect it to be true of all kinds of non-mathematical language. Not for poetry, for example.

A speaker of a language has mastered conventions which define functions from conversational context to meaning, so that in a sufficiently well-dcfined situation there will be a unique best interprctation of the speaker's words. The important feature of this for our purposes is the disposability of the actual words: the translation goes 'all the way down' to whatever form of representation is right for that discourse. Nothing is kept at any intermediate symbolic mode and once the hearer has made the translation into mentalese she can throw away the original sentence.

Thus the familiar fact that a few moments after you have digested a bit of narrative you can reproduce very few of the original words; if asked what was said you paraphrase, retranslate upwards.

It follows that when people are attending to language in narrative or descriptive mode they operate with a particular selection of the linguistic information that has been provided. There is the immediate linguistic input, the very words just uttered and the syntax which surrounds them. And there is the accumulated information: the story so far or the partially completed taxonomy, translated into a form quite distant fromt the original language. Very little else is available. No record of syntax of past utterances, no record of the actual words uttered more than a moment ago.

Mathematics as Language

Are there mathematical modes of language, which are different from narrative and descriptive modes? Some of the common characterizations of mathematics as language do not seem right. In particular, it is not true that mathematical assertions are less dependent on context than assertions in more normal discourse. Mathematical language is in fact thoroughly context-dependent. For example there is lexical ambiguity which has to be resolved contextually. Epsilon means membership in set theory and a variable over very small differences in calculus. *e* can be Euler's constant or the identity element of a group. (Lexical ambiguity does tend to occur and bc rcsolved in a somewhat difrerent way. The core ambiguity is 'what part of mathematics is this', and the resolution is usually immediate once this is specified. But this accompanies a tendency to a kind of ambiguity that is less common outside mathematics: the different meanings a symbol can have may of completely different syntactical categories. Nouns and verbs are often the same in English, but logical connectives and predicate terms generally have no phonetic overlap. But in mathematics an occurrence of $\supset$, for example, can be a connective, if, or a relation, contains, depending on the context.)[2]

Moreover in much of mathematics as in natural language the scopes of quantifiers have to be disambiguated by context. An equation like $x = \square t$ may occur on a line by itself, and the experienced mathematical reader knows that $\square$ is often used as a 'parameter' in relation to the use of x and t as variables, so that whatever implicit quantifier binds $\square$ is likely to have wider scope than the quantifier binding x and t. But for the relative scopes of x and t, and the nature of the quantifiers binding them one will have to go to the surrounding text. Even then the final assignment of quantifiers and scopes will often depend on what makes the most sense in the wider context, just as it usually does in non-mathematical language. In most mathematical language the purely mathematical symbols cannot be understood without the surrounding 'normal' prose which sets up the topic, disambiguates the terms, determines the scopes of the quantifiers, and so on.

So there is a lot of context dependence. But it is typically a different pattern of dependence on different features of the context. I shall illustrate this with three examples.

Pons asinorum

One reason why people's heads swim when they work through mathematics, why they feel dizzy and confused and wish they were back in the familiar land of narrative, is the way in which name-introducing contexts are at once longer and shorter than in more familiar discourse. In archetypal narrative a name is introduced for a protagonist at some stage in the story, and then that name is linked to that protagonist - to that heading in the narrative database - for the whole of the story. The number of names is usually fairly limited, but the length of time we keep them in mind can be very long, hundreds of pages. On the other hand in typical mathematics symbols are introduced and given designations for contexts that usually are fairly short. A term will be given a designation for the duration of a proof, which will rarely occupy more than a page. Then a new set of terms will be given designations. And many of these designations will not be simple abbreviations of descriptions involving the previously introduced terms. So the effect is that there is more assigning and re-assigning of designations, but the length of time a term keeps its designation is *less*. As a result if you try to understand mathematics in terms of a basically narrative way, opening a file on every individual you are introduced to and then updating the file as the story progresses, you find things very tiresome and confusing. You seem to spend all your time opening and closing files.

A good example of this is the *pons asinorum*. This is the traditional name of the first difficult theorem in Euclid, the asses' bridge. The donkeys can't get across it. They stop at that point, because it is the first point at which some mathematical sense is required. There's no change in symbolism or subject matter at this point. And the content of the theorem is no more recondite than what precedes it. We just run into a proof that baffles the innocent. Why might that be?

The theorem (Euclid's *Elements* Book 1 prop 5) asserts that the angles at the base of an isosceles triangle are equal. Equal sides, equal angles. The proof assumes a previous theorem: if two triangles have two sides and the angle between them equal then they are congruent. It then runs as follows. The isosceles triangle ABC is extended as in the diagram below. Then by the previous theorem the triangles AEB and ADC are congruent. So the angles AEB and ADC are equal, and the lines EB and DC are equal. Then by the previous theorem again the triangles CEB and CDB are congruent. So the angles BCE and CBD are equal. But the angles ABC and ACB are equal to a straight line (180°) minus CBD and BCE respectively, so they are equal.
Q.E.D.

Diagram for the *pons ansinorium*

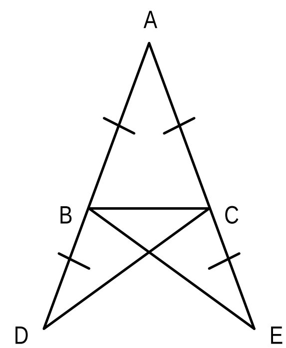

Why is this proof hard to follow? (I find it confusing. If I have to reproduce it I have to concentrate.) The reason is pretty clear. To understand it you have to fix in your mind a number of triangles and angles, and then you have to go through some reasoning in terms of them. And you have to fix them in your mind in the correct orientation in order to see the right congruences, so that just imagining the diagram vividly will not do the trick. At several points in the proof you have to go back from the reasoning to the diagram, to define a new term. The result is a conflict between that you have to keep in mind throughout - the diagram and the purpose of the proof - and the frequently changing definitions. You zig-zag back and forth between syntax and image, and in the process it is easy to loose your grip.

Maxwell's equations

Imagine a very pure mathematician. She has learned a lot of analysis, through multigrade calculus, in a very rigorous foundations-respecting way. But she has never been exposed to the corruptions of any applied mathematics. She sees Maxwell's equations for the first time. She has no trouble understanding the notation, ∇ etc, and sets about thinking out what the equations say about the electrical and magnetic fields. If understanding mathematical language were like understanding narrative language what she should do is translate the del notation first into combinations of partial derivatives, and then translate these down into assertions about limits in four dimensions. Probably she should translate even these down into horrendously complex quantifications over integers and sets of integers, spelling out the definitions of the limits. (Lim fx, $x \rightarrow a$, = b iff $\forall\varepsilon\exists\delta\forall x$ ($|a-x| < \delta$) $\Rightarrow$ $|f(x) - b| < \varepsilon$). Now imagine that definition iterated with variations to give the content of Maxwell's equations.) But that is obviously not what a mathematically sophisticated person would do. It is obviously not what anything much like a human being could do, without blowing a mental fuse. She would translate down to partial

derivatives, although someone of equal sophistication with experience of applied analysis might not have to do this in order to get the intended behaviour of the fields, and then try to deduce facts about
the fields. Only if this failed, or if some paradox seemed to be in the offing, would she go down a level. It would take some extraordinary reason to induce her to make even a partial translation into the very bottom level of assertions about limits.[3] (Limits are normally off limits.)

The contrast with narrative language here is not just that the mathematician would not translate assertions down into the most basic terms. For in many non-mathematical contexts, also, we do not. And in characterizing narrative and descriptive modes of language I was careful to say only that for each conversation in each mode there is an appropriate level of representation. The contrast is rather that the appropriate level of representation is not lixed by the linguistic context. In understanding Maxwell's equations the mathematician would register the *possibilities*
of decomposing the senses into different levels of meaning, of drawing consequences to different depths, but would take them apart only very superficially until she saw what the task in hand required. There is no fixed STOP instruction.

Mental arithmetic

Andrew has just flown from Bristol to Marseilles, via Gatwick airport and is describing the arduous journey, involving fog bound motorways, missed and delayed flights, and crashed taxis. He says that it took 25 hours to travel 450 miles. Belinda replies that on the same day she went by bike from Totterdown to Keynsham (both near Bristol), 5 miles, in 15 minutes. (She's a good cyclist!) She goes on to say 'so actually I went faster than you'. Andrew is amazed and intimidated. He imagines that she calculated his average speed and then hers, to see that hers was greater. What feats of mental arithmetic! But actually the mathematical spirit is opposed to the calculating spirit. What she did was think: he took one hundred times as long as me and went less than one hundred times as far. (So since v = (d/t my v is bigger than his v.)

This may seem quite a natural way of thinking to some people. But it wouldn't be to most. For what it requires is for Belinda to think of Andrew's speed and hers in a way that leaves them unreduced to basic terms. His is 450/250, but she doesn't divide it out. She stores it in a not-quite basic form and compares it with a similarly unreduced form of her speed, 5/0.25. And that's exactly what the mathematician in the previous example did with Maxwell's equations. There is no fixed STOP point; you have to draw consequences to whatever depth the problem requires.

This example brings out another characteristic of mathematical language. Our heroine doesn't calculate. She keeps those quotients undivided. Suppose she had been incapable of calculating, or had a friend, or an idiot savant twin sister, or a separate specialized lump of her brain that did such things. (A co-processor!) It would all work just the same way. It is as if a properly mathematical way of talking and thinking allows for references to (inputs from or demonstrative reference to) calculating capacities or geometrical intuitions that are not themselves properly linguistic. (Like a word processing program that allows you to say: picture referred to by such and such a file goes in the text at such and such a place.) And when that is the case it is no wonder that the linguistic context is not tracked in terms of the ultimate decomposition of the utterances: they are not in a form that our linguistic capacities recognize They aren't propositional.

The following three examples are meant to bring into focus observations every mathematically literate person must have made about the difference between understanding mathematics and understanding other kinds of language. Three apparent differences stand out:

1. It can seem that in mathematics we have to be ready to load our minds with an impossible accumulation of detail: all the terms and their definitions, the full decomposition of Maxwell's equations into their ultimate definienda, the results of all the relevant calculations.
2. It can seem that in mathematics we have to keep in mind the syntax of a whole long passage while trying to understanding the latest thing heard or read. We cannot just keep in mind something as simple as the story so far. You have to remember the detailed definitions of all the terms: you have to he ready to go back to Maxwell's equations as written to extract a more fundamental content from them; you have to keep all the terms of a bit of applied arithmetic in mind so that you can enter them into calculation if need be.
3 It can seem that in mathematics we have to refer to capacities which are not themselves cognitive or propositional, so that a linguistic form has to somehow include something inherently non-linguistic. To keep in mind the definitions behind a geometrical proof you have to fix them with a spatial intuition; to understand Maxwell's equations you have to bring to mind a capacity to re-interpret them on several levels; to do applied arithmetic you have to be able to manipulate within a propositional context a calculating skill.

The first two of these examples are just the way it seems: you do not really preserve all the syntax in mind, and you do not really think in terms of a superhumanly detailed representation. But the third is real. In thinklng mathematically you manipulate allusions to skills, capacities, or Intuitions, as if they were proper propositional contents. And this explains the appearance of the other two. The reason that a mathematical expression seems to have consequences and interpretations that go beyond any manageable level of semantical representation, is that its content is not really any of these representations, instead it consists in a capacity to expand to whichever one of them is appropriate to the problem posed. The reason that it may seem necessary to keep a long stretch of syntax in mind is that if you have drawn conclusions to any given depth (or of any given character) you may find that the problem requires you to draw conclusions to a different depth (or character), and this may require going back to the original symbols and starting again. So it can seem either that there is no semantics, that it is all syntax, or that there is an elusive and remote semantical content, if only you were capable of getting it into your mind. But in fact the content is neither grounded at level zero, syntax, nor at a determinate remote level. It is quite ungrounded.

Maps and Phone Boxes

Here's a metaphor. Speaking and thinking are like reading a map in a phone box. The map is too big to unfold in the box. So you need a way of getting at parts of it at a time. The narrativc strategy is to roll the map up, and unscroll it from one end to the other. The descriptive strategy is to do everything from the index. The mathematical strategy is to unfold one bit of it at a time, perhaps to unfold bits of several different maps at dlfferent scales, and to have a little notepad giving references to which bits of whlch map are to be consulted in which order. That's a less straightforward way to solve the problem, but becomes more natural when the maps arc very big, relative to the box, or numerous, or simply when the task to be solved cannot be done by unrolling or checking against the index.

Another metaphor. Consider demonstratives 'this' and 'that' referring to pictures which might accompany a text. So we can say 'if you have a picture like this and another like that then you can put them together side by side to make a bigger one, this and that together'. (A body of assertions of this form might generate their own notation, perhaps things like '□ & □ → □□'.) Someone might marvel that people who understood such assertions could handle really complex thoughts in which the pictures are intricate Breugels and Boschs, or which have pictures inside pictures inside pictures. Someone else might wonder what kinds of things 'this one' and 'that one' (or □) are. Explanations might be given in terms of powerful capacities to think in terms of pictures or to handle syntactical forms so effectively that the pictures can be ignored. But in fact to understand these idioms we would need only a capacity to refer to pictures (imagined or presented or however), a capacity to stick pictures together (as images, or physically, or however), and the capacity not to get distracted into looking too hard at the pictures. ('A capacity' to refer and stick together, note, not 'the capacity': new ways of linking the idiom to pictures and corresponding new ways of concatenating pictures might always be found.) Mathematics is like this picture-referring language. The pictures are an endless variety of capacities - counting, calculation, geometrical intuition, formal symbol manipulation - which get invoked and combined according to the demands of the task.

The Mathematical Mode

What I ought to provide at this point would be a description of some mathematical mode of language, for example to give a program like the one I gave above for understanding narrative language, which would glve instructions ror undcrstanding some significant part of mathematics. That would be great to have, it would be a big step forward in pedagogy. I can only give a tiny sample now, a taste of what it might be possible to do. My language will consist of the summation sign $\sum$ plus simple arithmetic. $\sum_a^b (x_i)$ where the x_i are integers (i in somc index set I) will be a determinate integer. Calculating this integer will usually be a lot of work, and for many purposes it is a bad idea to make the calculation. How then should we understand the formula? As the possibility of calculating. Consider first a little proof about $\sum$.

Claim For any function F of integers, define $\delta_n F$ as $F(n) - F(n - 1)$.
Let $G(n)$ be 2^n.
Then $\delta_{100}\sum_0^{100} G(s) = G(100)$.
Proof $\sum_0^{100} G(s) = 1 + 2 + . . . + 2^{100}$.
Therefore $\delta\sum_0^{100} G(s) = (1 + 2 + . . . + 2^{99} + 2^{100} - (1 + 2 + . . . + 2^{99})) = 2^{100}$

Instructions for understanding the proof

1. Set $\sum_o^n G(s)$ as process of adding up all the powers of 2 from 0 to 100 (but don't add them).
2. Set $\sum_0^{n-1} G(s)$ as process of adding up all the powers of 2 from 0 to 99 (but don't add them).
3. Set $\delta\sum$ as result of add up powers of 2 from 0 to 100 minus result of adding up powers of 2 from 0 to 99. But don't subtract them.
4. Set calculation for $\sum_0^{99} G(s)$ going whilc simultaneously setting calculation for $\sum_0^{99} G(s)$ going and subtracting partial sums.
5. Note result.

Comments

Re 1 The reader might know the formula for the sum of a geometrical series, and apply it here, but that would be a waste of effort. The instructions must suspend any evaluation of the formula until it is clear what kind of a calculation is worthwhile.

Re 3 This sets up a complex calculation. But it is essential not to perform the calculation in a literal way

Re 4 This presupposes some algebra, that a difference of sums is the sum of the differences.

Re 5 And why is the result worth getting? Because it is the fundamental theorem of calculus in disguise. For any function C of integers $\delta_n\sum_0^n G(s) = G(n)$, and the proof of this is really no harder. But it can't be read by my instructions.

A comprehension program for a $\sum$-formula *s*.

Interpret s, against a background of instructions I about how much and how far to calculate, as σ
Set $\tau = \sigma$
FOR *v* = immediate consequence of τ
IF *v* is of the form 'description of a calculation procedure' AND I applies to *v* THEN perform the calculation in accordance with I and STOP
ELSE set $\tau = v$

Note several things about this program:

- There is an analogue to the story-so-far, namely the background of calculations and instructions already set up.
- Calculations are performed only to the extent required by previously interpreted instructions.
- There are STOP points grounding the interpretation. But they consist not in arithmetical facts but in the determination of instructions for calculations procedures.
- The 'instructions' for calculations have been left unspecified. This is the aspect it would be most important to represent explicitly. It represents a central feature of mathematical thinking, appearing in a different form in, for example, the links between the definitions of angles and lincs in a geometrical construction and the diagram or geometrical intuition. (See the 'labels' in the counting program of the next section.)

What Mathematics Might Be

I have been arguing that mathematical language is inherently unlike narrative or descriptive language. It isn't the sort of idiom we would use to tell a story or describe anything. Instead it is the kind of idiom that we would use to apply a large variety of non-cognitive skills (calculation, visual/spatial intuition, skills of symbol manipulation) to an infinite variety of objects. We count, measure, and subdivide, and mathematics tells us how to do this. Mathematical assertions tell us how to count, measure, and subdivide. They are knowledge how not knowledge that.

Consider one very simple such skill, counting. Suppose you have some things of two kinds and you want to answer the question 'are there exactly as many As as Bs?' A natural procedure for finding the answer is as follows:

> (Background: As and Bs alrcady found and 'labels' already used.)
> Search for an A and a B distinct from those already found.
> IF there are no such A and B THEN STOP at YES
> IF there is such an A and no such B THEN STOP at No
> IF there is such an A and such a B THEN give them a label *l* distinct from all previous labels.
> REPEAT

Notice the labels *l*. They can bc exploited to answer other questions, such as 'If I take away one A will there then be exactly as many As as Bs?' The basic procedure of assigning labels to objects can be adapted for many counting tasks. But what are these labels? Anything that can be used for the purpose. They should be attachable to things and you shouldn't run out. Attachment is easiest in the mind, and not running out is a matter of whether you are counting big collections or small. If the labels had their home in the mind then you could pretend that there was a single set of lhem you used for all counting lasks. Even then it would not follow that there was a unique such set that would do the job. In fact, this would be obviously false. The labels are the ancestors of the integers; their lack of unique identities, based on the fact that they are not objects but ways in which we control and interrogate calculating processes, is intrinsic to the role they play in our thought.[4]

It does not follow from this that there could not be a world which arithmetic, by a sort of miracle, happened to describe. What does follow is that to interpret mathematical language as describing such a world is not to understand it the way we do in normal mathematical practice. For one thing, if the most basic mathematical facts were facts about the relations between numbers, sets and the like, then the best comprehension strategy would be to derive consequences from mathematical assertions until you arrive at mathematical bedrock. But as the examples above should show, this would be a disastrous strategy for any real human mathematician to use. The STOP points would come in the wrong places for following proofs and deriving applicable results. If you take the fundamental arithmetic facts - those in terms of which utterances and tasks have to be understood - to be the basic unquantified numerical sentences, then this set, although recursive, is very little use for deriving the quantified sentenccs that are needed to make real calculations.

In actual mathematical life we take ourselves to have understood a mathematical assertion when we can link it successfully with a structured intuitive or symbol manipulating skill. That is where is the real STOP points come. (And it is usually when we try to unpack the skills and interpret them as descriptions of objects that the Platonic objects appear.) The skills are such that - under favourable circumstances and covering our tracks to some extent - we get the same answers whenever we employ them. If a skill is not in this way reliable and replicable it does not get into the mathematical repertoire. It may seem a marvel that there are any such skills. It is. But it is a marvel that has to be faced by any analysis that pays attention not just to mathematical truth but the reality of mathematical practice.[5]

---

[1] The idea that hearers process speech until they arrive at something sufficiently relevant to the assumed purposes of the conversation comes from Dan Sperber and Deirdre Wilson *Relevance* (Basil Blackwell, 1986).

[2] Mathematical language does not seem to have much of the kind of context dependence that results in vagueness. But there are analogs of metaphors: for example $2^{1/2}$.

[3] On the other hand there is the famous story of von Neumann solving the trick question about the bee and the motorcycle by instantly summing an infinite series. But what I say is meant to apply only to mortals.

[4] There is an obvious connection with Paul Benacerraf's 'What numbers could not be' (*Philosophical Review*, Vol. 74,1965) here. But note a more indirect connection with his 'Mathematical Truth' (*Journal of Philosophy*, Vol. 74 1973) also, in that there is a tension between traditional truth conditions and obvious knowledge conditions (which for very simple assertions amount to calculation.) To see the conceptual origins of number words in the internal economy of activities like counting is to make a more radical 'structuralist' position than writers like Penelope Maddy, for example in chapters 3 and 5 of *Realism in Mathematics* (Oxford University Press, 1990), and Michael Resnick, for example in 'Mathematics as a science of Patterns: Ontology an Reference' (*Nous,* Vol. 15, 1981), who see mathematical concepts as applying to physically real properties.

[5] This analysis would go very naturally with a suggestion of non-traditional truth conditions, in terms of reliable psychological processes which we can correlate on empirical grounds with the physical world, as hinted at in Chapter 8 of my *A Guide Through the Theory of Knowledge* (Dickenson Press 1977). The theme of the indirect connection between mathematical formalism and empirical fact is developed a bit further in my 'Mathematical models: questions of trustworthiness' *(British Journal for the Philosophy of Science*,Vol 44, No. 4, 1993.)