Logique & Analyse **165–166** (1999), 107–126

WHY THE LARGEST NUMBER IMAGINABLE IS STILL A FINITE NUMBER*

JEAN PAUL VAN BENDEGEM

1. *Introduction*

The object of this paper is straightforwardly philosophical. Its subject is strict finitism: the view that there are such things as the largest (and the smallest) number and hence that infinities play no part whatsoever in (applicable) mathematics. One of the standard criticisms directed against strict finitism concerns both the ontological and epistemological status of these limit numbers. After all, so the basic remark goes, if I can imagine a largest number N, then I can equally imagine N+1, so N isn’t the largest number. If, however, I cannot imagine such N, then how can I claim that it is the largest number? The argument seems as devastating as Herman Weyl’s classic argument directed against a discrete geometry if it is meant to be an approximation of Euclidean (infinite) geometry. In his own words: “If a square is built up of miniature tiles, then there are as many tiles along the diagonal as there are along the sides; thus the diagonal should be equal in length to the side.” (Weyl, [1949], p. 43). As this argument has no longer the strength it once had,[1] and therefore no longer the status of a “killer” argument against strict finitism, it is necessary to deal with the largest number argument as well (if only for symmetry’s sake). Which is the philosophical aim of this paper.

A different way of formulating what is done here is by referring to Priest’s paper [1994], where he writes the following: “As long as the greatest numeral possible with our notation system is *so large that it has no physical*

*My thanks to one of the referees of this paper, who made a very thorough analysis and has allowed me to make clear improvements. Of course, if the end result is still not convincing, this is entirely my responsibility.

[1] There are actually quite simple considerations to counter the argument. See my [1987] and Forrest [1995], to name but two. In my attempt I introduced the notion of the width of line, that is supposed to be large compared to the size of miniature tiles. Defining the length of a line as its surface in terms of the number of tiles divided by the constant width, leads to a distance function that approximates the Euclidean distance function as close as necessary. Forrest on the other hand defines a line as a set of touching circles with a fixed radius and defines the length in terms of the number of circles needed to make up the line.

or psychological significance ..." (p. 8, my emphasis). One might expect, although it is not a trivial matter, as I will show in this paper, that for a strict finitist the idea of "no physical significance" seems quite meaningful. It seems less likely that the same holds for "no psychological significance". After all, does it not mean that psychological considerations are introduced in a discussion about the nature of mathematics? Is this not psychologism all over again? It is my aim to show that a clear understanding can be given of this idea without ending up in a form of psychologism. This paper can thus be read as an elaboration of Priest's proposal.

Finally, the paper can also be read as (I hope) an improved presentation of the first chapter of my [1987] where I tried to justify strict finitism, based on similar ideas as presented here. However that approach did not seem to carry an awful amount of convincing power. Hence this second attempt. In the next paragraph I really start from scratch —I will discuss the properties and particularities of the act of labelling— as I firmly believe that the roots of the misconceptions about the (mere) possibility of strict finitism are to be found at this fundamental level.

2. *Basic assumptions about labels*

Let us start by assuming that the world is indeed finite:[2]

(A1) A finite world W is given. That is, it consists of a finite set or totality of "objects" w_1, w_2, ..., w_n. The "objects" w_i are considered to be atoms.

Comments on (A1):

In the first place, it is obvious that this formulation is meant to be as general as possible. Therefore nothing is said about (i) what precisely these "objects" are (space points, time points, space-time points, regions of space, time, objects with extension, sharply delineated objects or vague objects, ...), (ii) how these objects are related to one another (in mereological terms, is there a part-whole relation?, is overlapping possible or not?, ...). The mere

[2] Note that in this assumption and others to follow terms will occur such as "finite", "set", "totality" and that indices are used. It should be understood that these do not have their standard mathematical meaning, i.e., I do not take these terms to be imbedded in a standard mathematical theory. The best way of looking at these words at this stage of the presentation is as vaguely defined terms that will get refined and sharpened during the presentation. This allows for a control mechanism: rereading the paper with the "final" definitions in mind should lead not to incoherencies or inconsistencies. In yet different terms, what I am looking for is self-consistency.

purpose of (A1) is to guarantee that the world W is finite "in the large", i.e., there are only a finite number of objects and that it is finite "in the small", i.e., the objects themselves are no further analysable.[3] One should note two things. The first is the asymmetry between the large and the small, the former one has to do with the number of objects, the latter one with the objects themselves. The second is that there is no necessity for the objects to be atomlike. One could imagine each object w_i being composed of some other objects w_j.[4]

In the second place, it must be clear that (A1) is an empirical claim. A thought experiment[5] to sustain this claim is the following: consider the possibility that the universe, as well as man has an infinite past. Infinitely long ago, a secret society was formed: *The Ancient Pythagorean Society (APS)*. What they have done is to check case-by-case Goldbach's conjecture (every even number is the sum of two primes) "starting" at infinity. Thus, n days ago, they checked case 2n. This implies that the *APS* can in fact decide today Goldbach's conjecture. In fact, yesterday they have checked the "last" case. Thus, if the *APS* were to exist, strict finitism in my view becomes pointless. Do note that, although the thought experiment suggests otherwise, the question is not a purely ontological one, i.e., whether or not an *APS* can exist, but also an epistemological one, in the sense that the experiment presupposes an extended form of mathematical proof, to be very specific, why would you believe the president of the *APS* when she announces that Goldbach's conjecture is true (or proved?).

It is tempting to believe that, if the world is (ontologically) infinite one way or another, then that means the end of strict finitism. But one has to take into account the question whether and, if so, how we come to know this. Either we have direct access —and, yes, in that case it does mean the end of strict finitism as we know it— or we do not. The former case can be put aside as no one, as far as I know, has directly experienced an infinite number of objects.[6]

[3] There is a close analogy here with Euclid's definition of a point in the first book of the *Elements*: a point is that which has no further parts.

[4] The idea should be explored further, for it opens the possibility of a curious combination of the finite and the infinite. Suppose we have three objects w_1, w_2 and w_3, such that each one is made up of the two others. Then at the same time there is a strictly finite model $W = \{w_1, w_2, w_3\}$ and an infinite model if we take into account the fact that $w_1 = \{w_2, w_3\} = \{\{w_1, w_3\}, \{w_1, w_2\}\} = ...$ However, in this paper I will restrict myself to the atomic case.

[5] With thanks to Gerald J. Massey who first put this problem to me.

[6] This statement is to be taken quite literally. I am not talking about, e.g., mystical experiences whereby the mystic claims to have experienced God as an infinite being. Although it is an important discussion to decide whether or not one can consider such mystic ability

So, unless humans acquire entirely new epistemic capacities, we can safely ignore this possibility. However, in the latter case there can still be arguments that support the idea of an infinity out there in the world. The question then is how definitive these arguments can be. Can they be necessary in some sense? It is my firm belief that the necessity of such arguments can always be doubted.[7] An illustration will perhaps make clear what I have in mind.

Suppose that —and one might forcefully argue that the present-day situation shows this supposition to be true— one were to tell me that our best, say, physical theories about the universe are such that the world is infinite, both in terms of space and time, and that no alternatives are available. In this particular case as well, I would argue that this need not mean the end of strict finitism. After all, given our epistemic capacities, only a finite part of the universe is accessible. Therefore all statements having to do with the global structure of the universe —and most definitely so, if it is supposed to be infinite— are extrapolations in one sense or another and thus presuppose principles that allow for these extrapolations, say, isotropy, homogeneity, to name the “classics”. However alternatives for such principles are imaginable. To give but one concrete example. Suppose we have a model M of the universe that is spatially infinite. Let M^* be the accessible part of the model M. Then it is always possible to construct a new model M' such that M' coincides with M^*, but outside of M^*, M' no longer obeys the same physical laws and nothing prevents one from introducing such laws that express, among other things, that M' is spatially and temporally finite, the extreme case being that where quite brutally $M' = \emptyset$. This raises the question of course why the laws would change from one part of the universe to another, but then it is instructive to read Paul Davies ([1994], especially chapter 10: *Sudden Death – and Rebirth*) where precisely such a scenario (with standard quantum mechanics in the background to indicate that the author is not exploring a pythonesque world) is presented.

I do realize that this sort of argumentation should raise doubts. After all, on the one hand, I am claiming that it is empirical, hence, roughly speaking, “facts” can decide the matter, but, on the other hand, our best theories of the moment, that rely on infinities, seem to have little or no impact on the very same matter. My defence should be read in first instance as a rejection of the

as an epistemic ability or accessibility, whether or not this experience can be communicated to others, this is not the issue at stake here. In terms of a classic example, it seems quite out of the question to see or experience *all* natural numbers at once. Thus the image of a row of numbers “disappearing on the horizon” is an inadequate image, in the very same sense that the notation 1, 2, 3, ... is inadequate.

[7] I am not claiming that this holds in general. In other contexts it can very well be that there are arguments that do have this compelling force.

idea that it is a *straightforward* empirical problem. In other words, I am not particularly impressed by the argument: "How can you be a strict finitist, if our best theories need infinities? Surely you must be wrong."

As the next step, since we are dealing with numerals and numbers,[8] let us introduce some counting mechanism in this finite world. Obviously, the simplest counting or notation system imaginable consists of a list of strokes, such that for every object to be counted a stroke is associated with it. Put otherwise, in this finite world labels are given to objects, the labels themselves being objects.

(A2) In this world W a labelling machine (or entity if a more neutral term is needed) is present, call it M. The actions of M are quite simple: what M does is to pick out particular objects w_1^*, w_2^*, ..., w_k^* and to use these objects as labels for other objects. It is assumed that (i) label and the labelled object are to be distinct (or distinguishable), (ii) labels are not labelled themselves, and (iii) one label labels one object and vice versa.

Comments on (A2):

In the first place —and, to a certain extent, this is the basic postulate of strict finitism— it must be clear that it will be impossible for M to label all objects in W.[9] As a label is itself an object w_i^*, distinct of the object it labels because of (i), and, because of condition (ii), self-labelling cannot occur,[10] at most half of the objects can be labelled by the other half. If we assume in addition (though this is not a necessary requirement) that the objects making up the machine M do not enter into the game, it is even less than half of it, namely half of W\M.

[8] This paper does not deal with the strict finitist approach to geometry. I refer the reader to my [1995], where such an approach is outlined. From a strict finitist perspective, arithmetic and geometry turn out to be quite distinct things.

[9] This trivial observation has a grave consequence. M cannot know that W is finite through labelling all objects. Hence the finitude of the world is just as abstract a statement as it's being infinite.

[10] This includes also the case of cyclical labelling. Suppose —I owe this example to the referee— a world consisting of three objects w_1, w_2 and w_3. w_1 labels w_2, w_2 labels w_3 and w_3 labels w_1. This contradicts the second condition, because w_1 is itself a label and is now labelled by w_3.

Short digression

It is important to realize that this rather "innocent" statement —at most half of the objects in the world can be labelled—, when written out in more detail is actually a very delicate matter. Take for instance the problem to find out what the minimal assumptions are needed to actually prove such a statement. Assume in formal terms that we have the following:

(a) a finite set W,
(b) a relation L over W, i.e., a subset of W x W (the labelling relation).
I will denote by
(c) D(L) the domain of L, i.e., $D(L) = \{x \mid (\exists y)xLy\}$,
(d) R(L) the range of L, i.e., $R(L) = \{x \mid (\exists y)yLx\}$.

A first problem is to reformulate such phrases as "labels are distinct from things labelled" in formal terms. At least the following sentences are potential candidates:

(i) $(\forall x)((\exists y)xLy \supset \sim(\exists z)(zLx))\&(\forall x)((\exists y)yLx \supset \sim(\exists z)(xLz))$,
(ii) $(\forall x)(\forall y)(\forall z)(xLy \supset \sim zLx)$ & $(\forall x)(\forall y)(\forall z)(yLx \supset \sim xLz)$,
(iii) $\sim(\exists x)(\exists y)(\exists z)(xLy\&zLx)$ & $\sim(\exists x)(\exists y)(\exists z)(yLx\&xLz)$,
(iv) $(\forall x)(\forall y)(xLy \supset x \neq y)$,
(v) $D(L) \cap R(L) = \emptyset$.

Note that (i), (ii) and (iii) do not require an identity predicate, whereas (iv) does, and (v) requires (a part of) the language of set theory. Even if we make a choice for the weakest formulation in terms of the expressibility of the language used, then (i), (ii) and (iii) have to be shown provably equivalent. From the viewpoint of classical logic, the three statements are indeed equivalent. In addition, if there are no language restrictions, then classically speaking, (v) is equivalent to (i), (ii) and (iii) as well. However, if one restricts oneself to, e.g., intuitionistic logic, then one cannot go from (iii) to (ii) or (i). This raises the difficult question what kind of logical principles and rules can be justified when thinking about the world and labels in such general terms as is proposed here. The obvious argument that, since everything is assumed to be finite, all classical rules are applicable, has to be rejected. As Wim Veldman in his [1995] shows, even for finite sets the excluded middle can fail.[11] To a certain extent

[11] Examples are easily produced. Suppose a set A consisting of two elements 0 and n, A = {0,n}. n is defined as follows: n = 0 if Goldbach's conjecture is true, $n \neq 0$ if Goldbach's

what we have here is a similar problem as with proofs for the existence of God. What logical inference rules can be used or not when talking about divine entities? What kind of justification is acceptable? I will not pursue this philosophical matter further, but it is clear that the matter is complicated.

A second problem concerns the exclusion of self-labelling. It is easy to demonstrate that all proposals, excluding (iv), do not allow for self-labelling.[12] The one-label-one-object condition is easily translated (if the identity predicate is available):

(α) $(\forall x)(\forall y)(\forall z)((xLy \,\&\, xLz) \supset (y = z))$
(β) $(\forall x)(\forall y)(\forall z)((xLz \,\&\, yLz) \supset (x = y))$

It remains a nice exercise to show that, given a set of objects and the principles (a)–(d), (v) and (α)–(β), at most half of them can be labels.[13] Perhaps this digression might appear as a trivial mathematical exercise —after all, I am only showing that, if L is one-to-one without loops, then the domain of L cannot be larger than half of W— but one should read it as a form of reverse mathematics, i.e., what are the minimal principles needed to prove such a statement.

End of digression

In the second place, it is also assumed in (A2) that all the objects w_i are accessible to the machine M. There is obviously no necessity at all to accept this condition, which will make things even worse. To be as concrete as possible: if the basic objects making up W are space-time regions with a

conjecture is false. Now take the statement that A contains either one element or more than one element. For the classical logician, this will be obviously true, but on an intuitionistic reading of the disjunction, this no longer holds, as we cannot decide which alternative is the case. It then follows that, since if n = 0, then A = {0}, and, since if $n \neq 0$, then A = {0,n}, we also cannot decide its cardinality, although the set is finite.

[12] Suppose it does happen, then we must have aLy, for some y, and zLa for some z. From the former it follows that $(\exists y)aLy$, but according to (i), it then follows that $\sim(\exists z)zLa$, contradicting yLa.

[13] Suppose that W = $\{w_1,w_2,w_3,w_4\}$. Take w_1 as the first label, w_1^*. Then there must be an object labelled by w_1^*, distinct from w_1, say w_2. Thus $w_1^*Lw_2$. Because of (β), w_1^* is no longer available as a label, and, because of (v), so is w_2. Hence, w_3 is the next candidate, thus w_3^* becomes a label for w_4, as w_2 is excluded because of (α), thus $w_3^*Lw_4$. Thus D(L) = $\{w_1^*,w_2^*\}$ and R(L) = $\{w_3,w_4\}$. Question: how problematic or not is the (at first sight) innocent phrase "take an element of W" in this proof sketch?

lower limit on size, say, in the order of Planck length and time, and if the machine M has capacities similar to human beings, then the objects needed for the labelling are on a much larger scale than the objects to be labelled, i.e., the space-time regions.

In the third place, I have assumed right from the start of this paper that the world is discrete: the objects w_i are distinct and make up the world W. Thus it is thereby assumed that, whatever a (mathematical) continuum might be, it cannot be "for real". This raises the interesting question how it is possible that we humans have this strong illusion of continuity in our daily experience. I have no doubt that a biological-psychological-epistemological explanation is possible, if not already available. However, it would represent an intriguing situation if there would be strong arguments that show that such an explanation is impossible. Then we would have a deep problem for strict finitism.

These two quite simple assumptions immediately lead to a first important thesis:

(T) M will not be able to give a label to all objects in W. If we assume that the act of counting involves at least the capacity to give labels distinct from the objects being counted, then the totality of objects in W is, although finite, not countable.

Observations on (T):

First of all, a clear advantage of the way (T) is formulated is that statements such as "If the universe is finite, then it must be possible for the strict finitist to count all the elementary particles" are far more complex than is usually assumed. It is, e.g., important to decide what the interpretation of the objects will be. If the objects are the elementary particles themselves, then, according to the analysis presented so far, it is clear that it is impossible to count them all (ignoring all other technical problems this choice will unavoidably generate). If, on the other hand, the objects are space-time regions, in such a way that a space-time region is much smaller than the size of an elementary particle, then it might well be possible to count them all. But then of course it will not be the case that "If the universe if finite, then it must be possible for the strict finitist to count all the space-time regions (ignoring once again all practical matters)". Furthermore, "counting" is understood here in the very narrow sense of one-one labelling, whereby label and object labelled are distinct. In short, it is important to realize that such statements that sound extremely plausible and simple at first hearing, in fact are quite complicated statements containing a number of hidden assumptions. This will also become clear(er) in the next chapter of this paper.

Secondly, any mathematics that will rely on the use of these labels —one can think, e.g., of elementary arithmetic in a stroke notation system or of elementary geometry where all points in the space are labelled— will not be able to give a full description of W. At best it will give us a description of a real part of W.

Thirdly, in (T) the phrase “not countable” is used. In a rather peculiar way, this reflects the same idea in standard (naive) set theory, where the set of ordinals cannot, if paradox is to be avoided, have an ordinality. Thus, although (T) mentions W, a “set” of “elements”, yet, it is not possible to specify #W, the cardinality of W.

The last point is really the core of the matter: there is nothing paradoxical about holding simultaneously the following two claims:

(C1) The world W consists of a set of objects w_i,
(C2) The world W does not have a cardinality, #W, expressible in a particular labelling system (such as the stroke notation system),
or, in a weaker version,
(C2′) Even if the world W did have a cardinality, this cardinality could never be expressed in a particular labelling system (such as the stroke notation system).

As soon as one thinks that (C1) must imply the negation of (C2), then of course strict finitism cannot make sense. By denying this implication at least the possibility exists for strict finitism.[14]

3. *Assumptions involving the structure of the world and of labels*

Let us complicate matters, as we must. There seems to be a rather simple counterargument to the above reasoning. Suppose that we agree that a space-time region of the size (Planck-length)$^3 \times$ Planck-time could very well represent the smallest meaningful object in the world. This is quite compatible with present-day physical theories. Suppose that in addition we can rely on cosmological models to give us information about the global volume-age of the universe. Let us fill in the details:[15]

[14] It is interesting to note that similar ideas (although differently motivated) have been expressed by the later Wittgenstein according to the work of Victor Rodych. See his [to appear].

[15] These numbers are to be found in Davies [1982], p. 39 and p. 45.

(a) Planck-length = 10^{-35} m
(b) Planck-time = 10^{-43} s
(c) basic space-time region = 10^{-148} m^3s
(d) volume of the universe = 10^{78} m^3
(e) age of the universe = 10^{18} s
(f) volume-age of the universe = 10^{96} m^3s

Finally, (g) number of objects in the world = (f)/(c) = 10^{244}.

Then, is the number in (g) not a perfect candidate for a specific value of #W?[16] This reasoning will not be in conflict with the argument above, showing that #W cannot be determined, once we realize that (C2) is restricted to a particular labelling system. As must be obvious, in the physicist's argument all kinds of additional properties are used in the argument.

This raises the question what these properties could be. One observation is crucial: it has at least to be possible to label objects that could not be labelled before (besides other considerations). We have been able to create more labels than in the initial situation. Or, in other words, there must exist labelling systems that are more efficient. This leads me straightforward to assumption (A3), that I will call an economy principle:

(A3) Because the labels w_i^* are themselves world objects, specific properties of these labels *as objects* can be invoked to optimise and thereby extend the use of labels.

The most obvious candidate for such properties are the spatial aspects of objects. To give but one example: since the labels are situated in space (and time) we can differentiate between an object w at a particular place and an identical object w' at another place, i.e. w and w' have all properties in common except spatial location. Thus we can construct a labelling system that uses a particular label and differentiates this label at different places. This produces (something similar to) a positional labelling system. It can then very well be that the set of labels is much smaller than the set of objects it can label. Just think of a stroke notational system compared with the decimal system (seen as a labelling system). It follows that of all the objects of W\M, we will need less than half of them to label the remaining ones. Thus we can extend the process beyond half of all the objects. However, this cannot

[16] I must add that this number should rather be considered as a lower limit, in the sense that if we allow a "set" of basic space-time regions to be an object on its own right, then of course we find numbers in the order of 2 to the power of #W. I restrict myself to the basic objects themselves at this stage, but when I deal with imagined numbers, I will be able to handle this problem.

be the end of argument. Because, no matter what we do, we will need *some* labels and, since self-labelling is prohibited, thus again not all objects can be labelled. So, the question remains: how is it possible in the above physicist's reasoning to give a number (and thus a label) for *all* objects in the universe, as this is what the argument seems to be doing? How is this possible?

The answer has to be found in the fact that, although the reasoning above seems to suggest that a direct labelling (and hence a direct count) of all the objects is taking place, actually a lot more is needed. Among other things, apart from the need of a labelling system, a counting and calculating system is required. In addition one needs a theoretical framework that allows one to talk about Planck time and length, it involves cosmological models to derive the age of the universe. In that sense the argument is indeed a highly theoretical argument. Putting aside for the moment the precise details of the process that allows one to give an estimate of #W, thus of all objects, it is clear that we must accept the following principle:

(A4) A theoretical framework or structure allows one to further reduce the set of labels needed to label the objects of the world W, including the possibility of constructing a label for all objects, i.e., #W.

Note that (A4) is not in contradiction with the prohibition of self-labelling. To give a concrete example: suppose that 2n objects are given. Then we will be able to label n of them at most, one by one. If we are now allowed to perform groupings of the objects, say in k groups, each of size n/k. Then we have to label just one group to get n/k and to label the k groups to arrive at n. This requires n/k+k labels, a number that is smaller than n, for n, k ≥ 2. In the same sense if we find out that n objects can be labelled, then doubling that figure gives us the total amount (more or less) of all objects. It is clear that what we have not done is to label each object separately, we actually "jumped" half of them. This is unfortunately not the end of the story.

We now face another problem. Apparently, if we are willing to be more generous in what we accept as labelling systems (and there seems little reason not to do so), then it is possible after all to give an explicit expression for the cardinality of W, #W. And, as one can see quite explicitly above, it does not take up all that much space to write #W down. Therefore nothing prevents us to write down the double of that number, or, if we agree that #W itself can function as a label, then we can consider $\#W^{\#W}$. This surely is a label that has no physical significance (in the sense of being realized at a particular space-time region). This must bring us to the following difficult question: how many labels are constructible, if we are so generous? Or, to reformulate the question: how many labels are *imaginable* under these conditions?

Before taking on this problem, let me note that the economy principle also generates "waste". As different labelling systems are introduced, some "sets" of objects will be characterized by a number of different labels according to the labelling system used. And if theoretical considerations are invoked, then the number of labels can become quite large. In standard classical infinitistic mathematics, the number can be actually infinite.[17] On the one hand, there is a cost reduction but one cannot thereby avoid the introduction of unnecessary complications, thus increasing the costs. One step further leads to the suggestive idea that in a strict finitist framework a labelling system of zero cost is impossible. I must at this moment leave open the problem what evidence could be given in favour of this idea.

4. *The role of imagination and communication*

The preceding analysis has brought us to the conclusion that the question to answer is not what is the label (if any) of the largest totality (such as W) existing, but what is the largest label *imaginable*? At first sight, this seems an impossible problem. Are there limits to human imagination? Are we forced to assume one or other position in the debate concerning the relations between body and mind? For, if all thoughts I have are reducible to (sets of) brainstates and these are counted in terms of the objects of W, then I will end up with quite a different story if instead I assumed that the world of thoughts, say Wt, is a separate world from W. What reasons do I have then for assuming that Wt is finite? In a quite desperate attempt to avoid these problems —as they lead to a total impasse— I will address the problem along the following lines.

Whatever goes on in one's mind is not the relevant issue here. Whatever it is that is imagined by someone, in order that I can judge whether or not what is imagined is a label (or an element of a labelling system) it has to be communicated to me. In short:

(A5) If different labelling systems M_1, M_2, ..., M_k are present in the world W, then it is necessary that whatever the labelling system M_i one comes up with, it has to be communicable to the other labelling systems.

What are the consequences of this assumption? First of all, if the labels have to be communicated, they must take up "public" space, thus the mere fact of

[17] Strictly speaking in standard mathematics, a grouping of two objects can be labelled by the numeral 2, but also by II or 4-2 or 8/4 or, for that matter, 2n/n for all n or ..., if the full mathematical framework is available. Obviously if one dwells in infinite realms, discussions about finite costs are inappropriate. Not so for the strict finitist.

having to write down the labels requires objects in the world and since by assumption W is finite, so are the labelling systems. Hence (A5) guarantees that no matter what it is that we have imagined, it will have to be finite when it is expressed.[18]

The second point is that the existence of different labelling machines —if human beings can be seen, among different things, as labelling machines, then this point is trivial— leads quite naturally to the existence of different labelling systems. This recaptures the idea of Yessenin-Volpin and its further development in the work of David Isles.[19]

The third point is that under these conditions it is straightforward to show that the procedure "Give me any numeral n you can imagine, I will give you the next one" has to break down at a certain point. Ask any person to imagine a very large numeral, say, in decimal presentation. Usually what we do is to form a picture, say, we see a blackboard and it is covered with ciphers all over. But that won't do. For once we have such a picture, it is obvious that is communicable, hence that it is finitely expressible and hence that there is room to imagine the next numeral and to communicate it. Thus, the alternative must be that the numeral is so large that it cannot be imagined, thereby making it senseless to talk about the next one. I will return to the implicit paradoxical nature of what I just wrote. What is being asked is to imagine a numeral so huge that it cannot be imagined.

Let me settle at this stage of the presentation an obvious counterargument: is it not a silly notion that I can imagine all numerals up to n, and then suddenly for the next one, my imagination fails me? What the argument shows is that the notion of the "largest numeral imaginable" must be a vague notion. This observation is supported by the fact that paradoxes concerning vague predicates also apply to this situation. One of the most straightforward connections is with the Wang paradox, itself a variant of the Sorites paradox. Suppose that Imag(x,n) is an abbreviation for "x is capable of imagining the

[18] Do note that I am talking here about finite expressibility. In the next chapter I will deal with what happens if the labels (or, by extension, the signs) talk about infinite structures. My aim here is to show how one can avoid the mind-body problem by relying on communicability.

[19] Yessenin-Volpin considers different number series or rather series of numerals, not just the unique natural number sequence. All these sequences are finite, but the largest numeral differs from series to series. Yessenin-Volpin does not speak of strict finitism, but of ultra-intuitionism instead. I do not elaborate his ideas any further in this paper as his aim is quite distinct from mine. He was looking for a kind of finitary consistency proof of ZF (do note the "kind of"), which is not my objective at all. In addition, the work of Yessenin-Volpin is rather cryptic and hence very difficult to understand, but David Isles has succeeded in "distilling" his version from these papers that make a lot of sense. A very nice summary is given in Epstein & Carnielli [2000], pp. 263–270, containing, of course, further references.

numeral n", then it is claimed that both:

(1) Imag(x,1)
(2) $(\forall n)(\text{Imag}(x,n) \supset \text{Imag}(x,n+1))$

are extremely plausible. After all, (2) is nothing but a reformulation of the idea that the next numeral can always be imagined. But, given (1) and (2), the conclusion

(3) $(\forall n)\text{Imag}(x,n)$

follows immediately by mathematical induction and that is nonsense. I will not enter into the discussion of vague predicates here, suffice it to mention that if the predicate 'Imag' is seen as a vague predicate, then there is no real problem to the imagination gradually failing to produce pictures of larger and larger numerals.[20]

However, our generosity might turn against ourselves and we might end up at the bottom of a slippery slope. For what about the following argument. Surely human beings have demonstrated that they can think about infinities and that they can imagine infinite situations. In addition, we have perfectly finite labels for such infinities, in the case of mathematical cardinal infinities, we use the labels $\aleph_0$, $\aleph_1$, $\aleph_2$, ... and we can write down finitely expressible statements such as: $\aleph_0 + \aleph_1 = \aleph_1 + \aleph_0$, that can therefore be communicated. All of this seems to satisfy the requirements presented above, so what then does strict finitism actually claim? Does is not inevitably reduce to finitism in (one of) the (many) sense(s) of Hilbert? The next chapter shows the way out.

5. *Tackling the general problem*

In general terms the problem we have to deal with is the following. Suppose that a finite set of sentences A_1, A_2, ..., A_n, all finitely expressible, is given.[21]

[20] The technicalities involved are unfortunately not all that simple. As far as I know, there is little or no consensus on an appropriate formalisation of vague predicates. See, e.g., Keefe & Smith [1997] for an overview. In my [2000], I have presented a form of vague mathematics, using a supervaluational method that obviously has the drawback that it supposes classical models "in the background" with non-vague predicates. It thereby suggests that vague predicates cannot be formally expressed "on their own", but require sharp predicates in the background. The aim however was merely to show that such a kind of mathematics is possible.

[21] I restrict myself to first-order predicate logic, as most of mathematics (and that is what the paper is about) is expressible in this language.

Suppose that what the sentences talk about requires infinities in some way. To be more precise, in standard logical terms, suppose that the set of sentences has models with infinite domains. Is the strict finitist not forced to accept this model, as all is finitely expressible, and hence to accept infinities? The answer I propose here is the following: in all such cases, there will always be finite (quasi-)models as well. I realize the implausibility of this claim, so let me present an example to indicate what I have in mind.

One of the “classic” examples is the following set of three sentences:

(a) $(\forall x)\sim xRx$
(b) $(\forall x)(\exists y)xRy$
(c) $(\forall x)(\forall y)(\forall z)((xRy \,\&\, yRz) \supset xRz)$.

It is easy to show that there cannot be a finite model for R. Hence any model of this set must have an infinite domain. However one should be clear about the notion of a model. If the question is whether a finite model can be formulated such that the syntactical elements occurring in (a), (b) and (c) refer *in one way or another* to elements in the domain and such that (a), (b) and (c) are true in the model, then actually the answer is trivial. As nothing else besides the given sentences comes into play, nothing prevents me from interpreting the universal quantifier as an existential quantifier and leave the existential quantifier as it is. Then (a), (b) and (c) are translated as:

(a)* $(\exists x)\sim xRx$
(b)* $(\exists x)(\exists y)xRy$
(c)* $(\exists x)(\exists y)(\exists z)((xRy \,\&\, yRz) \supset xRz)$.

This set obviously has a finite model, to be specific: take a two-element domain $\{a,b\}$ and stipulate the interpretation of the relation R as: $\{<a,b>\}$. In a way the position that I am defending here is to take Hilbert’s formalist view seriously. What I have in mind for R is not important, the problem is to find a model that satisfies certain conditions.

If this procedure were available at all times, then the solution to our problem would be very simple indeed. Replace all universal quantifiers by existential quantifiers, thereby obtaining a finite set of existential sentences and that set, of course, if satisfiable at all, can be satisfied in a finite domain.

Unfortunately, this is only half of the story. We are not particularly interested in finite sets of sentences in isolation, but we also want to reason about them. Thus I need logical rules, say the rule that $(\forall x)A(x) \equiv \sim(\exists x)\sim A(x)$, but as is clear, under the given interpretation, this does not hold.[22] Fortunately, several strategies are possible. I will sketch one such possibility. Suppose for simplicity that all sentences can be formulated in prenex normal form. Then, replace from left to right each occurrence of a universal quantifier $(\forall x)A(x)$ by $(\forall x)(x \leq K \supset A(x))$, for K some finite number, and likewise for the existential quantifier, i.e., $(\exists x)A(x)$ is replaced by $(\exists x)(x \leq K \supset A(x))$.

As an illustration, let us have another look at the classic example. (a), (b) and (c) are in prenex normal form and to keep matters simple assume that K = 2. Thus we have a two-element domain with names a and b. Then the three sentences are rewritten as follows, following the outlined procedure and some additional simplifications:[23]

(a)+ $(\forall x)(((x = a) \vee (x = b)) \supset \sim xRx)$
(b)+ $(\forall x)(\exists y)(((x = a) \vee (x = b)) \& (y = a) \vee (y = b)) \supset xRy)$
(c)+ $(\forall x)(\forall y)(\forall z)((((x = a) \vee (x = b)) \& ((y = a) \vee (y = b)) \&$
$((z = a) \vee (z = b))) \supset ((xRy \& yRz) \supset xRz))$.

It is easy to see that a three-element model M = {a, b, c} will do the job. R is interpreted as {<a,b>}. (a)+ and (c)+ present no problem, and note for (b)+, that if x = a, then pick y = b and, if x = b, then pick y = c. At first sight, the only thing that is happening here is to take the infinite model and to look at a finite part of it through the introduction of an unnamed element in the domain —c in the above case— and to use this element as "the rest of the infinite domain reduced to one element".

There is, however, more to it. The simple translation technique guarantees that if $A(x_1, x_2, ..., x_n)$ holds in the classical model, then $A+(x_1, x_2, ..., x_n)$, i.e., the formula A after all quantifiers have been replaced as above, holds as well. So no "truths" are lost. Quite to the contrary, additional truths will be found. E.g., the sentence $(\exists x)xRx$ does not hold in the classical infinite model, but after rewriting, this becomes $(\exists x)(((x = a) \vee (x = b)) \supset xRx)$.

[22] Under the given interpretation $(\forall x)A(x) \equiv \sim(\exists x)\sim A(x)$ reduces to $(\exists x)A(x) \equiv \sim(\exists x)\sim A(x)$. Take for 'A(x)' x = a, then it says that $(\exists x)(x = a) \equiv \sim(\exists x)(x \neq a)$. This says there can only be one element in the domain. In the given case, such a model is inconsistent, as (a) says that R does not hold between a and itself, whereas (b) says that it does as both x and y are interpreted as a.

[23] If I abbreviate $(x = a) \vee (x = b)$ by S(x), then (b) is first rewritten as $(\forall x)(S(x) \supset (\exists y)xRy)$ and then rewritten as $(\forall x)(S(x) \supset (\exists y)(S(y) \supset xRy))$. Bring out the quantifier and relying on classical logic, namely that $A \supset (B \supset C)$ is equivalent with $(A \& B) \supset C)$, (b)+ is obtained. The same holds for (c)+.

Take c for x and the formula is true. This is the reason why the term quasi-model seems more appropriate. Do note that, if we are interested above all in truths, then nothing is lost.

It is an entirely different matter if the additional requirement is that some falsehoods should also remain (exclusive) falsehoods in the quasi-model. E.g., if the requirement is that for no element of the domain xRx is to hold, then we have a problem with the above quasi-model. It is perhaps not an ideal strategy to suggest that $(\forall x){\sim}xRx$ should also hold in the quasi-model, thus that it is true that there is no x such that ${\sim}(\forall x){\sim}xRx$, i.e., $(\exists x)xRx$ holds, as required.

Some comments:

In the first place, a classical mathematician or logician will surely remark that, however clever this procedure might be, it still implies a reinterpretation of the universal quantifier. "For all x, ..." does not have its classical meaning, for, in all cases, we are supposed to read "For all $x \leq K$, ...". But what if the classical interpretation is required? I wish to be rather brief about this problem. One ends up in exactly the same situation as when discussing discrete versus continuous time (and similarly discrete versus continuous space). Namely, if time is discrete, then if two consecutive time moments are taken, it does not make sense to ask what is in between, as, if time is discrete, there is no in between. Asking what is in between carries the presupposition that it makes sense to ask what is in between and precisely that does not make sense. In the very same way, if one asks for a standard interpretation of the universal quantifier, then one presupposes the possibility of an infinite domain, hence one can never have such an interpretation in a finite domain. This may seem a rather crude way of solving a problem, but occasionally it just so happens that the solution is that straightforward.

In the second place, related to the first point, an argument based on compactness does not apply either. Suppose we have a quasi-model for every K. Does it then not follow that we can put together all these finite quasi-models into one model that will necessarily be infinite? For that is precisely what compactness tells us. The argument is the same: the expression "for every K" presupposes a classical reading of the universal quantifier. Hence, "for every K" has to be thought in terms of "for every $K \leq L$", where L is some finite number determined by the particular circumstances one is considering.

In the third place, the kind of interpretation presented here is just one of a quite extended set of possibilities. In the above example, I have replaced the infinite domain by a finite partition, consisting of two elements a and b and a special name c for the (infinite) remaining part of the domain. It is easy to see that any finite partition of the domain will do. This actually turns out to be a very powerful technique indeed. Do note that some connections can be made with other types of logics than the ones that have been discussed so far in this paper, i.e., classical and intuitionistic logic. If we restrict ourselves to the translated sentences and define two new quantifiers $\forall^*$ and $\exists^*$ as an abbreviation, then we get pretty close in the neighbourhood of paraconsistency, for then both $(\forall^*x)\sim xRx$ and $(\exists^*)xRx$ hold in the model. What comes out of this has been presented and discussed in my [1994].

All these considerations justify the following claim: no matter what theory we are dealing with, it is always possible to deal with finite (quasi-)models, thus no infinities whatsoever are needed, even if we leave room for all the things the human imagination can come up with, the only restriction being that whatever the imagination has produced, it should be communicable using only a finite number of signs or labels.

6. *Summary*

What conclusions can be drawn from all this? The following cases can be distinguished:
(a) Labels are used merely as labels: if the world is finite, so is the set of labels, and it is impossible the label all “objects” in the world.
(b) Labels form a structured set. In this case the labelling process can become more economical and more efficient, but it remains the case that the set of labels stays finite.
(c) Labels form a structured set inserted in a theoretical framework.
Here two subcases can be distinguished:
(c1) There are interpretations of the theoretical framework that refer to “objects” in the world. Obviously in this case everything remains finite again on the assumption that the world is finite.
(c2) There are no specific interpretations that refer to “objects” in the world. It is then always possible to find finite quasi-models that are derived from the classical infinite models of the theoretical framework. In some cases (as shown in the example above) these quasi-models can be seen as extensions of the classical model since it is possible to keep all classically true statements true in the quasi-model. Thus in those cases no truths are lost.

The last case also applies to all labels that can be imagined by a labelling machine, if the requirement is that the labels should be communicable. Hence, if it is representable, it is obvious that we can imagine something larger, as we usually represent something in an environment, hence additional space is available. What we have to imagine, is a label such that if we try to represent it, we should fail to do so. Hence the agreement with Priest’s description quoted at the beginning of this paper: “*so large that it has no physical or psychological significance* ...”. It is paradoxical to be sure. If formulated in terms of questions, the problem becomes immediately obvious. The question "What is the largest label or numeral that is not imaginable?", should not be answered by "The label so-and-so with properties such-and-such", because then it has been imagined, thereby not answering the question. The answer must be: "Whatever it is, that label". An alternative reply would be: "The largest label is that label about which questions such as the question posed cannot be asked". It is that label that ceases to be that label as soon as something is said about it. A conclusion that fits in nicely with the argued for vagueness of the largest label.

Vrije Universiteit Brussel
Centrum voor Logica en Wetenschapsfilosofie
Universiteit Gent

REFERENCES

Paul DAVIES: *The Last Three Minutes. Conjectures about the Ultimate Fate of the Universe.* New York: Basic Books, 1994.

Paul DAVIES: *The Accidental Universe.* Cambridge: Cambridge University Press, 1982.

Richard L. EPSTEIN & Walter A. CARNIELLI: *Computability. Computable Functions, Logic, and the Foundations of Mathematics with Computability: A Timeline.* London: Wadsworth, 2000, (2nd edition).

Peter FORREST: "Is Space-Time Discrete or Continuous? – An Empirical Question". *Synthese*, vol. 103, no. 3, 1995, pp. 327–354.

Rosanna KEEFE & Peter SMITH (eds.): *Vagueness: A Reader.* London: MIT, 1997.

Graham PRIEST: "What Could the Least Inconsistent Number Be?" *Logique et Analyse*, vol. 37, 145, 1994 (published 1996), pp. 3–12.

Victor RODYCH: "Wittgenstein's Anti-Modal Finitism". *Logique et Analyse*, to appear.

Jean Paul VAN BENDEGEM: *Finite, Empirical Mathematics: Outline of a Model.* Gent: Werken uitgegeven door de Faculteit Letteren en Wijsbegeerte, R.U. Gent, volume 174, 1987.

Jean Paul VAN BENDEGEM: "Zeno's Paradoxes and the Weyl Tile Argument". *Philosophy of Science*, East Lansing, Michigan, 1987, 54, 2, pp. 295–302.

Jean Paul VAN BENDEGEM: "Strict Finitism as a Viable Alternative in the Foundations of Mathematics". *Logique et Analyse*, volume 37, nr. 145, 1994 (date of publication: 1996), pp. 23–40.

Jean Paul VAN BENDEGEM: "In defence of discrete space and time". *Logique et Analyse*, volume 38, nrs. 150–151–152, 1995 (date of publication: 1997), pp. 127–150.

Jean Paul VAN BENDEGEM: "Alternative Mathematics: The Vague Way". In: Décio Krause, Steven French & Francisco A. Doria (eds.), Festschrift in Honour of Newton C.A. da Costa on the Occasion of his Seventieth Birthday, *Synthese*, vol. 125, nos. 1–2, 2000, pp. 19–31.

Wim VELDMAN: *Some Intuitionistic Variations on the Notion of a Finite Set of Natural Numbers.* Nijmegen: Dept. of Mathematics, University of Nijmegen, 1995, (Report no. 9520).

Hermann WEYL: *Philosophy of Mathematics and Natural Sciences.* Princeton: Princeton University Press, 1949.