17

Mathematics is megethology

Mereology is the theory of the relation of part to whole, and kindred notions. Megethology is the result of adding plural quantification, as advocated by George Boolos in [1] and [2], to the language of mereology. It is so-called because it turns out to have enough expressive power to let us express interesting hypotheses about the size of Reality. It also has the power, as John P. Burgess and A. P. Hazen have shown in [3], to simulate quantification over relations.

It is generally accepted that mathematics reduces to set theory. In my book *Parts of Classes,* [6], I argued that set theory in turn reduces, with the aid of mereology, to the theory of singleton functions. I also argued (somewhat reluctantly) for a 'structuralist' approach to the theory of singleton functions. We need not think we have somehow achieved a primitive grasp of some one special singleton function. Rather, we can take the theory of singleton functions, and hence set theory, and hence mathematics, to consist of generalisations about all singleton functions. We need only assume, to avoid vacuity, that there exists at least one singleton function.

But we need not assume even that. For it now turns out that if the size of Reality is right, there must exist a singleton function. All we need as a foundation for mathematics, apart from the framework of megethology, are some hypotheses about the size of Reality.

First published in this form in *Philosophia Mathematica* 3 (1993), 3–23. This paper consists in part of near-verbatim excerpts from David Lewis, *Parts of Classes* (Oxford, Blackwell, 1991). Reprinted with kind permission from *Philosophia Mathematica* and from Blackwell Publishers.

(Megethology can have no complete axiom system; and it would serve little purpose to fix upon some one official choice of an incomplete fragment. Rather we shall proceed informally, and help ourselves to principles of megethology as need arises. Only the less obvious principles will be noted.)

This article is an abridgement of parts of *Parts of Classes* not as it is, but as it would have been had I known sooner what I know now. It begins by repeating some material from the early sections of the book, mostly by means of near-verbatim excerpts. It mostly skips the philosophical middle sections. It gives a new presentation of some of the technical material from the late sections, simplified by using simulated quantification over relations. And it proves the new result that if the size of Reality is right, then there exists a singleton function.

One mereological notion is that of a *fusion* or *sum:* the whole composed of some given parts (see [5], II. 4). The fusion of all cats is that large, scattered chunk of cat-stuff which is composed of all the cats there are, and nothing else. It has all cats as parts. There are other things that have all cats as parts. But the cat-fusion is the least such thing: it is included as a part in any other one.

It does have other parts too: all cat-parts are parts of it, for instance cat-whiskers, cat-quarks. For parthood is transitive; whatever is part of a cat is thereby part of a part of the cat-fusion, and so must itself be part of the cat-fusion.

The cat-fusion has still other parts. We count it as a part of itself: an *improper* part, a part identical to the whole. But also it has plenty of *proper parts* – parts not identical to the whole – besides the cats and cat-parts already mentioned. Lesser fusions of cats, for instance the fusion of my two cats Magpie and Possum, are proper parts of the grand fusion of all cats. Fusions of cat-parts are parts of it too, for instance the fusion of Possum's paws plus Magpie's whiskers, or the fusion of all cat-tails wherever they be. Fusions of several cats plus several cat-parts are parts of it. And yet the cat-fusion is made of nothing but cats, in this sense: it has no part that is entirely distinct from each and every cat. Rather, every part of it overlaps some cat.

We could equivalently define the cat-fusion as the thing that overlaps all and only those things that overlap some cat. Since all and only overlappers of cats are overlappers of cat-parts, the fusion of all cats is the same as the fusion of all cat-parts. It is also the fusion of all cat-molecules, the fusion of all cat-particles, and the fusion of all things that are either cat-front-halves or cat-back-halves. And since all and only overlappers of cats are overlappers of cat-fusions, the fusion of all cats is the same as the fusion of all cat-fusions.

The *class* of all cats is something else. It has all and only cats as members. Cat-parts such as whiskers or cells or quarks are *parts of* members of it, but they are not themselves members of it, because they are not whole cats. Cat-parts are indeed members of the class of all cat-parts, but that's a different class. Fusions of several cats are *fusions of* members of the class of all cats, but again they are not themselves members of it. They are members of the class of cat-fusions, but again that's a different class.

The class of *A*'s and the class of *B*'s are identical only if the *A*'s are all and only the *B*'s; but the fusion of the *A*'s and the fusion of the *B*'s can be identical even when no *A* is a *B*. Therefore we learn not to identify the class of *A*'s with the fusion of *A*'s, and the class of *B*'s with the fusion of *B*'s, lest we identify two different classes with one single fusion.

A member of a member of something is not, in general, a member of it; but a part of a part of something is always a part of it. Therefore we learn not to identify membership with the relation of part to whole.

So far, so good. But I used to think, and so perhaps did you, that we learned more. We learned to distinguish two entirely different ways to make one thing out of many: the way that made one fusion out of many parts, *versus* the way that made one class out of many members. We learned that fusions and classes were two quite different kinds of things, so that no class was ever a fusion. We learned that the part-whole relation applies to individuals, not sets. We even learned to call mereology 'The Calculus of Individuals'!

All that was a big mistake. Just because a class isn't the mereological fusion of its members, we needn't conclude that it isn't a fusion. Just because one class isn't composed mereologically out of its many members, we needn't conclude that there must be some unmereological way to make one out of many. Just because a class doesn't have all and only its members as parts, we needn't conclude that it has no parts.

Mereology does apply to classes. They do have parts: their subclasses. (It remains to be seen whether they have other parts as well.) But for now we have this

First Thesis. *One class is part of another iff the first is a subclass of the second.*

To explain what the First Thesis means, I must hasten to tell you that my usage is a little bit idiosyncratic. By 'classes' I mean things that have members. By 'individuals' I mean things that are members, but do not themselves have members. Therefore there is no such class as the null class. I don't mind calling some memberless thing – some individual – the null *set*. But that doesn't make it a memberless class. Rather, that makes it a 'set' that is not a class. Standardly, all sets are classes and none are individuals. I am sorry to stray, but I must if I am to mark the line that matters: the line between the membered and the memberless. Besides, we had more than enough words. I can hijack 'class' and 'individual', and still leave other words unmolested to keep their standard meanings. As follows: a *proper class* is a class that is not a member of anything; a *set* is either the null set or else a class that is not a proper class.

My First Thesis, therefore, has nothing to say yet about the null set. It does not say whether the null set is part of any classes, nor whether any classes are part of the null set. I shall take up those questions later. Now that you understand what the First Thesis means, what can I say in its favour?

First, it conforms to common speech. It does come natural to say that a subclass is part of a class: the class of women is part of the class of human beings, the class of even numbers is part of the class of natural numbers, and so on. Likewise it comes natural to say that a hyperbola has two separate parts – and not to take that back when we go on to say that the hyperbola is a class of *x*-*y* pairs. The devious explanation of what we say is that we speak metaphorically, guided by an analogy of formal character between the part–whole relation and the subclass relation. The straightforward explanation is that subclasses just are parts of classes, we know it, we speak accordingly.

Second, the First Thesis faces no formal obstacles. We learned, rightly, that membership could not be (a special case of) the part-whole relation because of a difference in formal character. But the subclass relation and the part-whole relation behave alike. Just as a part of a part is itself a part, so a subclass of a subclass is itself a subclass; whereas a member of a member is not in general a member. Just as a whole divides exhaustively into parts in many different ways, so a class divides exhaustively into subclasses in many different ways; whereas a class divides exhaustively into members in only one way. We have at very least an analogy of formal character, wherefore we are free to claim that there is more than a mere analogy.

Finally, I hope to show you that the First Thesis will prove fruitful. Set theory is peculiar. It all seems so innocent at first! We need only accept that when there are many things, then also there is one thing – the class – which is just the many taken together. It's hard to object to that. But it turns out later that this many-into-one can't always work, on pain of contradiction – yet it's just as hard to object to it when it doesn't work as when it does. What's more, the innocent business of making many into one somehow transforms itself into a remarkable making of one into many. Given just one single individual, Magpie or Possum or the null set or what you will, suddenly we find ourselves committed to a vast hierarchy of classes built up from it. Not so innocent after all! This ontological extravagance is just what gives set theory its welcome mathematical power. But, like it or not, it's far from what we bargained for when we first agreed that many can be taken together as one. We could understand set theory much better if we could separate the innocent Dr. Jekyll from the extravagant and powerful Mr. Hyde. The First Thesis is our first, and principal, step toward that separation.

The First Thesis leaves it open that classes might have other parts as well, besides their subclasses. Maybe classes sometimes, or always, have individuals as additional parts: the null set, cat Magpie, Possum's tail (and with it all the tail-segments, cells, quarks, and what-not that are parts of Possum's tail). To settle the question, I advance this

Second Thesis. *No class has any part that is not a class.*

The conjunction of the First and Second Theses is our

Main Thesis. *The parts of a class are all and only its subclasses.*

But the Second Thesis seems to me far less evident than the First; it needs an argument. The needed premises are the First Thesis plus three more.

Division Thesis. *Reality divides into individuals and classes.**

Priority Thesis. *No class is part of any individual.*

Fusion Thesis. *Any fusion of individuals is itself an individual.*

Roughly speaking, the Division Thesis says that there is nothing else except individuals and classes. But that is not exactly right. If we thought that Reality divided exhaustively into animal, vegetable, and mineral, that would not mean that there was no such thing as a salt beef sandwich. The sandwich is no counterexample, because the sandwich itself divides: the beef is animal, the bread is vegetable, and the salt is mineral. Likewise, the Division Thesis permits there to be a mixed thing which is neither an individual nor a class, so long as it divides exhaustively into individuals and classes. I accept a principle of *Unrestricted Composition:* whenever there are some things, no matter how many or how unrelated or how disparate in character they may be, they have a mereological fusion. That means that if I accept individuals and I accept classes, I have to accept mereological fusions of individuals and classes. Like the mereological fusion of the front half of a trout plus the back half of a turkey, which is neither fish nor fowl, these things can be mostly ignored. They can be left out of the domains of all but our most unrestricted quantifying. They resist

* [Added 1996] The Division Thesis is badly worded. The meaning that I intended, and that is required by subsequent discussions here and in *Parts of Classes*, is better expressed as follows: everything is either an individual, or a class, or a fusion of an individual and a class. I thank Daniel Nolan for pointing out the problem to me.

concise classification: all we can say is that the salt beef sandwich is part animal, part vegetable, part mineral; the trout-turkey is part fish and part fowl; and the mereological fusion of Possum plus the class of all cat-whiskers is part individual and part class. Likewise, Reality itself – the mereological fusion of everything – is mixed. It is neither individual nor class, but it divides exhaustively into individuals and classes. Indeed, it divides into one part which is the most inclusive individual and another which is the most inclusive class.

(If we accept the mixed fusions of individuals and classes, must we also posit some previously ignored classes that have these mixed fusions as members? No; mixed fusions are forced upon us by the principle of Unrestricted Composition, but classes having mixed fusions as members are not forced upon us by any otherwise acceptable principle. Let us indulge our offhand reluctance to believe in them.)

All I can say to defend the Division Thesis, and it's weak, is that as yet we have no idea of any third sort of thing that is neither individual nor class nor mixture of the two. Remember what an individual is: not necessarily a commonplace individual like Magpie or Possum, or a quark, or a spacetime point, but anything whatever that has no members but is a member. If you believe in some remarkable non-classes – universals, tropes, abstract simple states of affairs, God, or what you will – it makes no difference. They're still individuals, however remarkable, so long as they're members of classes and not themselves classes. Rejecting the Division Thesis means positing some new and hitherto unheard-of disqualification from membership, applicable to things that neither are classes nor have classes as parts. I wouldn't object to such a novel proposal, if there were some good theoretical reason for it. But so long as we have no good reason to innovate, let conservatism rule.

The Priority Thesis and the Fusion Thesis reflect our vague notion that somehow the individuals are 'basic' and 'self-contained' and that the classes are somehow a 'superstructure'; 'first' we have individuals and the classes come 'later'. (In some sense. But it's not that God made the individuals on the first day and the classes not until the second.) Indeed, these two theses may be all the sense that we can extract from that notion. We don't know what classes are made of – that's what we want to figure out. But we do know what individuals are made of: various smaller individuals, and nothing else.

From the First Thesis, the Division Thesis, the Priority Thesis, and the Fusion Thesis, our Second Thesis follows.

Proof. If the Thesis fails, some class x has a part y that is not a class. By the Division Thesis, y is either an individual or a mixed fusion, and either way, x has an individual as part. Let z be the fusion of all individuals that are part of x; then z is an individual, by the Fusion Thesis. Now consider the difference $x - z$ (the fusion of all parts of x that do not overlap z). Since $x - z$ has no individuals as parts, it is not an individual or a mixed fusion. By the Division Thesis, it must be a class. We now have that x is the fusion of class $x - z$ with an individual z. Since $x - z$ is part of x, and not the whole of x (else there wouldn't have been any z to remove), we have that the class $x - z$ is a proper part of the class x. So, by the First Thesis, $x - z$ must be a proper subclass of x. Then we have v, a member of x but not of $x - z$. According to standard set theory, we then have u, the class with v as its only member. By the First Thesis, u is part of x but not of $x - z$; by the Priority Thesis, u is not part of z; so u has some proper part w that does not overlap z. No individual is part of w; so by the Division Thesis, w is a class. By the First Thesis, w is a proper subclass of u. But u, being one-membered, has no proper subclass. This completes a *reductio*. QED

A consequence of our Second Thesis is that classes do not have the null set as part. Because it was a memberless member, we counted it as an individual, not a class; therefore it falls under our denial that individuals ever are parts of classes. (To be sure, it is *included* in any class, because all its members – all none of them – are members of that class. But it never can be a *sub*class if it is not even a class.) Were we hasty? Should we amend the Second Thesis, and the premises whence we derived it, to let the null set be a part of classes after all? I think not.

Or should we perhaps reject the null set? Is it a misguided posit, meant to streamline the formulation of set theory by behaving in peculiar ways? Again, I think not. Its behaviour is not, after all, so very peculiar. It is included in every class just because it lacks members – and lacking members is not so queer, all individuals do it. And it serves two useful purposes. It is a denotation of last resort for class abstracts that denote no nonempty class. And it is an individual of last resort: we can count on its existence in a way we could count on the existence of nothing else, and once we have it we can fearlessly build up the hierarchy of pure sets: the null set, the singleton of the null set, the singleton of the singleton of the null set, the set of the null set and its singleton, . . . *ad infinitum* and beyond, until we have enough modelling clay to build the whole of mathematics.

Or should we accept the null set as a most extraordinary individual, a little speck of sheer nothingness, a sort of black hole in the fabric of Reality itself? Not that either, I think.

We want a null set, but we needn't be ontologically serious about it. It's useful to have a name that's guaranteed to denote some individual, but it needn't be a special individual with a whiff of nothingness about it. An ordinary individual – *any* ordinary individual – will suffice. Any individual has the first qualification for the job – memberlessness. As for the second qualification, guaranteed existence, that is not really a qualification of the job-holder itself, rather it is a requirement on our method of choice. To guarantee that we'll choose some existing individual to be the null set, we needn't choose something that's guaranteed to exist. It's enough to make sure that we choose from among whatever things happen to exist. The choice is arbitrary. I make it – arbitrarily! – as follows: let the null set be the fusion of all individuals. If any individuals exist, this selects one of them for the job. We also get a handy name for one of the main subdivisions of Reality: if the null set is the fusion of all individuals, then by our Theses, the individuals are all and only the parts of the null set. It makes the null set omnipresent, and thereby respects our 'intuition' that it is no more one place than another. It's far from the notion that the null set is a speck of nothingness, and that's all to the good.

Our Main Thesis says that the parts of a class are all and only its subclasses. This applies, in particular, to one-membered classes: *unit classes,* or *singletons.* Possum's singleton has Possum as its sole member. It has no subclasses except itself. Therefore it is a mereological *atom:* it has no parts except itself, no proper parts. Likewise the singleton of Possum's singleton is an atom; and likewise for any other singleton.

Anything that can be a member of a class has a singleton: every individual has a singleton, and so does every set. The only things that lack singletons are the proper classes – classes that are not members of anything, and *a fortiori* not members of singletons – and those mixed things that are part individual and part class. And, of course, nothing has two singletons. So the singletons correspond one-one with the individuals and sets.

A class has its singleton subclasses as atomic parts, one for each of its members. Its larger parts (unless it is a singleton) are its non-singleton subclasses. A class is the union, and hence the fusion, of the singletons of its members. For example, the class of the two cats Possum and Magpie is the fusion of Possum's singleton and Magpie's singleton. The class whose three members are Possum, Magpie's singleton, and the aforementioned class is the fusion of Possum's singleton, Magpie's singleton's singleton, and the singleton of the aforementioned class. And so it goes.

Taking the notion of a singleton henceforth as primitive, and appealing to our several theses, we get these new definitions. Membership, hitherto primitive, shall be so no longer.

Classes are fusions of singletons.

Members of a class are things whose singletons are parts of that class.

(We must add that only classes have members. A mixed fusion has singletons as parts, but we probably would not want to say that it had members.)

Individuals are things that have no singletons as parts.

The fusion of all individuals, our choice to be the null set, is therefore the mereological difference, Reality minus all the singletons.

Sets are the null set, and all classes that have singletons.

Proper classes are classes that have no singletons.

The class of all sets that are non-self-members had better not be a set, on pain of Russell's paradox. Although it is indeed a non-self-member, that won't make it a self-member unless it is a set. So it isn't; it is a proper class, it has no singleton, and it cannot be a member of anything.

We dare not allow a *set* of all sets that are non-self-members, but there are two alternative ways to avoid it. One way is to restrict composition: we have all the sets that are non-self-members and we have a singleton of each of these sets, but somehow we have no fusion of all these singletons, so we have no class of all sets that are non-self-members. But there is no good independent reason to restrict composition. Mereology *per se* is unproblematic, and not to blame for the set-theoretical paradoxes; so it would be unduly drastic to stop the paradoxes by mutilating mereology, if there is any other remedy. (Just as it would be unduly drastic to solve problems in quantum physics by mutilating logic, or problems in the philosophy of mind and language by mutilating mathematics.) The better remedy, which I have adopted, is to restrict not composition, but rather the making of singletons. We *can* fuse all the singletons of all the sets that are non-self-members, thereby obtaining a class, but this class does not in turn have a singleton; it is proper. Unlike composition, the making of singletons is ill-understood to begin with, so we should not be surprised or disturbed to find that it needs restricting.

I do not say, note well, that we must posit the proper classes for the sake of their theoretical utility. We have them willy-nilly, be they useful or be they useless. We do not go out of our way to posit them. We just can't keep them away, given our Main Thesis and Unrestricted Composition.

The proper classes aren't much use, in fact. For George Boolos has argued convincingly in [1] and [2] that we do not require proper classes in order to formulate powerful systems of set theory. We can get the needed power instead by resorting to irreducibly plural quantification, something well-known to us as speakers of ordinary language. Suppose we say 'If there are some people such that each of Peter's parents is one of them, and every parent of one of them is one of them, then each of Peter's ancestors is one of them'; here it seems, *prima facie,* that we are quantifying just over people – not over classes. (And not over any class-like entities that differ from classes only in name.) And if we say 'Some things are all and only the classes that are non-self-members', this seems to be a trivial truth (given that there are such classes), not a paradoxical assertion of the existence of the Russell class. (Note that this time I said 'classes', not 'sets', to block the solution which invokes a proper class.) I join Boolos in concluding that what's true is just what seems to be true. Plural quantification is not class quantification in sheep's clothing. It is innocent of set theory, except of course when we quantify plurally over classes themselves. I shall make free with plural quantification in presenting mereology. In fact, I've done so already. When I said 'whenever there are some things they have a fusion', my plural quantifier could not have been read without loss as a class quantifier or as a substitutional 'quantifier'. I meant that *whenever* there are some things, regardless of whether they are members of some class, and regardless of whether they are the satisfiers of some formula, they have a fusion.

So we could rebuild set theory within mereology, if only we had the primitive notion of singleton. But that, I fear, is a tall order.

Cantor taught that a set is a 'many, which can be thought of as one, *i.e.,* a totality of definite elements that can be combined into a whole by a law'. To this day, when a student is first introduced to set theory, he is apt to be told something similar. He is told that a set is formed by combining or collecting or gathering several objects, or by thinking of them together. Maybe also he will be given some familiar examples: Halmos's textbook mentions packs of wolves, bunches of grapes, and flocks of pigeons. But after a time, the unfortunate student is told that some classes – the singletons – have only a single member. Here is just cause for student protest, if ever there was one. This time, he has no 'many'. He has no elements or objects to be 'combined' or 'gathered together' into one, or to be 'thought of together as one'. Rather, he has just one single thing, the element, and he has another single thing, the singleton, and nothing he was told gives him the slightest guidance about what the one thing has to do with the other. Nor did any of those familiar examples concern single-membered sets. His introductory lesson just does not apply.

He might think: whatever it is that you do, in action or in thought, to make several things into a class, just do that same thing to a single thing and you make it into a singleton. How do you make several paintings into an art collection? Maybe you make a plan, you buy the paintings, you hang them in a special room, you even publish a catalogue. If you do the same thing, but your money runs out after you buy the first painting on your list, you have a collection that consists of a single painting. – But this thought is worse than useless. For all those allusions to human activity in the forming of sets are a bum steer. Sooner or later our student will hear that there are countless classes, most of them infinite and miscellaneous, so that the vast majority of them must have somehow got 'formed' with absolutely no attention or assistance from us. Maybe we've formed a general concept of classes, or a theory of them, or some sort of sketchy mental map of the whole of set-theoretical Reality. Maybe we've formed a few mental representations of a few very special classes. But there just cannot be anything that we've done to all the classes one at a time. The job is far too big for us. Must set theory rest on theology? – Cantor thought so!

We were told nothing about the nature of the singletons, and nothing about the nature of their relation to their elements. That might not be quite so bad if the singletons were a very special case. At least we'd know about the rest of the classes. But since all classes are fusions of singletons, and nothing over and above the singletons they're made of, our utter ignorance about the nature of the singletons amounts to utter ignorance about the nature of classes generally. We understand how bigger classes are composed of their singleton atoms. That's the easy part: just mereology. *That's* where we get the many into one, the combining or collecting or gathering. Those introductory remarks (apart from the misguided allusions to human activity) introduced us only to the *mereology* in set theory. But as to what is distinctively set-theoretical – the singletons that are the building blocks of all classes – they were entirely silent. Dr. Jekyll was there to welcome us. Mr. Hyde kept hidden. What do we know about singletons when we know only that they are atoms, and wholly distinct from the familiar individuals? What do we know about other classes, when we know only that they are composed of these atoms about which we know next to nothing?

Set theory has its unofficial axioms, traditional remarks about the nature of classes. They are never argued, but are passed along heedlessly from one author to another. Some are acceptable enough, I suppose: they do nothing to characterise the classes positively, but limit themselves to a *via negativa*. They say that such things as cats and quarks and spacetime regions should come out as individuals, not classes or mixed fusions.

Other unofficial axioms are bolder. One of them says that the classes are nowhere: they are outside of space and time. But why think this? Because we never see them or stumble over them? But maybe they are invisible and intangible. Maybe they can share their locations with other things. Maybe Possum's singleton is just where Possum is; maybe every singleton is just where its member is. Since members of singletons occupy extended spatiotemporal regions, and singletons are atoms, that would have to mean that something can occupy an extended region otherwise than by having different parts that occupy different parts of the region, and that would certainly be peculiar. But not more peculiar, I think, than being nowhere at all – we get a choice of equal evils, and cannot reject either hypothesis by pointing to the repugnancy of the other. I don't say the classes are in space and time. I don't say they aren't. I say we're in the sad fix that we haven't a clue whether they are or whether they aren't. We go much too fast from not knowing whether they are to thinking we know they are not.

Another unofficial axiom says that classes have nothing much by way of intrinsic character. That's not quite right: to be an atom, or to be a fusion of atoms, or to be a fusion of seventeen atoms, are matters of intrinsic character. However, these are not matters that distinguish one singleton from another, or one seventeen-membered class (a seventeen-fold fusion of singletons) from another. Are all singletons exact intrinsic duplicates? Or do they sometimes, or do they always, differ in their intrinsic character? If they do, do those differences in any way reflect differences between the character of their members? Do they involve any of the same qualities that distinguish individuals from one another? Again we haven't a clue.

Sometimes our offhand opinions about the nature of classes don't even agree with one another. When Nelson Goodman ([5], II.2) finds the notion of classes 'essentially incomprehensible' and refuses to 'use apparatus that peoples his world with a host of ethereal, platonic, pseudo entities' we should ask which are they: ethereal or platonic? The ether is everywhere, and one bit of it is pretty much like another; whereas the forms are nowhere, and each of them is unique. Ethereal entities are 'light, airy or tenuous', says the dictionary, whereas the forms are changeless and most fully real (whatever that means). If we knew better whether the classes were more fittingly called 'ethereal' or 'platonic', that would be no small advance!

I suppose we could bear up under our utter ignorance of the character and whereabouts (or lack thereof) of the singletons. Who ever said we could know everything? But there is worse to come. Because we know so little about the singletons, we are ill-placed even to begin to understand the relation of a thing to its singleton. We know what to call it, of course – membership – but that is all. Is it an external relation, like a relation of distance? An internal relation, like a relation of intrinsic similarity or difference? A combination of the two? Something else altogether?

It's no good saying that a singleton has *x* as its member because it shares the location of *x*. If singletons do share the location of their members, then *x* and *x*'s singleton and *x*'s singleton's singleton all three share a location; so *x* shares that location as much with one singleton as with the other. It's no good saying that a singleton has *x* as its member because of some sort of similarity between the singleton and *x*. For two perfect duplicates may have different singletons. It's no good saying that a singleton has *x* as its member because it has the property: being the singleton of *x*. That's just to go in a circle. We've named a property; but all we know about the property that bears this name is that it's the property, we know not what, that distinguishes the singleton of *x* from all other singletons.

It seems that we have no alternative but to suppose that the relation of member to singleton holds in virtue of qualities or external relations of which we have no conception whatsoever. Do we really understand what it means for a singleton to have a member?

Singletons, hence all classes, and worst of all the relation of membership, are profoundly mysterious. Mysteries are an onerous burden. Should we dump the burden by dumping the classes? If classes do not exist, we needn't puzzle over them. Renounce classes, and we are set free.

No; for set theory pervades modern mathematics. Some special branches and some special styles of mathematics can perhaps do without, but most of mathematics is into set theory up to its ears. If there are no classes, then there are no Dedekind cuts, there are no homeomorphisms, there are no complemented lattices, there are no probability distributions,. . . . For all these things are standardly defined as one or another sort of class. If there are no classes, then our mathematics textbooks are works of fiction, full of false 'theorems'. Renouncing classes means rejecting mathematics. That will not do. Mathematics is an established, going concern. Philosophy is as shaky as can be. To reject mathematics for philosophical reasons would be absurd. If we philosophers are sorely puzzled by the classes that constitute mathematical reality, that's our problem. We shouldn't expect mathematics to go away to make our life easier. Even if we reject mathematics gently – explaining how it can be a most useful fiction, 'good without being true' – we still reject it, and that's still absurd. Even if we hold onto some mutilated fragments of mathematics that can be reconstructed without classes, if we reject the bulk of mathematics that's still absurd.

That's not an argument, I know. But I laugh to think how *presumptuous* it would be to reject mathematics for philosophical reasons. How would *you* like to go and tell the mathematicians that they must change their ways, and abjure countless errors, now that *philosophy* has discovered that there are no classes? Will you tell them, with a straight face, to follow philosophical argument wherever it leads? If they challenge your credentials, will you boast of philosophy's other great discoveries: that motion is impossible, that a being than which no greater can be conceived cannot be conceived not to exist, that it is unthinkable that anything exists outside the mind, that time is unreal, that no theory has ever been made at all probable by evidence (but on the other hand that an empirically ideal theory can't possibly be false), that it is a wide-open scientific question whether anyone has ever believed anything, and so on, and on *ad nauseam?* Not me!

There is a way out of our dilemma. Or part-way out, and that may have to be good enough. We can take a 'structuralist' line about the theory of singleton functions; and derivatively, a structuralist line about set theory (like that of Paul Fitzgerald, [4]) and about all the mathematics that reduces to set theory. Even if we don't grasp *the* member-to-singleton function, we can still understand what it is to be *a* singleton function: a function that has the right formal character, and also obeys whatever 'unofficial' axioms we see fit to accept. So we might take set theory to be not the theory of *the* singleton function, plus mereology, but rather the general theory of all singleton functions. A set-theoretical truth would then have the covertly universal form: for any singleton function *s*, —*s*—*s*—. We need not say that any one singleton function has any special status.

Compare an algebraist's answer to a protesting student who says he hasn't been told what the Klein 4-group is just by being shown the table for it:

	e	*a*	*b*	*c*
e	*e*	*a*	*b*	*c*
a	*a*	*e*	*c*	*b*
b	*b*	*c*	*e*	*a*
c	*c*	*b*	*a*	*e*

What *are* these four things *e, a, b, c?* The Prof may answer: 'They're anything you like. No one thing is *the* Klein 4-group; rather, any function (or equivalently, any four-things-and-a-function) that obeys the table is *a* Klein 4-group. Anything I tell you about "the" Klein 4-group is tacitly general. For instance, when I said that any permutation of the non-identity elements is an automorphism, I meant this to go for any Klein 4-group, no matter what its elements might be. And there do exist Klein 4-groups' – and here he offers an example or two – 'so there's no fear that generalisations about Klein 4-groups will come out vacuously true'.

Similarly, we might well be attracted to a 'structuralist' philosophy of arithmetic. It says that there's no one sequence that is *the* number sequence; rather, arithmetic is the general theory of omega-sequences. Each sequence has its own zero, its own successor function, and so on. The successor function and the sequence are interdefinable, so equivalently we could say that arithmetic is the general theory of successor functions. A successor function is characterised by means of the Peano axioms, as follows.

A *successor function* is any unary one-one function s such that

(1) the domain of s consists of its range and one other thing (called the *s-zero*); and
(2) all things in the domain of s are generated from the s-zero by iterated applications of s.

(Clause (2) means: if there are some things, and the s-zero is one of them, and when x is one of them so is $s(x)$, then everything in the domain of s is one of them.) Arithmetical truths are general; they apply to all successor functions. An arithmetical truth has the covertly universal form: for any successor function s, —s—s—. There do exist successor functions, as witness the familiar set-theoretical models of arithmetic: the Zermelo numbers, the von Neumann numbers, etc. So again there is no fear that our generalisations will turn out vacuous and make the wrong things come out as arithmetical truths.

(By 'structuralism' I don't mean to suggest that the subject matter of arithmetic is some interesting entity, an 'abstract structure', common to all the many successor functions. I suspect such entities are trouble, but in any case, they're an optional extra. We needn't believe in 'abstract structures' to have general structural truths about all successor functions.)

Similarly, *mutatis mutandis* for singleton functions. We can characterise them by means of a modified, mereologized form of the Peano axioms:

A *singleton function* is any unary one-one function s such that

(0) the range of s consists of atoms (called *s-singletons*);
(1) the domain of s consists of all small fusions of s-singletons together with all things (called *s-individuals*) that have no s-singletons as parts; and
(2) all things are generated from the s-individuals by iterated applications of s and of fusion.

(Something is *small* iff its atoms do not correspond one-one with all the atoms. Clause (2) means: if there are some things, and every s-individual is one of them, and when x is one of them so is $s(x)$, and every fusion of some of them is one of them, then everything is one of them.)

Suppose we have some 'unofficial axioms' (credible ones, not the overbold ones that say the sets lack whereabouts and character) that tell us which are the things – cats, puddles, spacetime points, souls if there are any – that ought to turn out to be individuals. Suppose also that these axioms endorse our Division, Priority, and Fusion Theses. Then we may call a singleton function s *correct* iff its demarcation of individuals – that is, its division between s-individuals and fusions of s-singletons – falls just where the unofficial axioms say it should. Since we have decided that the null set shall be the fusion of all individuals, we can also put it this way: the unofficial axioms determine what ought to be the null set, and a correct singleton function is one that defines the null set in agreement with that determination. Let us ignore all other singleton functions, and say that set theory consists of the general truths about all correct singleton functions.

Structuralist set theory is nominalistic set theory, in the special sense of Goodman ([5], II.3). It has no set-theoretical primitive, neither the general relation of membership nor the special case of membership in singletons. All the 'combining' or 'collecting' or 'gathering together' in set theory is purely mereological. But we do not renounce classes. We still have things that are classes relative to all the many correct singleton functions. Although these functions disagree about which classes have which members, at least they agree about which things are the classes and which are the individuals.

If we go structuralist about singleton functions, we may bid farewell to all our worries about how we understand 'the' singleton function. That which is not there need not be understood. Sad to say, we do *not* bid farewell to our lamentable ignorance of the whereabouts and character of the classes.

If we go structuralist, do we rebel against established set theory, and all of set-theoretical mathematics? Well, we challenge no proofs and we deny no theorems. But we do rebuke the mathematicians for a foundational error: they think they've fixed on one particular singleton function when really all they've got is the general notion of *a* singleton function. Even this much of a philosophical challenge to established mathematics is presumptuous and suspect. It would be much better if we could find a way to take mathematics just as we find it. But I have no way to offer, and the structuralist reinterpretation is the best fallback I know.

You may smell a rat. Here we are, quantifying over functions. How can we do that before we already have set theory, where 'having' set theory includes understanding its primitive notions? Isn't a function a class of ordered pairs? And aren't ordered pairs, in turn, set-theoretical constructions: Kuratowski pairs, or something similar? And when I characterised singleton functions, didn't I quantify over relations again when I spoke of one-one correspondence? And wasn't it set-theoretic modelling that gave us our needed assurance that there did exist Klein 4-groups and successor functions, and don't we need a parallel assurance that there exist singleton functions? In short, structuralism about set theory seems to presuppose set theory. Only if we don't need it can we have it.

Recent work by John P. Burgess and A. P. Hazen saves the day. They've shown how we can simulate quantification over relations using only the framework of megethology: mereology plus plural quantification. Roughly speaking, a quantifier over relations is a plural quantifier over things that encode ordered pairs by mereological means.

Here we shall take a hybrid method, beginning Hazen's way and finishing Burgess's way. We must assume that there are infinitely many atoms; at this stage of the game, we can express that by saying that there are some things, each of which is part of another. For simplicity, we shall also assume that Reality consists entirely of atoms. (But if it did not, indeed even if there were no atoms at all, we could assume instead just that there are infinitely many mutually non-overlapping things – call them 'quasi-atoms' – whose fusion is all of Reality. Then all that follows could be unchanged, except that 'atom' would mean 'quasi-atom' and all quantifications would be tacitly restricted to things that were fusions of quasi-atoms.)

First step (Zermelo). We can encode orderings of atoms. Given some things O, atom b *O-precedes* atom c iff whenever c is part of one of O then b is too, but not conversely. Say that O *order the atoms* iff O-precedence is transitive, asymmetric, and connected. We assume as a principle of megethology that some O order the atoms.

Second step (Hazen). We can encode relations of atoms. We can say that O, A, B, C *relate atoms* iff O order the atoms, A are some diatoms (two-atom fusions), B are some diatoms, and C are some atoms. Given some such O, A, B, C, say that atom b is *OABC-related* to atom c iff either b O-precedes c and $b + c$ is one of A, or c O-precedes b and $b + c$ is one of B, or b and c are identical and b is one of C. The idea is that the extraneous ordering O takes the place of the ordering normally built into ordered pairs, so that we can get by with *un*ordered diatoms. (We would need to make special provision for some relations: in the case of the identity relation, for instance, both A and B would be missing. But we can encode the two relations needed for the next step, and that is enough.) Consider a special case; O, A, B, C *map all atoms one-one* iff O, A, B, C relate atoms; any atom is $OABC$-related to exactly one atom, called its *OABC-image;* and no two atoms have the same $OABC$-image. We define the *OABC-image* of anything as the fusion of $OABC$-images of its atoms.

Third step (Burgess). We can encode ordered pairs. We assume as a principle of megethology that if there are infinitely many atoms, then there are O, A, B, C, D, E, F such that O, A, B, C map all atoms one-one, and so do O, D, E, F; and such that no $OABC$-image ever overlaps any $ODEF$-image. Thus all of Reality is mapped into two non-overlapping microcosms, and every part of Reality has an image in each microcosm. The *ordered pair* of x and y (with respect to $O \ldots F$) is the fusion of the $OABC$-image of x and the $ODEF$-image of y. We recover the first term of an ordered pair z as the fusion of all atoms whose $OABC$-images are parts of z; and the second term as the fusion of all atoms whose $ODEF$-images are parts of z.

Final step. We simulate a quantifier over relations (binary relations will suffice for present purposes) by a plural quantifier over ordered pairs, preceded by a string of plural quantifiers over the wherewithal for decoding such pairs. Example. 'For some relation r, ——$r(x,y)$ ——' becomes: 'For some $O \ldots F$ meeting the conditions above, for some R such that each one of R is an ordered pair (with respect of $O \dots F$), ——the ordered pair (with respect to $O \dots F$) of x and y is one of R——.' The case of a universal quantifier over relations is similar, except that all the quantifiers in the string are universal.

Megethology therefore suffices as a framework for structuralist set theory. Once we have simulated quantification over relations, we can define singleton functions; we can quantify over them; and with respect to any given one of them, we can define membership, classhood, etc. as above.

Now we shall formulate some hypotheses about the size of Reality, and thereby see how megethology deserves its name. Recall that something is *small* iff its atoms correspond one-one with some but not all the atoms; otherwise *large*. Similarly in the plural: some things are *few* iff they correspond one-one with some but not all the atoms, otherwise *many;* and they are *barely many* iff they correspond one-one with all the atoms. An *infinite* thing is one whose atoms correspond one-one with only some of its atoms.

Hypothesis U. *The fusion of a few small things is small.*

Hypothesis P−. *The parts of a small thing are few or barely many.*

Hypothesis P. *The parts of a small thing are few.*

Hypothesis I. *Something small is infinite.*

Hypotheses U and P− together imply this

Existence Thesis. *For any small thing, n, there is a singleton function s such that n is the fusion of the s-individuals.*

First Step. Given Hypotheses U and P−, we have that the small parts of Reality are barely many. Proof of the first step: Assume, as a principle of megethology, that we have a well-ordering of all the atoms. If so, then we have an *initial* well-ordering of all the atoms: that is, one in which each atom is preceded by only a few others. For, given a well-ordering that isn't already initial, we take the first atom that is preceded by many others. Imaging under the one-one correspondence between these many preceding atoms and all the atoms, we have a new well-ordering of all the atoms; and this new well-ordering, being isomorphic to a segment of the old one in which each atom is preceded by only a few others, is initial. Atom *b* ***bounds*** *x* iff every atom of *x* precedes *b* in our initial well-ordering. Any small *x* is bounded: for any atom *c* of *x*, let [*c*] be the fusion of *c* and all atoms that precede *c*; these [*c*]'s are small and few; their fusion is small, by Hypothesis U; some atom falls outside that fusion, and it bounds x. If any atom bounds *x*, there is a first atom that bounds *x*, called the ***limit*** of *x*. Any small *x* has a limit. When *c* is the limit of *x*, *x* is part of [*c*]. By P − , since [*c*] is small, [*c*]'s parts are at most barely many. So each of barely many atoms is the limit of at most barely many small things. So, by a principle of megethology akin to the multiplication rule for cardinal numbers, there are at most barely many small things. And each of barely many atoms is small, so there are at least barely many small things.

Second Step. Given that the small parts of Reality are barely many, we have that for any small thing *n*, there exists a unary one-one function *f* such that (0) the range of *f* consists of atoms of −*n*; (1) the domain of *f* consists of all small parts of −*n* together with all parts of *n*. Proof of the second step: By a principle of megethology, all parts of *n* are themselves small. So the small parts of −*n* together with the parts of *n* are some of the small things, so they are at most barely many. And they are at least barely many, since all atoms are either small parts of −*n* or parts of *n*. By another principle of megethology and the infinity of atoms, since n is small, −*n* is large. So the domain and range of the desired function both stand in one-one correspondence with all the atoms, so they stand in one-one correspondence with each other.

Third Step (due partly to Burgess). Given *f* as specified in the second step, there is a singleton function *s* such that the *s*-individuals are exactly the parts of *n*. Proof of the third step: Though *f* needn't be a singleton function, call $f(y)$ the ****-singleton*** of *y*; define ****-classes***, ****-sets***, and ****-membership*** accordingly. Let *G* be the things generated from the parts of *n* by iterated applications of *f* and of fusion. Let *H* be the *-singletons among *G*. Now consider the *-class of all *-sets among *G* that are non-self-*-members. It cannot itself be a *-set among *G*, by Russell's paradox. But if it were small, it would be a *-set among *G*. So it is large. So some of *H* are many. We saw that −*n* is large; so *H* stand in one-one correspondence with all the atoms of −*n*. Extend this to a correspondence that also maps all atoms of n to themselves. The image of f under this correspondence is the desired function s. QED

That is, there is a correct singleton function, no matter what small thing the unofficial axioms may deem to be the null set. And we have this

Categoricity Thesis. *For any small thing n, if we have two singleton functions s and t such that n is the fusion of the s-individuals and also the fusion of the t-individuals, then s and t differ only by a permutation of singletons.*

Proof. Define relation p as follows: if x is an individual (part of n), $s(x)$ bears p to $t(x)$; if x and y are small fusions of singletons (atoms of $-n$), and every atom of x bears p to some atom of y, and for every atom of y, some atom of x bears p to it, then $s(x)$ bears p to $t(y)$. Call a singleton *good*-1 iff it bears p to exactly one singleton; and call anything *good*-1 iff every singleton that is part of is is good-1. Likewise, call a singleton *good*-2 iff exactly one singleton bears p to it; and call anything *good*-2 iff every singleton that is part of it is good-2. Using clause (2) in the definition of a singleton function, we can show by induction that everything is good-1, and that everything is good-2. So p is a permutation of singletons; and t is the image of s under p. QED

Take any mathematical – that is, set-theoretical – sentence. Its only vocabulary, after we eliminate defined terms, will be logical and mereological vocabulary, and 'singleton'. Replace 'singleton' by a variable to obtain a formula —*s*—*s*—. The original sentence is mathematically true iff this formula holds for all correct singleton functions; mathematically false iff it holds for none; and indeterminate iff it holds for some but not others. The Existence Thesis guarantees that nothing is both mathematically true and mathematically false. The Categoricity Thesis guarantees that there will be no indeterminacy; for if two singleton functions differ only by a permutation,then both or neither will satisfy the formula. Despite our structuralist reinterpretation, mathematical sentences come out bivalent, just as if we had been able to fix on one particular singleton function.

Given our definitions and some principles of megethology, we can show that all but two of the standard axioms of set theory hold for any singleton function: Null Set, Extensionality, Pair Sets, *Aussonderung,* Replacement, *Fundierung,* Choice, and Unions. We get the last two standard axioms only conditionally: Power Sets given Hypothesis P, and Infinity given Hypothesis I.

Proof. Null Set holds because, by definition, there are individuals; their fusion exists, by the principle of Unrestricted Composition; and by our definitions it is a set with no members.

Extensionality holds in virtue of a principle of Uniqueness of Composition: it never happens that the same things have two different fusions. Apply it to classes: there are never two fusions of the same singletons, hence no coextensive classes. Apply it to individuals: there are not two fusions of the individuals, hence there is only one null set.

Pair Sets follows from Unrestricted Composition, plus the fact that there are at least three atoms (singleton of the null set, singleton of that, singleton of that) and hence any fusion of two singletons is small.

Aussonderung follows from Unrestricted Composition plus the fact that any part of a small thing is small (the Null Set axiom covers the case in which the required set is empty).

Replacement follows from a corresponding principle of megethology: if there is a function from atoms of x to all atoms of y, then y is small if x is.

Fundierung. Something is *grounded* iff it belongs to no class that intersects each of its own members and thereby violates *Fundierung.* Using clause (2) of the definition of a singleton function, we can show by induction that everything is grounded, hence nothing violates *Fundierung.*

Choice follows from a corresponding principle of megethology: given some non-overlapping things, something shares exactly one atom with each of them.

Unions. First we need a Lemma: the existence of a singleton function implies Hypothesis U. Proof of the Lemma: If not, we have a few small things *R* such that their fusion is not small. Then the atoms of the fusion of R correspond one-one with all the atoms. Let the things S be the images of R under this correspondence; then S are a few small things, but their fusion is all of Reality. Given a singleton function, let the things T be as follows: whenever a class or individual is one of S, it is one of T; whenever a mixed fusion is one of S, its largest individual part is one of T and its largest class part is one of T. T also are a few small things, their fusion is all of Reality, and further each one of them has a singleton. Now we can encode anything x, unambiguously, by an atom: first we have the intersections of x with each of T, and these are a few individuals and sets; then we have the set of these intersections; and finally we have the singleton of that set. So there are no more fusions of atoms than atoms, which is impossible by the reasoning of Cantor's theorem.

Now, take any set. If it has no classes as members, its union is the null set. Otherwise we have its union class, by Unrestricted Composition; the Lemma gives us Hypothesis U, whereby this class is small and hence a set.

Power Sets. For any set, we have its power class, by Unrestricted Composition. Given Hypothesis P, this class is small, hence a set.

Infinity. Given Hypothesis I, something infinite is small. We have the class of its atoms, by Unrestricted Composition and the fact that all atoms are either individuals or singletons, and this class is an infinite set. QED

So to regain set theory we need to assume Hypotheses U, P, and I. (We needn't list P– separately, since it follows from P.) These constrain the size of Reality, as measured by the total number of atoms. It is easy to see how any two of the constraints can hold. U and P hold, but I fails, if there are countably many atoms, so that 'small' means 'finite'. U and I hold, but P fails, if there are aleph-one atoms, so that 'small' means 'countable'. P and I hold, but U fails, if there are beth-omega atoms. Making all three hold together is harder. That takes a (strongly) 'inaccessible' infinity of atoms – an infinity that transcends our commonplace alephs and beths in much the same way that any infinity transcends finitude. There will be inaccessibly many atoms, inaccessibly many singletons, and inaccessibly many sets.

Do you find it extravagant to posit so many things? Especially when you know nothing about their whereabouts and character? Beware! The inaccessibly many atoms are not the wages of a newfangled mereological-cum-structuralist reconstruction of set theory. It is orthodox set theory itself that incurs the commitment to an inaccessible domain of sets. Like this: (1) The largest classes are proper classes, smaller ones are sets. (2) The size of the proper classes cannot be reached from below by taking unions: the union of a set of sets is still a mere set. (3) The size of the proper classes cannot be reached from below by taking powers: the class of all subsets of a set is still a mere set. (4) The size of the proper classes is not the smallest infinite size: some mere set is infinite. So says orthodoxy. I have faithfully reconstructed this aspect of orthodoxy along with the rest. Will you tell the mathematicians to abjure their errors? Not me!

REFERENCES

1 George Boolos, 'To be is to be the value of a variable (or to be some values of some variables)', *Journal of Philosophy*, **81** (1984), 430–49.

2 ———, 'Nominalist Platonism', *Philosophical Review*, **94**, (1985), 327–44.

3 John P. Burgess, A. P. Hazen, David Lewis, 'Appendix on pairing' in [6].

4 Paul Fitzgerald, 'Meaning in science and mathematics', *PSA 1974*, R. S. Cohen *et al.*, eds. (Reidel, 1976), section IV.

5 Nelson Goodman, *The structure of appearance* (Harvard University Press, 1951).

6 David Lewis, *Parts of classes* (Blackwell, 1991).