Papers in metaphysics and epistemology

DAVID LEWIS

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521582483

First published 1999

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Lewis, David K., 1941–
Papers in metaphysics and epistemology / David Lewis.
p. cm. – (Cambridge studies in philosophy)
A collection of the author's papers, most of which were previously published.
Includes bibliographical references and index.
ISBN 0-521-58248-2 (hardbound). – ISBN 0-521-58787-5 (pbk.)
1. Metaphysics. 2. Knowledge, theory of. I. Title. II. Series.
BD111.L45 1999 98-25689
110 – dc21 CIP

ISBN-13 978-0-521-58248-3 hardback
ISBN-10 0-521-58248-2 hardback

ISBN-13 978-0-521-58787-7 paperback
ISBN-10 0-521-58787-5 paperback

Transferred to digital printing 2006

For the philosophers, past and present,
of Sydney and Canberra

9

Many, but almost one

THE PROBLEM OF THE MANY

Think of a cloud – just one cloud, and around it clear blue sky. Seen from the ground, the cloud may seem to have a sharp boundary. Not so. The cloud is a swarm of water droplets. At the outskirts of the cloud the density of the droplets falls off. Eventually they are so few and far between that we may hesitate to say that the outlying droplets are still part of the cloud at all; perhaps we might better say only that they are near the cloud. But the transition is gradual. Many surfaces are equally good candidates to be the boundary of the cloud. Therefore many aggregates of droplets, some more inclusive and some less inclusive (and some inclusive in different ways than others), are equally good candidates to be the cloud. Since they have equal claim, how can we say that the cloud is one of these aggregates rather than another? But if all of them count as clouds, then we have many clouds rather than one. And if none of them count, each one being ruled out because of the competition from the others, then we have no cloud. How is it, then, that we have just one cloud? And yet we do.

This is Unger's (1980) 'problem of the many'. Once noticed, we can see that it is everywhere, for all things are swarms of particles.

First published in Keith Campbell, John Bacon, and Lloyd Reinhardt (eds.), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong* (Cambridge, Cambridge University Press, 1993.

There are always outlying particles, questionably parts of the thing, not definitely included and not definitely not included. So there are always many aggregates, differing by a little bit here and a little bit there, with equal claim to be the thing. We have many things or we have none, but anyway not the one thing we thought we had. That is absurd.

Think of a rusty nail, and the gradual transition from steel, to steel with bits of rust scattered through, to rust adhering to the nail, to rust merely resting on the nail. Or think of a cathode, and its departing electrons. Or think of anything that undergoes evaporation or erosion or abrasion. Or think of yourself, or any organism, with parts that gradually come loose in metabolism or excretion or perspiration or shedding of dead skin. In each case, a thing has questionable parts, and therefore is subject to the problem of the many.

If, as I think, things perdure through time by having temporal parts, then questionable temporal parts add to the problem of the many. If a person comes into existence gradually (whether over weeks or over years or over nanoseconds doesn't matter for our present purpose) then there are questionable temporal parts at the beginning of every human life. Likewise at the end, even in the most sudden death imaginable. Do you think you are one person? – No, there are many aggregates of temporal parts, differing just a little at the ends, with equal claim to count as persons, and equal claim to count as you. Are all those equally good claims good enough? If so, you are many. If not, you are none. Either way we get the wrong answer. For undeniably you are one.

If, as some think but I do not,[1] ordinary things extend through other possible worlds, then the problem of the many takes on still another dimension. Here in this world we have a ship, the *Enigma*; there in another world is a ship built at about the same time, to plans that are nearly the same but not quite, using many of the same planks and some that are not the same. It is questionable whether the ship in that other world is *Enigma* herself, or just a substitute. If *Enigma* is a thing that extends through worlds, then the question is whether *Enigma* includes as a part what's in that other world. We have two versions of *Enigma*, one that includes this questionable other-worldly

1 See Lewis (1986a, pp. 210–20).

part and one that excludes it. They have equal claim to count as ships, and equal claim to count as *Enigma*. We have two ships, coinciding in this world but differing in their full extent. Or else we have none; but anyway not the one ship we thought we had.

THE PARADOX OF 1001 CATS

Cat Tibbles is alone on the mat. Tibbles has hairs $h_1, h_2, \ldots, h_{1000}$. Let c be Tibbles including all these hairs; let c_1 be all of Tibbles except for h_1; and similarly for $c_2, \ldots, c_{1000}$. Each of these c's is a cat. So instead of one cat on the mat, Tibbles, we have at least 1001 cats – which is absurd. This is P. T. Geach's (1980, pp. 215–16) paradox of 1001 cats.

Why should we think that each c_n is a cat? Because, says Geach, 'c_n would clearly be a cat were the hair h_n plucked out, and we cannot reasonably suppose that plucking out a hair *generates* a cat, so c_n must already have been a cat' (p. 215). This need not convince us. We can reply that plucking out h_n turns c_n from a mere proper part of cat Tibbles into the whole of a cat. No new cat is generated, since the cat that c_n becomes the whole of is none other than Tibbles. Nor do c_n and Tibbles ever become identical *simpliciter* – of course not, since what's true about c_n's past still differs from what's true about Tibbles's past. Rather, c_n becomes the whole of cat Tibbles in the sense that c_n's post-plucking temporal part is identical with Tibbles's post-plucking temporal part. So far, so good; except for those, like Geach, who reject the idea of temporal parts. The rest of us have no paradox yet.

But suppose it is spring, and Tibbles is shedding. When a cat sheds, the hairs do not come popping off; they become gradually looser, until finally they are held in place only by the hairs around them. By the end of this gradual process, the loose hairs are no longer parts of the cat. Sometime before the end, they are questionable parts: not definitely still parts of the cat, not definitely not. Suppose each of $h_1, h_2, \ldots, h_{1000}$ is at this questionable stage. Now indeed all of $c_1, c_2, \ldots, c_{1000}$, and also c which includes all the questionable hairs, have equal claim to be a cat, and equal claim to be Tibbles. So now we have 1001 cats. (Indeed, we have many more than that. For instance there is the

cat that includes all but the four hairs h_6, h_{408}, h_{882}, and h_{907}.) The paradox of 1001 cats, insofar as it is a real paradox, is another instance of Unger's problem of the many.

To deny that there are many cats on the mat, we must either deny that the many are cats, or else deny that the cats are many. We may solve the paradox by finding a way to disqualify candidates for cathood: there are the many, sure enough, but the many are not all cats. At most one of them is. Perhaps the true cat is one of the many; or perhaps it is something else altogether, and none of the many are cats. Or else, if we grant that all the candidates are truly cats, we must find a way to say that these cats are not truly different from one another. I think both alternatives lead to successful solutions, but we shall see some unsuccessful solutions as well.

TWO SOLUTIONS BY DISQUALIFICATION: NONE OF THE MANY ARE CATS

We could try saying that not one of the *c*'s is a cat; they are many, sure enough, but not many cats. Tibbles, the only genuine cat on the mat, is something else, different from all of them.

One way to disqualify the many is to invoke the alleged distinction between things and the parcels of matter that constitute them. We could try saying that the *c*'s are not cats. Rather, they are cat-constituting parcels of matter. Tibbles is the cat that each of them constitutes.[2]

This dualism of things and their constituters is unparsimonious and unnecessary. It was invented to solve a certain problem, but a better solution to that problem lies elsewhere, as follows. We know that the matter of a thing may exist before and after the thing does; and we know that a thing may gain and lose matter while it still exists, as a cat does, or a wave or a flame. The dualists conclude that the matter is not the thing; constitution is not identity; there are things, there are the parcels of matter that temporarily constitute those things; these are items of two different categories, related by the special relation of

2 This is the solution advanced in Lowe (1982).

constitution. We must agree, at least, that the temporally extended thing is not the temporally extended parcel of matter that temporarily constitutes that thing. But constitution may be identity, all the same, if it is identity between temporal parts. If some matter constitutes a cat for one minute, then a minute-long temporal segment of the cat is identical to a minute-long temporal segment of the matter. The cat consists entirely of the matter that constitutes it, in this sense: The whole of the cat, throughout the time it lives, consists entirely of temporal segments of various parcels of matter. At any moment, if we disregard everything not located at that moment, the cat and the matter that then constitutes it are identical.[3] So only those who reject the notion of temporal parts have any need for the dualism of things and constituters. But suppose we accept it all the same. At best, this just transforms the paradox of 1001 cats into the paradox of 1001 cat-constituters. Is that an improvement? We all thought there was only one cat on the mat. After distinguishing Tibbles from her constituter, would we not still want to think there was only one cat-constituter on the mat?

Further, even granted that Tibbles has many constituters, I still question whether Tibbles is the only cat present. The constituters are cat-like in size, shape, weight, inner structure, and motion. They vibrate and set the air in motion – in short, they purr (especially when you pat them). Any way a cat can be at a moment, cat-constituters also can be; anything a cat can do at a moment, cat-constituters also can do. They are all too cat-like not to be cats. Indeed, they may have unfeline pasts and futures, but that doesn't show that they are never cats; it only shows that they do not remain cats for very long. Now we have the paradox of 1002 cats: Tibbles the constituted cat, and also the 1001 all-too-feline cat-constituters. Nothing has been gained.

3 The dualism of things and their constituters is also meant to solve a modal problem: Even at one moment, the thing might have been made of different matter, so what might have been true of it differs from what might have been true of its matter, so constitution cannot be identity. This problem too has a better solution. We should allow that what is true of a given thing at a given world is a vague and inconstant matter. Conflicting answers, equally correct, may be evoked by different ways of referring to the same thing, e.g., as cat or as cat-constituter. My counterpart theory affords this desirable inconstancy; many rival theories do also. See Lewis (1986a, pp. 248–63).

I conclude that invoking the dualism of cats and cat-constituters to solve the paradox of 1001 cats does not succeed.

A different way to disqualify the many appeals to a doctrine of vagueness in nature. We could try saying that cat Tibbles is a vague object, and that the *c*'s are not cats but rather alternative precisifications of a cat.

In one way, at least, this solution works better than the one before. This time, I cannot complain that at best we only transform the paradox of 1001 cats into the paradox of 1001 cat-precisifications, because that is no paradox. If indeed there are vague objects and precisifications, it is only to be expected that one vague object will have many precisifications.

If the proposal is meant to solve our paradox, it must be meant as serious metaphysics. It cannot just be a way of saying 'in the material mode' that the words 'Tibbles' and 'cat' are vague, and that this vagueness makes it indefinite just which hairs are part of the cat Tibbles. Rather, the idea must be that material objects come in two varieties, vague and precise; cats are vague, the *c*'s are precise, and that is why none of the *c*'s is a cat.

This new dualism of vague objects and their precisifications is, again, unparsimonious and unnecessary. The problem it was made to solve might better be solved another way. It is absurd to think that we have decided to apply the name 'Tibbles' to a certain precisely delimited object; or that we have decided to apply the term 'cat' to each of certain precisely delimited objects. But we needn't conclude that these words must rather apply to certain *im*precisely delimited, vague objects. Instead we should conclude that we never quite made up our minds just what these words apply to. We have made up our minds that 'Tibbles' is to name one or another Tibbles-precisification, but we never decided just which one; we decided that 'cat' was to apply to some and only some cat-precisifications, but again we never decided just which ones. (Nor did we ever decide just which things our new-found terms 'Tibbles-precisification' and 'cat-precisification' were to apply to.) It was very sensible of us not to decide. We probably couldn't have done it if we'd tried; and even if we could have, doing it would have been useless folly.

Semantic indecision will suffice to explain the phenomenon of vagueness.[4] We need no vague objects.

Further, I doubt that I have any correct conception of a vague object. How, for instance, shall I think of an object that is vague in its spatial extent? The closest I can come is to superimpose three pictures. There is the *multiplicity* picture, in which the vague object gives way to its many precisifications, and the vagueness of the object gives way to differences between precisifications. There is the *ignorance* picture, in which the object has some definite but secret extent. And there is the *fadeaway* picture, in which the presence of the object admits of degree, in much the way that the presence of a spot of illumination admits of degree, and the degree diminishes as a function of the distance from the region where the object is most intensely present. None of the three pictures is right. Each one in its own way replaces the alleged vagueness of the object by precision. But if I cannot think of a vague object except by juggling these mistaken pictures, I have no correct conception.[5]

I can complain as before that we end up with a paradox of 1002 cats: Tibbles the vague cat, and also the 1001 precise cats. Once again,

4 Provided that there exist the many precisifications for us to be undecided between. If you deny this, you will indeed have need of vague objects. See van Inwagen (1990, pp. 213–83).

5 I grant that the hypothesis of vague objects, for all its faults, can at least be made consistent. If there are vague objects, no doubt they sometimes stand in relations of 'vague identity' to one another. We might think that when a and b are vaguely identical vague objects, the identity statement $a = b$ suffers a truth-value gap; but in fact this conception of vague identity belongs to the theory of vagueness as semantic indecision. As Gareth Evans showed, it doesn't mix with the idea that vague identity is due to vagueness in nature. For if a and b are vaguely identical, they differ in respect of vague identity to a; but nothing, however peculiar it may be, differs in any way from itself; so the identity $a = b$ is definitely false. See Evans (1978). (Evans's too-concise paper invites misunderstanding, but his own testimony confirms my interpretation. See Lewis 1988.) To get a consistent theory of vague objects, different from the bastard theory that is Evans's target, we must disconnect 'vague identity' from truth-value gaps in identity statements. Even if $a = b$ is definitely false, a and b can still be 'vaguely identical' in the sense of sharing some but not all of their precisifications.

the cat-precisifications are all too cat-like. More so than the cat-constituters, in fact: The precisifications are cat-like not just in what they can do and how they can be at a moment, but also over time. They would make good pets – especially since 1001 of them will not eat you out of house and home!

Don't say that the precisifications cannot be cats because cats cannot be precise objects. Surely there could be cats in a world where nature is so much less gradual that the problem of the many goes away. It could happen that cats have no questionable parts at all, neither spatial nor temporal. (In this world, when cats shed in the spring, the hairs *do* come popping off.) So it is at least possible that cat-like precise objects are genuine cats. If so, how can the presence of one vague cat spoil their cathood?

I conclude that invoking the dualism of vague objects and their precisifications to solve the paradox of 1001 cats does not succeed.

A BETTER SOLUTION BY DISQUALIFICATION: ONE OF THE MANY IS A CAT

Since all of the many are so cat-like, there is only one credible way to deny that all of them are cats. When is something very cat-like, yet not a cat? – When it is just a little less than a whole cat, almost all of a cat with just one little bit left out. Or when it is just a little more than a cat, a cat plus a little something extra. Or when it is both a little more and a little less.

Suppose we say that one of our many is exactly a cat, no more and no less; and that each of the rest is disqualified because it is a little less than a cat, or a little more, or both more and less. This invokes no unparsimonious and unnecessary dualisms; it disqualifies all but one of the many without denying that they are very cat-like; it leaves us with just one cat. All very satisfactory.

The trouble, so it seems, is that there is no saying which one is a cat. That is left altogether arbitrary. Settling it takes a semantic decision, and that is the decision we never made (and shouldn't have made, and maybe couldn't have made). No secret fact could answer the question, for we never decided how the answer would depend on secret facts.

Which one deserves the name 'cat' is up to us. If we decline to settle the question, nothing else will settle it for us.[6]

We cannot deny the arbitrariness. What we can deny, though, is that it is trouble. What shall we do, if semantic indecision is inescapable, and yet we wish to carry on talking? The answer, surely, is to exploit the fact that very often our unmade semantic decisions don't matter. Often, what you want to say will be true under all different ways of making the unmade decision. Then if you say it, even if by choice or by necessity you leave the decision forever unmade, you still speak truthfully. It makes no difference just what you meant, what you say is true regardless. And if it makes no difference just what you meant, likewise it makes no difference that you never made up your mind just what to mean. You say that a famous architect designed Fred's house; it never crossed your mind to think whether by 'house' you meant something that did or that didn't include the attached garage; neither does some established convention or secret fact decide the issue; no matter, you knew that what you said was true either way.

This plan for coping with semantic indecision is van Fraassen's (1966) method of *supervaluations*. Call a sentence *super-true* if and only if it is true under all ways of making the unmade semantic decisions; *super-false* if and only if it is false under all ways of making those decisions; and if it is true under some ways and false under others, then it suffers a super-truth-value gap. Super-truth, with respect to a language interpreted in an imperfectly decisive way, replaces truth *simpliciter* as the goal of a cooperative speaker attempting to impart information. We can put it another way: Whatever it is that we do to determine the 'intended' interpretation of our language determines not one interpretation but a range of interpretations. (The range depends on context, and is itself somewhat indeterminate.) What we try for, in imparting information, is truth of what we say under all the intended interpretations.

6 I do not think reference is entirely up to our choice. Some things are by their nature more eligible than others to be referents or objects of thought, and when we do nothing to settle the contest in favour of the less eligible, then the more eligible wins by default; see Lewis (1984). That's no help here: nature is gradual, no handy joint in nature picks out one of the *c*'s from all the rest.

Each intended interpretation of our language puts one of the cat-candidates on the mat into the extension of the word 'cat', and excludes all the rest. Likewise each intended interpretation picks out one cat-candidate, the same one, as the referent of 'Tibbles'. Therefore it is super-true that there is just one cat, Tibbles, on the mat. Because it is super-true, you are entitled to affirm it. And so you may say what you want to say: there is one cat. That is how the method of supervaluations solves the paradox of 1001 cats.

Objection. Just one of the candidates is a cat, no more and no less. But don't try to say which one it is. Nothing you might say would be super-true. For it is exactly this semantic decision that remains unmade; it is exactly in this respect that the intended interpretations differ. Although it is super-true that something is a cat on the mat, there is nothing such that it is super-true of it that *it* is a cat on the mat. (It's like the old puzzle: I owe you a horse, but there's no horse such that I owe you that horse.) This is peculiar.

Reply. So it is. But once you know the reason why, you can learn to accept it.

Objection.[7] Supervaluationism works too well: it stops us from ever stating the problem in the first place. The problem supposedly was that all the many candidates had equal claim to cathood. But under the supervaluationist rule, that may not be said. For under any one way of making the unmade decision, one candidate is picked as a cat. So under any one way of making the decision, the candidates do *not* have equal claim. What's true under all ways of making the decision is super-true. So what's super-true, and what we should have said, is that the candidates do *not* have equal claim. Then what's the problem? And yet the problem was stated. So supervaluationism is mistaken.

Reply. What's mistaken is a fanatical supervaluationism, which automatically applies the supervaluationist rule to any statement whatever, never mind that the statement makes no sense that way. The rule

7 Here I'm indebted to remarks of Saul Kripke many years ago. At his request, I note that what I have written here may not correspond exactly to the whole of what he said on that occasion.

should instead be taken as a defeasible presumption. What defeats it, sometimes, is the cardinal principle of pragmatics: The right way to take what is said, if at all possible, is the way that makes sense of the message. Since the supervaluationist rule would have made hash of our statement of the problem, straightway the rule was suspended. We are good at making these accommodations; we don't even notice when we do it. Under the supervaluationist rule, it's right to say that there's only one cat, and so the candidates have unequal claim. Suspending the rule, it's right to say that the candidates have equal claim, and that all of them alike are not definitely not cats. Suspending the rule, it's even right to say that they are all cats! Is this capitulation to the paradox? – No; it's no harm to admit that in *some* sense there are many cats. What's intolerable is to be without any good and natural sense in which there is only one cat.

Objection.[8] The supervaluationist's notion of indeterminate reference is conceptually derivative from the prior notion of reference *simpliciter.* But if the problem of the many is everywhere, and semantic indecision is inescapable, then reference *simpliciter* never happens. To the extent that we gain concepts by 'fixing the reference' on actual examples, we are in no position to have the concept of reference. Then neither are we in a position to have the derivative concept of indeterminate reference due to semantic indecision.

Reply. We don't need actual examples to have the concept. We have plenty of imaginary examples of reference *simpliciter,* uncomplicated by semantic indecision. These examples are set in sharper worlds than ours: worlds where clouds have no outlying droplets, where cats shed their hairs instantaneously, and so on. When we picked up the concept of reference, in childhood, we probably took for granted that our own world was sharp in just that way. (When not puzzling over the problem of the many, maybe we half-believe it still.) We fixed the reference of 'reference' on these imaginary examples in the sharp world we thought we lived in – and if any theory of reference says that cannot be done, so much the worse for it.

8 Here I'm indebted to Andrew Strauss (personal communication, 1989).

I conclude that the supervaluationist solution to the paradox of 1001 cats, and to the problem of the many generally, is successful. But is it the only successful solution? – I think not. I turn now to the other sort of solution: the kind which concedes that the many are cats, but seeks to deny that the cats are really many.

RELATIVE IDENTITY: THE MANY ARE NOT DIFFERENT CATS

Geach himself favours one such solution. The paradox of 1001 cats serves as a showcase for his doctrine of relative identity.

> Everything falls into place if we realize that the number of cats on the mat is the number of *different* cats on the mat; and c_{13}, c_{279}, and c are not three different cats, they are one and the same cat. Though none of these 1001 lumps of feline tissue is the same lump of feline tissue as another, each is the same cat as any other: each of them, then, is a cat, but there is only one cat on the mat, and our original story stands. . . . The price to pay is that we must regard '——is the same cat as———' as expressing only a certain equivalence relation, not an absolute identity restricted to cats; but this price, I have elsewhere argued, must be paid anyhow, for there is no such absolute identity as logicians have assumed. (1980, p. 216)

'Same cat' is a relation of partial indiscernibility, restricted to respects of comparison somehow associated with the term 'cat', and discernibility by just a few hairs doesn't count. 'Same lump of feline tissue' is a different relation of partial indiscernibility, and a more discerning one.

I agree that sometimes we say 'same', and mean by it not 'absolute identity' but just some relation of partial indiscernibility. I also agree that sometimes we count by relations of partial indiscernibility. As I once wrote:

> If an infirm man wishes to know how many roads he must cross to reach his destination, I will count by identity-along-his-path rather than by identity. By crossing the Chester A. Arthur Parkway and Route 137 at the brief stretch where they have merged, he can cross both by crossing only one road. (1976, p. 27)

I'll happily add that for that brief stretch, the two roads are the same. But though I don't object to this positive part of Geach's view, it doesn't ring true to apply it as he does to the case of the cats.

If you ask me to say whether c_{13}, c_{279}, and c are the same or different, I may indeed be of two minds about how to answer. I might say they're different – after all, I know how they differ! Or I might say they're the same, because the difference is negligible, so I duly ignore it. (Not easy to do while attending to the example as I now am; if I attend to my ignoring of something, *ipso facto* I no longer ignore it.) But if you add the noun phrase, either 'same cat' or 'same lump of feline tissue', it seems to me that I am no less hesitant than before. Just as I was of two minds about 'same', so I am still of two minds about 'same cat' and 'same lump of feline tissue'.

Other cases are different. If you ask me 'same or different?' when you hold Monday's *Melbourne Age* in one hand and Tuesday's *Age* in the other, or when you hold one Monday *Age* in each hand, again I won't know how to answer. But if you ask me 'same or different newspaper?' or 'same or different issue?' or 'same or different copy?' then I'll know just what to say. We can dispute his explanation of what happens, but at least the phenomenon happens exactly as Geach says it does. Not so, I think, for the case of 'same cat' versus 'same lump'.

Something else is lacking in Geach's solution. In other cases where it comes natural to count by a relation other than identity, it seems that identity itself – 'absolute identity' – is not far away. Local identity, as between the Arthur Parkway and Route 137 for the stretch where they have merged, is identity *simpliciter* of spatial parts. Likewise temporary identity, as between a thing and the matter that temporarily constitutes it, is identity *simpliciter* of temporal parts. Qualitative identity is identity *simpliciter* of qualitative character. The newspaper that Monday's *Age* is an issue of and the newspaper that Tuesday's *Age* is an issue of are identical *simpliciter;* likewise my copy and your copy of Monday's *Age* are copies of the identical issue. But Geach never tells us what the 'same cat' relation has to do with identity *simpliciter.*

He wouldn't, of course, because he thinks 'there is no such absolute identity as logicians have assumed'. (Nor would he accept all my examples above; certainly not the one about temporary identity and identity of temporal parts.) But Geach's case against absolute identity

is unconvincing. It seems to come down to a challenge: If Geach is determined to construe all that I say in terms of relations of partial indiscernibility, is there any way I can stop him? Can I *force* him to understand? (What's more, can I do it with one hand tied behind my back? Can I do it, for instance, without ever using the second-order quantification that Geach (1967) also challenges?) I suppose not. But I don't see why that should make me doubt that I know the difference between identity and indiscernibility.

We have the concept of identity, *pace* Geach; and if we are to justify denying that the cats are many, we need to show that they are interrelated by a relation closely akin to identity itself. Geach has not shown this, and wouldn't wish to show it. Nevertheless it can be shown, as we shall soon see. But at that point we shall have a solution that bypasses Geach's doctrine of relative identity altogether.

PARTIAL IDENTITY: THE MANY ARE ALMOST ONE

What is the opposite of identity? *Non*-identity, we'd offhand say. Anything is identical to itself; otherwise we have two 'different' things, two 'distinct' things; that is, two non-identical things. Of course it's true that things are either identical or non-identical, and never both. But the real opposite of identity is distinctness: not distinctness in the sense of non-identity, but rather distinctness in the sense of non-overlap (what is called 'disjointness' in the jargon of those who reserve 'distinct' to mean 'non-identical'). We have a spectrum of cases. At one end we find the complete identity of a thing with itself: it and itself are entirely identical, not at all distinct. At the opposite end we find the case of two things that are entirely distinct: They have no part in common. In between we find all the cases of partial overlap: things with parts in common and other parts not in common. (Sometimes one of the overlappers is part of the other, sometimes not.) The things are not entirely identical, not entirely distinct, but some of each. They are partially identical, partially distinct. There may be more overlap or less. Some cases are close to the distinctness end of the spectrum: Siamese twins who share only a finger are almost completely distinct, but not quite. Other cases are close to the identity end. For instance, any two of our cat candidates overlap almost completely. They differ

by only a few hairs. They are not quite completely identical, but they are almost completely identical and very far from completely distinct.

It's strange how philosophers have fixed their attention on one end of the spectrum and forgotten how we ordinarily think of identity and distinctness. You'd think the philosophers of common sense and ordinary language would have set us right long ago, but in fact it was Armstrong (1978, Vol. 2, pp. 37–8) who did the job. Overshadowed though it is by Armstrong's still more noteworthy accomplishments, this service still deserves our attention and gratitude.

Assume our cat-candidates are genuine cats. (Set aside, for now, the supervaluationist solution.) Then, strictly speaking, the cats are many. No two of them are completely identical. But any two of them are almost completely identical; their differences are negligible, as I said before. We have many cats, each one almost identical to all the rest.

Remember how we translate statements of number into the language of identity and quantification. 'There is one cat on the mat' becomes 'For some x, x is a cat on the mat, and every cat on the mat is identical to x'. That's false, if we take 'identical' to express the complete and strict identity that lies at the end of the spectrum. But the very extensive overlap of the cats does approximate to complete identity. So what's true is that for some x, x is a cat on the mat, and every cat on the mat is almost identical to x. In this way, the statement that there is one cat on the mat is almost true. The cats are many, but almost one. By a blameless approximation, we may say simply that there is one cat on the mat. Is that true? – Sometimes we'll insist on stricter standards, sometimes we'll be ambivalent, but for most contexts it's true enough. Thus the idea of partial and approximate identity affords another solution to the paradox of 1001 cats.

The added noun phrase has nothing to do with it. Because of their extensive overlap, the many are almost the same cat; they are almost the same lump of feline tissue; and so on for any other noun phrase that applies to them all. Further, the relation of almost-identity, closely akin to the complete identity that we call identity *simpliciter*, is not a relation of partial indiscernibility. Of course we can expect almost-identical things to be very similar in a great many ways: size, shape, location, weight, purring, behaviour, not to mention relational properties like location and ownership. But it is hard to think of any very

salient respect in which almost-identical things are guaranteed to be entirely indiscernible. Finally, the relation of almost-identity, in other words extensive overlap, is not in general an equivalence relation. Many steps of almost-identity can take us from one thing to another thing that is entirely distinct from the first. We may hope that almost-identity, when restricted to the many cats as they actually are, will be an equivalence relation; but even that is not entirely guaranteed. It depends on the extent to which the cats differ, and on the threshold for almost-identity (and both of these are matters that we will, very sensibly, leave undecided). What this solution has in common with Geach's is just that we count the cats by a relation other than strict, 'absolute' identity. Beyond that, the theories differ greatly.[9]

ONE SOLUTION TOO MANY?

We find ourselves with two solutions, and that is one more than we needed. Shall we now choose between the way of supervaluation and the way of partial identity? I think not. We might better combine them. We shall see how each can assist the other.

Here is how to combine them. In the first place, there are two kinds of intended interpretations of our language. Given many almost-identical cat-candidates, some will put every (good enough) candidate into the extension of 'cat'; others will put exactly one. Context will favour one sort of interpretation or the other, though not every context will settle the matter. Sometimes, especially in our offhand and unphilosophical moments, context will favour the second, one-cat sort of interpretation; and then the supervaluation rule, with nothing to defeat it, will entitle us to say that there is only one cat. But sometimes,

9 There is another way we sometimes count by a relation other than strict identity. You draw two diagonals in a square; you ask me how many triangles; I say there are four; you deride me for ignoring the four large triangles and counting only the small ones. But the joke is on you. For I was within my rights as a speaker of ordinary language, and you couldn't see it because you insisted on counting by strict identity. I meant that, for some w, x, y, z, (1) w, x, y, and z are triangles; (2) w and x are distinct, and . . . and so are y and z (six clauses); and (3) for any triangle t, either t and w are not distinct, or . . . or t and z are not distinct (four clauses). And by 'distinct' I meant non-overlap rather than non-identity, so what I said was true.

for instance when we have been explicitly attending to the many candidates and noting that they are equally cat-like, context will favour the first, many-cat sort of interpretation. (If we start with one-cat interpretations, and we say things that the supervaluation rule would make hash of, not only is the rule suspended but also the many-cat interpretations come into play.) But even then, we still want some good sense in which there is just one cat (though we may want a way to say the opposite as well). That is what almost-identity offers.

This is one way that almost-identity helps a combined solution. It is still there even when we discuss the paradox of 1001 cats, and we explicitly choose to say that the many are all cats, and we thereby make the supervaluation solution go away.

Perhaps it helps in another way too. The supervaluation rule is more natural in some applications than in others. For instance it seems artificial to apply it to a case of unrelated homonyms. 'You said you were going to the bank. Is that true? No worries, you bank at the ANZ, it's right down by the river, so what you said was true either way!' – I don't think such a response is utterly forbidden, but it's peculiar in a way that other applications of the supervaluation rule are not. The two interpretations of 'bank' are so different that presumably you did make up your mind which one you meant. So the means for coping with semantic indecision are out of place. The supervaluation rule comes natural only when the alternative interpretations don't differ too much. If they are one-cat interpretations that differ only by picking almost-identical cats, that's one way for them not to differ much.

How, on the other hand, do supervaluations help the combined solution? Why not let almost-identity do the whole job?

For one thing, not every case of the problem of the many is like the paradox of 1001 cats. The almost-identity solution won't always work well.[10] We've touched on one atypical case already: if not a problem of the many, at least a problem of two. Fred's house taken as including the garage, and taken as not including the garage, have equal claim to be his house. The claim had better be good enough, else he has no house. So Fred has two houses. No! We've already seen how

10 Here I'm indebted to Phillip Bricker (personal communication, 1990).

to solve this problem by the method of supervaluations. (If that seemed good to you, it shows that the difference between the interpretations was not yet enough to make the supervaluation rule artificial.) But although the two house-candidates overlap very substantially, having all but the garage in common, they do not overlap nearly as extensively as the cats do. Though they are closer to the identity end of the spectrum than the distinctness end, we cannot really say they're almost identical. So likewise we cannot say that the two houses are almost one.

For another thing, take a statement different from the statements of identity and number that have concerned us so far. Introduce a definite description: 'The cat on the mat includes hair h_{17}'. The obvious response to this statement, I suppose, is that it is gappy. It has no definite truth-value, or no definite super-truth-value, as the case may be. But how can we get that answer if we decide that all the cat-candidates are cats, forsake supervaluations, and ask almost-identity to do the whole job? We might subject the definite description to Russellian translation:

(R1) There is something that is identical to all and only cats on the mat, and that includes h_{17}.

Or equivalently:

(R2) Something is identical to all and only cats on the mat, and every cat on the mat includes h_{17}.

Both these translations come out false, because nothing is strictly identical to all and only cats on the mat. That's not the answer we wanted. So we might relax 'identical' to 'almost identical'. When we do, the translations are no longer equivalent: (R1)-relaxed is true, (R2)-relaxed is false. Maybe we're in a state of semantic indecision between (R1)-relaxed and (R2)-relaxed; if so, we could apply the supervaluation rule to get the desired gappiness. Or we might apply the supervaluation rule more directly. Different one-cat interpretations pick out different things as the cat, some that include h_{17} and some that don't. Under any particular one-cat interpretation the Russellian translations are again equivalent, and different one-cat interpretations give them different truth values; so the translations, and likewise the original sen-

tence, suffer super-truth-value gaps. Or more simply, different one-cat interpretations differ in the referent of 'the cat'; some of these referents satisfy 'includes h_{17}' and some don't, so again we get a super-truth-value gap. Whichever way we go, supervaluations give us the gappiness we want. It's hard to see how else to get it.

REFERENCES

Armstrong, D. M. 1978. *Universals and Scientific Realism*, 2 vols. (Cambridge University Press).

Evans, Gareth. 1978. 'Can There be Vague Objects?', *Analysis* 38: 208. Reprinted in *Collected Papers* (Oxford University Press, 1985).

Geach, P. T. 1967. 'Identity', *Review of Metaphysics* 21: 3–12. Reprinted in *Logic Matters* (Oxford: Blackwell, 1972), pp. 238–47.

1980. *Reference and Generality*, 3rd ed. (Ithaca, NY: Cornell University Press).

Lewis, David. 1976. 'Survival and Identity'. In *The Identities of Persons*, ed. Amélie Rorty (Berkeley: University of California Press), pp. 17–40. Reprinted in Lewis, *Philosophical Papers*, vol. 1 (Oxford University Press, 1983), pp. 55–72.

1984. 'Putnam's Paradox', *Australasian Journal of Philosophy* 62: 221–36 (reprinted in this volume as Chapter 2).

1986a. *On the Plurality of Worlds* (Oxford: Blackwell).

1988. 'Vague Identity: Evans Misunderstood', *Analysis* 48: 128–30.

Lowe, E. J. 1982. 'The Paradox of the 1,001 Cats', *Analysis* 42: 27–30.

Unger, Peter. 1980. 'The Problem of the Many', *Midwest Studies in Philosophy* 5: 411–67.

van Fraassen, Bas C. 1966. 'Singular Terms, Truth-Value Gaps, and Free Logic', *Journal of Philosophy* 63: 481–95.

van Inwagen, Peter. 1990. *Material Beings* (Ithaca, NY: Cornell University Press).