Lewis vs Lewis on the Problem of the Many

Dan López de Sa

May 16, 2006

Abstract Consider a cat on a mat. On the one hand, there seem to be just one cat, but on the other there seem to be many things with as good a claim to being a cat, and there seems to be nothing in the vicinity with a better claim. Hence, the problem of the many. In his 'Many, but Almost One,' David Lewis offered two solutions. According to the first, only one of the many is indeed a cat, although it is indeterminate exactly which one. According to the second, the many are all cats, but they are almost identical to each other, and hence they are almost one. For Lewis, the two solutions do not compete with each other but are mutually complementary, as each can assist the other. This paper has two aims: first to argue against the first of these two solutions, and then to defend the second as a self-standing solution from Lewis's considerations to the contrary. In both parts I will assume the certainly plausible but also controversial view on the nature of vagueness, having it that vagueness is a kind of semantic indecision—of which Lewis himself is one of the main defenders.

Keywords problem of the many, vagueness as semantic indecision, almost indentity, counting , David Lewis

Consider a cat on a mat, Tibbles. Most of the hairs, attached to her body, are clearly part of her, and a lot of others, spread out over the mat, are clearly not. But for a number of them, neither of these is the case: they are borderline cases of being hairs of her. Take 1000 of such questionable hairs, and consider the 1000 entities that have just one of these hairs as part of them (and also have all clear parts of Tibbles). Each one of them has equal claim to being a cat, and as good a claim as anything does, or so it seems. So they are all cats, which, together with the one that has all of the borderline hairs as a part, make up 1001 cats, when we were inclined to say there was just one!

This is a version of the paradox of 1001 cats, as an instance of the problem of the many. Roughly, where there seemed to be just one thing of a kind, there turns out to be a number of candidates with as good a claim as anything to be something of that kind. A solution by disqualification aims to disqualify most or all of the candidates, thus rejecting and explaining (away) the intuition according to which all of them have an equal claim by having the features relevant for being of the kind. An egalitarian solution aims to vindicate the good claim of the plurality, thus rejecting and explaining (away) the intuition according to which there is just one of

the kind.

In his 'Many, but Almost One,' David Lewis (1993) offered two solutions to the problem of the many, one of each of the types mentioned. According to the first, only one of the many is indeed a cat, although it is indeterminate exactly which one. According to the second, the many are all cats, but they are almost identical to each other, and hence they are almost one. Lewis thought that these two solutions do not compete with each other but are mutually complementary, as each can assist the other.

This paper has two aims: first to argue against the first of these two solutions, and then to defend the second as a self-standing solution from Lewis's considerations to the contrary. In both parts I will assume the certainly plausible but also controversial view on the nature of vagueness, having it that vagueness is a kind of semantic indecision—of which Lewis himself is one of the main defenders.

I The Problem of the Many

The problem of the many is ubiquitous. Take what seems to be a single cloud in an otherwise clear blue sky. Clouds are made of water droplets, standing in appropriate relations to each other. There are relatively isolated droplets that are clearly not part of the cloud, but there are also plenty of droplets in the border. Thus there will be many aggregates of droplets, some more inclusive, some less inclusive, some inclusive in different ways than others, all of which seem to have as a good claim as anything in the vicinity to be cloud. Each of them is made of water droplets in the way clouds are, look, move and interact with other meteorological phenomena in the characteristic way clouds do, including eventually raining. So it seems, all of them are clouds—but, there was only a single one to begin with!

Mutatis mutandis for mountains, chairs, coins, persons, and so on. In each case, we have a paradox, in that we seem to have strong intuitions pulling in different, inconsistent directions. On the one hand in our examples it is clear that we are considering cases where there seems to be just one cat, or just one cloud, and so on. But on the other hand, once we attend to the plurality of candidate entities, we realize that each of them has the features required for being a cat, a cloud, and so on. So each of them has an equal claim, and indeed as good a claim as anything in the vicinity, to be a cat, a cloud, and so on. I will speak of the first sort as *counting* intuitions, as they concern the number of things of the kind we seem to be dealing with, and of the second sort as *grounding* intuitions, for they concern the thought that each of the many does exhibit the appropriate grounds to be of the kind.

Most people dealing with the problem of the many ultimately reject the grounding intuitions, thus defending a solution by disqualification that tries to explain them away.[1] Contrary to appearances, they hold, it is not the case that there are many

[1] One exception is precisely Unger, in the seminal 'The Problem of the Many' (1980): by holding that it turns out there are no cats, nor clouds, nor people, etc., he qualifies as proposing a (nihilist but) egalitarian solution, rejecting the counting intuition—which has it that there *is* (just) one of the kind, in the relevant scenarios.

things each of which has an equal claim to be of the kind, and as good a claim as anything in the vicinity. This, in turn, in virtue either of there not being the many things, or of not all of them having an equal claim, or of there being something else that has a better claim, after all.

II Some Solutions by Disqualification

A solution by disqualification thus consists in denying that the many are all really cats, clouds, and so on.

One extreme form such a solution might take is just to refuse to acknowledge that there *are* many cat-candidates, to begin with—regardless of the issue as to how good their claim is, or would be. Some people suggest that this need not be such an extreme position, as they seem to understand that the problem of the many requires the principle of unrestricted mereological composition—which, though plausible, some people reject.[2] Not so. As the examples make clear, all *moderate* restrictions on composition would plausibly generate the required plurality of things, as the cases are minutely different to each other, and to a paradigm case of composition, if such there is at all. Hence, the only option here seems to be to appeal to a *brutal* restriction of composition, which most foes of unrestricted composition would also reject—or invoke vagueness *in rebus*, as considered below.

Another more popular way of rejecting that each of the many candidates is a cat is actually to claim that *none* of them is a cat. There are all these many candidates, and each has an equal claim to be a cat—but there is something else in the vicinity with a much better claim to be a cat: the one and only cat. The many candidates are merely "lumps of feline tissue" which constitute, without being identical to, cats. What the problem of the many shows is that this relation of constitution is vague: the many are not *cat*-candidates, but cat-*constituter*-candidates, each of which can be said to (indeterminately) constitute the one and only cat on the mat. (See (Johnston, 1992), (Lowe, 1995).) Arguably, any such view entails that there is vagueness *in rebus*: entities like the one and only cat on the mat which are (determinately) non-identical to any entity specifiable in a precise way[3]—such as, we may assume, the many candidates. Many (even if not all) philosophers dealing with this

[2] A *sum* (or *fusion*) of some things is something such that anything overlaps the sum if and only if it overlaps some of those things—one thing overlaps another if there is something that is part of both. There is clearly something that is the sum of my head, torso and limbs: my body. But, to use Lewis's example, is there something that is the sum of the right half of my left shoe plus the Moon plus the sum of all Her Majesty's ear-rings? According to the *principle of unrestricted mereological composition*, whenever there are some things, no matter how disparate and unrelated, there is something that is a sum of them or, as it is also said, composition takes place. See (Lewis, 1986) for further discussion, and for an argument against moderate restrictions on composition—i.e., restrictions that aim to offer an explanation in more basic terms of why in some but not all cases composition occurs.

[3] This is only a rough negative characterization, for further discussion see (Barnes, MS). That initial working characterizations are all negative is crucial, given that, as we are about to see, it might be at stake precisely whether it is ultimately possible to have a correct conception of what a vague entity (that does not represent things being a certain way) would be.

issue agree in rejecting a metaphysical view on vagueness of this kind. (Lewis, 1993) contains what are—for some, myself included—the main reasons against vagueness *in rebus*: it is hard to have a correct conception of what a vague entity would be, and the phenomena allegedly motivating the view are neatly accommodated by the alternative views. Regarding the problem of the many, it is not even clear that the view would actually provide a solution. To begin with, it is not clear why being, at least possibly, vague *in rebus* should be a feature of things like cats at all: hence, adding a vague object to the 1001 candidates only makes it a paradox of the 1002 cats. Furthermore, and perhaps even more importantly, it is not clear why there would not be 1001 *vague* cat-candidates. See (Hudson, 2001), (Morreau, 2002) and (Eklund, 2005) for further recent discussion. Although I tend to regard these difficulties as definitive, this view of mine will play no role here.

III Vagueness as Semantic Indecision

I will assume here the view of vagueness as *semantic indecision*. It holds (roughly) that, whatever it is that in the thoughts, experiences and practices of language users determines the meaning of expressions, it fails to determine any one in particular of a given range of candidates for vague expressions. Every way of ("arbitrarily") fixing what is left semantically indeterminate gives rise to a "precisification" or "sharpening" of the original vague expression.(Strictly speaking, sharpenings are of the language as a whole, and not of isolated expressions.)

Although all such sharpenings are, by their essence, arbitrary to a certain extent, not all of them are *admissible*. At least some features of the indeterminate meaning should be respected by them. In particular, for the case of predicates, they should preserve *clear cases*, both of application and of non-application. To illustrate: no sharpening according to which 'Ronaldinho is bald' is counted as true would be admissible for 'is bald,' nor would it be a sharpening according to which 'Ronaldo is bald' is not counted as true. Furthermore, *penumbral connections* should also be preserved. To illustrate: no admissible sharpening for 'is rich' according to which 'Ferran is rich' counts as true could fail to count 'Patri is rich' as true, if Patri is richer than Ferran. And no admissible sharpening for 'is red' and 'is pink' counts 'Rose Fifí is red and pink' as true, despite Fifí's being a borderline case of each.[4]

Finally, what one says by means of a vague expression is true if it counts as true according to all admissible sharpenings; is false if it counts as false according to all of them; and is indeterminate otherwise. Statements involving vagueness thus become supervaluated.[5]

[4]See (Fine, 1975). How to characterize in an explicit satisfactory way the notion of *admissible* that these connections (possibly among others) constitute, though central to a full defense of the view of vagueness as semantic indecision, is not crucial for our present concerns. Notice that 'is admissible' is, of course, itself vague: this is arguably part of what accounts, in this framework, for the phenomenon of "higher-order" vagueness. Complications involving this will be set aside here.

[5]Thus it is false according to the view that, in general, an utterance that says something is true or false. (Williamson, 1994) contains an argument for the incompatibility of this feature with Tarskian views about truth and falsity, which apparently convinced most people in the field. In

IV The So-Called "Supervaluationist" Solution

The view of vagueness as semantic indecision provides a solution to the main paradox involving vagueness. Consider a sorites series going from our paradigm cat Tibbles to pet-robot Tama, such that each individual has one more natural organ replaced by an artificial prosthesis.[6] As both Tibbles and Tama are paradigm cases with respect to 'is a cat,' 'Tibbles is a cat' turns out to be true, as all sharpenings count it as such, whereas 'Tama is a cat' turns out to be false, as all sharpenings count it as such. With respect to the sorites premise itself, 'If one of the individuals is a cat, so is one which differs from the first by having just one organ replaced by an artificial prosthesis,' it turns out to be false, as all sharpenings count it as such. Each fixes a sharp boundary for cathood "arbitrarily" at a particular juncture, albeit a different one for each. Hence, there is a cut-off point, although exactly where it is indeterminate.

The first solution offered by Lewis to the problem of the many aims to disqualify all but one of the candidates in a somewhat similar fashion. As we have just seen, 'is a cat' is vague and hence, assuming vagueness as semantic indecision, it has different admissible sharpenings. Suppose that each of them is such that one and only one of the different candidates counts as a cat, according to it—a different one, according to different sharpenings. Therefore, one and only one of the many is indeed a cat, although it is indeterminate exactly which one. This is how the solution rejects, and aims to explain away, the grounding intuition: the arbitrariness felt in having one of the candidates with a better claim than the rest comes from the fact that it is indeterminate which of them is the one with this better claim.

This solution is sometimes referred to as "supervaluationist," given the use of supervaluations over admissible sharpenings. The labelling might be misleading, as it is sometimes also used to refer to the view about the nature of vagueness as semantic indecision which, with Lewis, I am also assuming here. Although proponents of the view of vagueness as semantic indecision have also favored the so-called "supervaluationist" solution to the problem of the many, they need not do so. The point of this paper is precisely that they *should* not.

---

my view, however, (Andjelković & Williamson, 2000) contains the key elements for resisting it, see for elaboration (López de Sa, MS).

[6]For further info about Tama, see http://news.bbc.co.uk/1/hi/sci/tech/652293.stm. As with the usual examples involving 'is bald,' 'is rich,' etc., this involves idealization regarding which feature is such that the predicates seem tolerant with respect to small changes in it, but not with respect to big ones: number of hairs, amount of money, and here replacement of natural organs by artificial prosthesis. If whole organs seem too big for this—when attending to particularly important ones such as brains, hearts and so on—replace the series by a longer one in which merely tiny bits thereof are replaced. I use, for illustrative purposes, the sorites for 'is a cat' because it will prove relevant to the discussion to come, but also to make the following point: the problem of the many and the sorites paradox arise also for predicates like 'is a cat' which arguably signify natural (enough) properties or kinds: even if some candidates are intrinsically more eligible than others, there might well be many equally eligible natural (enough) candidates for predicates like 'is a cat.' For this point, see also (Lewis, 1993, footnote 6)

v Penumbral Connections

My main worry about the so-called "supervaluationist" solution to the problem of the many can be stated quite simply. Any sharpening of 'is a cat' that distinguishes one of the cat-candidates from the rest would thereby violate penumbral connections, and this renders it inadmissible.

One way of stating the relevant violated principle is Unger's "Principle of Minute Differences:"

(PMD) If something is a *paradigm* case of an f, and something else is very similar to the former with respect to the features relevant for something being an f, then the latter is also an f;

where 'f' is substituted by a common noun of the sort of those that concern us here—'cat,' 'cloud,' 'coin,' 'mountain,' 'person,' and so on.[7]

The cat on the mat is a paradigmatic cat, if anything is: Tibbles figures at one of the clear ends of the considered sorites series, could be used as the relevant sample in an ostensive definition of 'is a cat,' and what have you. Now according to the defender of the solution we are considering, in each sharpening just one of the cat-candidates is indeed the one and only cat—and thus a paradigmatic cat. But clearly all the others are very similar to a paradigmatic cat, with respect to the features relevant for something being a cat—after all, they all count as the one and only (paradigmatic) cat according to some other of the sharpenings. Paraphrasing Lewis himself (1993, 168), the cat-candidates are all 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-candidates also can be; anything a cat can do at a moment, cat-candidates also can do. They are too cat-like not to be cats! Still, they are not counted as cats by the envisaged sharpening. The sharpening thus violates (PMD) for cats, and is thereby rendered inadmissible.

Another example, adapted from a related discussion (McKinnon, 2002), may help to reinforce the point. Suppose that next to Tibbles, there is also—what we are inclined to count as—(just) one small golden coin, an old peseta. According to the solution, just one of the 1001 coin-candidates in the vicinity of the coin is indeed a coin, although each one has an equal claim to be one, and thus it is indeterminate exactly which one is. But the coin may be a gift to visitors at the exhibition *The 1001 Pesetas*, consisting of 1001 (distinct) coins, each of which happens to be a duplicate of one of the coin-candidates on the mat (the increase in determination might be the product of the special conditions of the Museum's exhibition hall.) Duplication should certainly count as sufficient for close similarity in the features that are relevant for coin-ness. (We need not take duplication to be restricted to intrinsic features, but can also suppose relevant parts of histories and the like to

[7] I have altered the formulation. The original one runs: "With respect to any *kind of ordinary things*, if something is a *typical member* of the kind, then, if there are entities that differ from that thing, in any respects relevant to being a member of the kind, quite *minutely*, then each of those entities is a member of that kind." (1980, 447).

be duplicated.) The 1001 coins in the Collection are quite paradigmatic, and hence admissible sharpenings of 'is a coin' preserve them all. Therefore, sharpenings that count just one of the coin-candidates on the mat would violate (PMD) for coins, inadmissibly.

Thus the so-called "supervaluationist" solution to the problem of the many requires that each sharpening selects just one of the candidates, but in so doing the sharpenings are rendered inadmissible, by violating penumbral connections such as (PMD).

VI Paradigmatic Cats

In the discussion I have just alluded to, Neil McKinnon (2002) offers a related criticism of the so-called "supervaluationist" solution to the problem of the many. Unfortunately, he claims that sharpenings are rendered inadmissible by violating the maxim of "Non-Arbitrary Differences," here stated for coins:

(NAD) For any coin and non-coin, there is a principled difference between them which forms the basis for one's being a coin and the other's being a non-coin;

which, he says, "imposes the following penumbral connection on every permissible sharpening: if d is a coin, then so is e unless it differs from d in a principled way" (2002, 333)—a principled difference is a relevant difference in the features that are relevant for something being a coin, so that the coin-candidates do not differ in a principled way.

(NAD) crucially differs from (PMD) in not being restricted to paradigmatic cases. And this is unfortunate because, without the restriction, there is every reason to reject the claim that (NAD) is a penumbral truth. Rather, it is an intuitively appealing but ultimately rejectable soritical principle, which is inconsistent with there being (paradigmatic) coins and (paradigmatic) non-coins, which can be connected in a sorites series made up of individuals that do not differ from their adjacent ones in a principled way.

By contrast, (PMD) is not soritical, as it is restricted to a paradigmatic case of the kind—and does not require that those things minutely differing from it are in turn paradigmatic, but merely that they are of the kind.[8] I will thereby stick to (PMD).

VII Determinate Cats

It has been suggested to me that the assumption I made, that the cat on the mat is a paradigmatic cat, might be problematic—if a paradigmatic cat is to be understood in the sense of a *determinate* cat. For the so-called "supervaluationist" solution to

[8]The restriction will also be important in the discussion of the Problem of the Two, in the second part of the paper.

the problem of the many would entail precisely that there are no determinate cats.[9]

Rightly so. But this just points to the other major difficulty that the so-called "supervaluationist" solution faces. As we have seen, the problem of the many arises also with respect to paradigmatic clear cases, such as Tibbles, as there is in the vicinity a plurality of equally good candidates. The solution requires that each sharpening counts just one of them as a cat. Hence there is nothing, however clear or paradigmatic, that is preserved in the extension of 'is a cat' across the different sharpenings. And yet that they preserved clear cases of both application and non-application was a constraint on which sharpenings of predicates like 'is a cat' were admissible!

Indeed there seems to be something deeply disturbing about the thought that there is no relevant difference between the individuals at the clear end of the sorites series and those in the middle ground between Tibbles and pet-robot Tama: all of them are, according to the solution, merely borderline cases with respect to 'is a cat.' They are similar in being counted as cats by some but not all of the sharpenings of the predicate. Emphatically:

> There isn't, anywhere in the world, anything of which it is determined that it satisfies 'mountain.' Forget about thinning hair. Nothing is determined to satisfy 'bald man,' because nothing is determined to satisfy 'man.' (McGee, 1998, 145)

Thus the so-called "supervaluationist" solution to the problem of the many requires that each of the candidates is rejected by some sharpening, but in so doing the sharpenings are again rendered inadmissible, if paradigmatic Tibbles should indeed be determinately a cat.[10]

viii Maximal Cats

It has also been suggested to me that one could, on behalf of the defender of the so-called "supervaluationist" solution to the problem of the many, argue against (PMD) being a penumbral truth by contending that 'is a cat' signifies a *maximal property*, in Ted Sider's (2001, 2003) sense. I think this would put the cart before the horse—in a way that Sider himself does not seem to do. Let me explain.

A property F is *maximal*, in this sense, iff (roughly) large parts of an F are

[9] At least, on standard ways of characterizing what it is for something to satisfy a determinately-involving matrix, see (McGee, 1998).

[10] Robbie Williams offers a criticism of the solution of this sort in (2006). Besides, he also claims that, in virtue of nothing determinately satisfying 'is a mountain,' the solution undermines the explanation offered by defenders of the view of vagueness as semantic indecision such as (Keefe, 2000) of the persuasiveness that the (false) sorites premise certainly has: "Our belief that there is no true instance of the quantification gets confused with a belief that the quantified statement is not true. ... The confusion ... is a confusion of scope, according to whether the truth predicate appears inside or outside the existential quantifier" (2000, 185). Insofar as I can see, the difference in scope in truth- (or determinate-) involving existential statements appealed to here is, however, compatible with nothing determinately satisfying 'is a mountain'—however disturbing the latter might be for other reasons.

not themselves Fs. According to Sider, the properties of being a cat (and a house, and a rock, and a person, etc.) are all maximal properties. This is supposed to be ultimately grounded in linguistic intuition dictating things like the following: House-minus—the mereological difference between a given house and one of its windows—is not a house, even if it would have been, had the window never existed. Similarly, Tibbles-minus—Tibbles minus her tail—is not a cat, but might have been one.

Now the (alleged) fact that 'is a cat' signifies a maximal property does not provide, by itself, any solution to the problem of the many. This is because, unlike Tibbles-minus, it is not true of the many different candidates that they lack parts that are determinately part of Tibbles. This is clearly acknowledged by Sider:

> Forget about counting; consult your linguistic intuitions about whether House-minus is a house directly. Mine say that it is not. Here there is an asymmetry with the problem of the many, for in that case linguistic intuition refuses to identify any of the many as being any more house-worthy than the rest. Not so for large parts of the house that exclude bits that are definitely part of the house. These seem clearly not to be houses. (2001, 359)

My own intuitions are here much less clear than Sider's regarding House-minus and Tibbles-minus, and I think one might explain appearances to the contrary by appealing to the pragmatics of counting, in a way that we are asked to forget about here. But never mind this now—it will concern us later, when discussing the Problem of the Two. What is important to emphasize is that appealing to maximality does not by itself solve the problem of the many, as many of the many candidates are equally maximal—none being a proper part of another, and none lacking parts that are determinate parts of the cat on the mat.

This notwithstanding, something structurally similar to maximality is actually presupposed by the so-called "supervaluationist" solution to the problem of the many:

> 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. (Lewis, 1993, 170)

As Sider says (2001, 358), the sharpenings required by the so-called "supervaluationist" solution all validate a "maximality maxim"—strictly speaking, an "exclusion principle" along the lines of:

(EP) If several things overlap massively, then at most one of them is a cat, or a mountain, or a house.

Now it would be to put the cart before the horse to simply take for granted that (EP) is intuitively true—let alone a penumbral connection.[11] We *do* have strong intuitions to the effect that there is *just* one cat on the mat. And something like (EP)

[11] This is done, in my view, by the recent defenders of the so-called "supervaluationist" solution to the problem of the many—(McGee & McLaughlin, 2000) just take exclusion principles like (EP) to be "penumbral constraints" (2000, 138), which are "presumably analytic" (2000, 142); for (Varzi, 2001), they are "part of our concept of a mountain" and the like (2001, xx); according to (Weatherson, 2003), the fact that something massively overlapping a coin is not a coin shows

is what one turns out to be committed to, if the counting intuition is to be respected by defending a solution by disqualification such as the so-called "supervaluationist" one. But this does not undermine at all the intuitive character of (PMD). Otherwise, there would be no *problem* of the many, as nothing would go, at least *prima facie*, against the counting solution.

IX Summing Up

It is of course a possibility that, its intuitive plausibility notwithstanding, the "very surprising" upshot of the problem of the many is that (PMD) turns out *not* to be true—in a similar way in which the upshot of the sorites is that the sorites premise turns out not to be true. This is the position of Mark Johnston, see (1992, 99). A *solution* to a paradox requires one to explain the rejected intuitions away, perhaps by identifying something in the vicinity of the surprising falsity that is still true after all—as is typically done with respect to the soritical premise in the Sorites paradox.

Johnston, for instance, attempts to do this by dwelling on views on relations of constitution other than identity and consequent distinctions between ordinary things and mere quantities or pieces of matter.[12] In defense of the "supervaluationist" solution, no such explanation of a similar sort has been, in my view, offered—nor, in fact, even seriously attempted. McGee & McLaughlin marginally mention the following:

> If something is a typical mountain, and if x is an entity that differs from it, in the respects that are relevant for something being a mountain, at most quite minutely, then there is, in x's mediate vicinity, a mountain composed of virtually the same matter as x. (2000, fn. 31)

But no attempt is made to show how intuitions favoring (PMD) could indeed be explained away in terms of this. And the prospects for it seem, initially, quite meagre. Or, at least, much less promising than with respect to the corresponding explanation regarding counting that will concern us for the rest of the paper.

Paradoxically enough, Lewis himself does not seem to disagree radically with such a judgment against the so-called "supervaluationist" solution to the problem of the many. When considering the sense in which the solution is merely partial, and in need of assistance, he says:

> When we have been explicitly attending to the many candidates and noting that they are equally catlike, context will favor [sharpenings that put every (good enough) candidate into the extension of 'cat,' and not sharpenings that put exactly one]. This is one way that almost-identity helps a combined

(NAD) (and presumably (PMD)) to fail (2003, 494–95). This is not done, as we are about to see, by Lewis himself.

[12] The true principle to replace (PMD) according to Johnston is this: "If y is a paradigm F and x is an entity that differs from y in any respect relevant to being an F only very minutely *and x is of the right category, i.e., is not a mere quantity or piece of matter*, then x is an F" (1992, 100), and he then argues that none of the candidates are cats or clouds or ... as they all are mere quantities or pieces of matter.

> solution. It is still there even when we discuss the paradox of the 1001 cats, and we explicitly choose to say that the many are all cats, and we thereby make the supervaluationist solution go away. (Lewis, 1993, 180)

To my mind, this seems close to acknowledging that the so-called "supervaluationist" solution goes away precisely when one is dealing with the problem of the many. One might then wonder in what sense it deserves to be called a solution to the problem at all. In fact, this, plus the sense we will discuss in which, for him, the alternative solution is in need of assistance, plus the very title of his paper itself, might suggest that he did not—even partially—hold a "supervaluationist" solution to the problem of the many. But I do not want to go into Lewisian exegesis here.

x The Almost-Identity Solution

Lewis's second solution does not aim to disqualify any of the candidates: it is an egalitarian solution, aiming to vindicate the good claim of the plurality. The many are indeed cats. It thus rejects and tries to explain (away) the counting intuition, having it that there is just one cat on the mat.

All the cat-candidates are certainly different things: their non-identity is actually obvious from the very beginning, as they differ in their parts. But *different* things need not be *distinct*, in the sense of non-overlapping. There may be something that is a common part of them, and thus be *partially identical*, as Lewis puts it following Armstrong.[13] All the cat-candidates are indeed partially identical to each other. They are all cats, the solution has it, but not distinct cats. Actually, as Lewis says, any two of the cat-candidates overlap almost completely, and thus they are *almost* identical. This provides the almost-identity solution to the problem of the many:

> 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 ... We have many cats, each one almost identical to all the rest. (Lewis, 1993, 178)

The many candidates have as good a claim as something can possibly have to being a cat. And thus they all *are* cats.[14]

[13]For Armstrong, however, things may be partially identical in virtue of sharing a "non-mereological constituent," like two different states of affairs involving the same universal but different (perhaps indeed distinct) particulars. In the present paper, as in Lewis's, two things are (at least) *partially identical* iff there is something that is part of both. Partial identity is therefore simply identity of parts.

[14]Lewis's second, almost-identity solution is thus an "over-population" solution, in (Weatherson, 2004)'s taxonomy, as it rejects the premise that there is at most one cat on the mat. According to Weatherson, however, this is a misattribution, and he quotes Lewis stating that the second solution is of a "kind which concedes that the many are cats, but seeks to deny that the cats are really many" (Lewis, 1993, 175). On this basis, Weatherson concludes that Lewis would rather reject the contention that if something is a cat, and something different is also a cat, then there are two cats. But, as we have just seen, Lewis is quite explicit in claiming that, strictly speaking, there are many cats. As to the quote provided by Weatherson, occurring before the discussion of the second solution has started properly, it can obviously be seen as a way of making, non-strictly, the point about non-strict counting to be considered shortly.

xi Counting Cats

Almost everybody in the debate regards the almost-identity solution to the problem of the many as just "implausibly counterintuitive," and as such it is normally dismissed with (at most) a couple of sentences to the effect that it is obvious that there is just one cat on the mat.

It is beyond doubt that *obviously* there seems to be just one cat on the mat, and the counterintuitive character of the solution is no less manifest. No surprise here: after all, we are dealing with a solution to a paradox, which as such needs to ultimately reject something that had strong intuitions in its favor. Not to acknowledge this would be an error analogous to that of not appreciating the counterintuitive character of the exclusion principles considered above—nor, thereby, appreciating the force of the problem of the many.

Now, whether the undoubted counterintuitive solution is *implausibly* so will depend on how plausible is the explanation (away) of the counting intuition on offer—in particular, how plausible it is in contrast to the plausibility of the alternatives. And this, it seems fair to say, is something which the swift dismissals of the almost-identity solution mentioned do not tend to pause on.

According to Lewis, in most conversational contexts, the intuitive (and appropriate) answer to 'How many cats are there on the mat?' is indeed 'Just one.' Thus, although strictly speaking false, it is loosely speaking true enough:

> 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. (Lewis, 1993, 178)

To elaborate, it is a generally acknowledged fact that in a lot of conversations, what is appropriate is to (non-strictly) count by relations other than (strict) identity. Sometimes one counts by relations of partial indiscernibility—which are equivalence relations on the relevant domains. One such case, mentioned by Lewis, is the following:

> If an infirm man wishes to know how many roads he must cross to reach 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 just one road. (Lewis, 1993, 175)

One other case, structurally more similar to our case in the relevant respects, is also given by Lewis:

> 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); (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. (Lewis, 1993, fn. 9)

Here one seems to be counting by the (non-transitive) relation of overlapping. This hints at a more general characterization of the phenomenon. In a given counting conversation, quantifiers are restricted by excluding the rest of different things that stand in the given salient relation to the one considered. Hence, in the road case, if a road is counted, the rest of partially indiscernible roads are disregarded. In the triangle case, if a triangle is counted, the rest of overlapping triangles are disregarded. In the limiting case in which every difference is relevant, if we count something, the rest of things that are identical to it are disregarded—to wit, nothing is disregarded.

If not the details, something along these lines seems to be motivated by considerations about the pragmatics of domain restriction in counting conversations, which arise independently of issues having to do with the problem of the many.

Now suppose that the many cat-candidates are indeed cats. The prediction is that the independently motivated conversational mechanism about counting would indeed make it the case that, in most conversations, one should count by almost-identity, delivering the answer 'Just one.' to the question of how many cats there are on the mat. For, once one is counted, the rest of cats almost identical to it are disregarded. This is how the counting intuition is explained (away), according to the almost-identity solution to the problem of the many.

Still, an opponent might stress, there is a sense in which, strictly speaking, the cats are indeed many, according to the solution. One would then expect that there are contexts, not of the most ordinary variety perhaps, in which the appropriate answer to the question, 'How many cats are there on the mat?', is indeed 'Many.' Such a context would be one in which the minute differences between the different almost identical cat-candidates are not negligible but, by contrast, relevant—so that the rest of almost identical cats are not ignored but included in the domain of quantification salient in the conversation. Not an ordinary context, indeed. But the context of considering the paradox seems to be precisely a context of this kind.

With these programmatic remarks I do not aim to dispel all the reasonable doubts that might arise in connection with the Lewisian explanation of the ultimately rejected counting intuition, but I *do* hope it will transpire that the view is not completely "insane"—and its prospects, for all we know, may easily be more promising than those to be offered on behalf of the so-called "supervaluationist" solution to the problem of the many. I want to devote the rest of this second part of the paper to discussing the sense in which, according to Lewis, the almost-identity solution is also in need of assistance, to be provided by supervaluations.

xii The Problem of the Two

According to Lewis, then, the almost-identity solution is also in need of assistance, to be provided by supervaluations. He mentions two considerations. Here is the first

one:

> Why not let almost-identity do the whole work? 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. We've touched on one atypical case already: if not a problem of the many, at least a problem of the two. Fred's house taken as including the garage, and taken as not including the garage, have equal claim to be his house. So Fred has two houses. No! ... 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. (Lewis, 1993, 180)

Agreed. But, on the face of it, this by itself does not prove the almost-identity solution to be substantially incomplete—or, rather, at best it shows that the appropriate "many" solution, if it is to take care both of instances of the problem of the many such as the paradox of the 1001 cats and the problem of the two, should take a more general form—having the almost-identity version as a special case, which suits the special case of the problem of the many constituted by the paradox of the 1001 cats and the like.[15] And we have just seen what this form might be, when considering the case of counting the triangles in the square: in the most natural contexts, the domain might be restricted by disregarding the things which overlap in a salient way with those counted. This respecting the equal claim that the different candidates have to be, in the case, houses.

Likewise, it seems to me, for the issue considered above about maximality. The contention that typically expressions of the sort of 'is a rock,' 'is a cat,' 'is a conscious being' and the like signify maximal properties has substantial consequences, including that, contrary to appearances, the properties of being a rock, a cat, or a conscious being turn out to intrinsic, in unexpected ways. Now the main consideration Sider provides in favor of the contention involves in effect the counting intuition:

> Otherwise in the vicinity of every house there would be a multitude of houses; in the vicinity of every cat there would be a multitude of cats. (Sider, 2003, 139)

Sider is of course very familiar with the mechanism of domain restriction that could alternatively explain (away) intuitions about counting—he himself makes substantial use of the strategy in other contexts, including that of defending the principle of unrestricted mereological composition from the charge that in most conversations people would not count "weird" sums as things over and above their more natural constituents. This is why, I take it, he aims to offer a further, independent motivation, which I quoted above, and repeat here:

> Forget about counting; consult your linguistic intuitions about whether House-minus is a house directly. Mine say that it is not. (2001, 359)

[15] I am indebted here to Robbie Williams.

But, as I suggested above, this further consideration seems to be rather weak. For what it is worth, mine say that it is.[16]

Suppose that someone, for whatever reason, is not convinced by the previous line of argument, but rather invokes supervaluations to deal with the problem of the two. In what sense, if any, would this undermine the self-standingness of the second, almost-identity solution to the problem of the many? Two thoughts might come to mind, but neither proves conclusive. First, one might think that the paradox of the 1001 cats and the problem of the two exhibit a relevant common, unifying form, in virtue of which they call for a unified solution. Indeed. But if someone has a reason for preferring to invoke supervaluations where the pragmatic explanation of the counting intuition is, as seen, available, this must be because she does not regard the cases as relevantly similar. And this undermines (for her) the request for a unified treatment of the cases.

Second, and more importantly, if one invokes supervaluations, the different sharpenings would count just one of the two house-candidates as a house, their equal claim to being houses notwithstanding. Why should this be any more acceptable than the sharpenings counting just one of the cat-candidates? Doesn't this then undermine my previous case against the so-called "supervaluationist" solution to the problem of the many? I don't think so. Let me explain.

Sharpenings of this sort are rendered inadmissible by McKinnon's soritical (NAD). Here is Weatherson:

> It might not be clear whether, for example, for example, we have one mountain with a southern and a northern peak, or two mountains, one of them a little north of the other. Whether there is one mountain here or two. Clearly the two peaks exist, and their fusion exists two. The real question is which of these three things is a mountain. However this question is resolved, a substitution instance of (NAD) with the two objects being the southern peak and the fusion of the two peaks will be false. (2003, 497)

The last statement is not, strictly speaking, true. As we have seen, on one way of resolving the question, the three of them are indeed mountains. And this is compatible with fluctuation between 'One' and 'Two' being the intuitively appropriate response to the question of how many mountains there are. Actually, as suggested above, I think this is the best option in this case. Compare: in a classroom, four small tables are aligned two-by-two to make a bigger one. How many tables are there? Fluctuation between 'One' and 'Four' is to be expected, and not 'Five,' let alone 'Nine' or others. Still, all the candidates would clearly be reckoned as tables on reflection, or so it seems to me.

Suppose, however, as I said, that for whatever reason someone is dissatisfied with this response, and invokes supervaluations as to make 'One or two, but not three,' not only conversationally appropriate in most contexts, but the only strictly speaking true answer to the question about the number of mountains. Then, as

[16]Something in the thought about maximality might still be right, even for those who reject the view: it seems plausible that the relevant domain restriction would first disregard things that are proper parts of cats, houses, and so on.

Weatherson observes, we would have a counterexample to (NAD) for mountains—which was false anyway, as we saw, due to its soritical character.

What it is important to notice is that this situation is perfectly compatible with (PMD), again due to the restriction to paradigm cases—on the understanding that houses, mountains and the like which are subject to the problem of the two are, as Lewis had it, atypical, or anyway not paradigmatic.

So either the problem of the two is indeed relevantly similar to the paradox of 1001 cats, or it is not. If the former, the almost-identity solution is only a particular version of a more general "many" solution to the problem of the many. If the latter, the appeal to supervaluations is compatible with the previous case against the "supervaluationist" solution. In either case, the almost-identity solution is not shown to be in need of substantial assistance.

XIII Which of the Cats is 'the Cat on the Mat'?

The other consideration that Lewis mentions concerns the semantics of singular definite descriptions such as 'the cat on the mat.' As I understand him, Lewis points to the fact that, assuming the general view of vagueness as semantic indecision, there should be a—possibly less than fully strict—sense in which 'the cat is on the mat is brown' counts as true, whereas 'the cat on the mat includes hair h_{17}' counts as neither true nor false—provided h_{17} is part of some but not all the many cats.

One mechanism for obtaining this result is offered by Lewis himself on behalf of the defender of the almost-identity solution:

> 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 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)-relax and (R2)-relaxed; if so, we could apply the supervaluation rule to get the desired gappiness. (Lewis, 1993, 181)

So, in general, the proposal is that 'the cat on the mat is F' receives, non-strictly, the value received by both 'there is something that is almost-identical to all and only cats on the mat, and it is F' and 'something is almost-identical to all and only cats on the mat, and every cat on the mat is F' if such there is, and is, non-strictly, indeterminate otherwise.

Of course, as Lewis observes, we would get the gappiness more directly—and indeed as the strict status—if things were the way they are according to the so-called "supervaluationist" solution to the problem of the many.

So far, so good. In what sense does all this tell against the "completeness" of the almost-identity solution? What Lewis says is this:

> Whichever way we go, supervaluations give us the gappiness we want. It's hard to see how else to get it. (Lewis, 1993, 182)

So tt seems to me that, rather than submitting a consideration for the view that the almost-identity solution to the problem of the many *as such* is in need of assistance, he is in effect observing that supervaluations might still be needed to account for *other* issues—like, here, involving the semantics of singular definite descriptions.

xiv Which of the Cats is 'Tibbles'?

And this seems to be certainly right—assuming the view of vagueness as semantic indecision. Take the expression 'Tibbles.' This is the name of the cat on the map, we were told. Which of the many cats is, strictly speaking, Tibbles? It is plausible to hold that whatever it is that in the thoughts, experiences and practices of language users determines the meaning of expressions, it fails to determine any one of the cats as the referent of 'Tibbles.' A view of vagueness as semantic indecision with respect to singular terms like 'Tibbles' seems as plausible as with respect to predicates like 'is bald' or 'is a cat.' Thus 'Tibbles' indeterminately refers to any of the cats, and statements containing it are to be regarded as true (false) if and only if they are counted as such by all admissible sharpenings. Hence 'Cat Tibbles is on the mat' is, strictly speaking, true; whereas 'Cat Tibbles has hair h_{17} as a part' is, strictly speaking, false.

To acknowledge the vagueness of 'Tibbles,' and hence the semantic indecision among the different candidates, is independent of whether, in any sharpening, the rest of the candidates that are not selected as the referent of 'Tibbles' are counted or not under the extension of 'is a cat.' In other words, the vagueness as semantic indecision of 'Tibbles' is independent of whether one adopts the so-called "supervaluationist" solution to the problem of the many, or the almost-identity one. Actually, if the overall argument of this paper is right, the defender of the view of vagueness as semantic indecision, with respect to both 'Tibbles' and 'is a cat,' not only *can* but indeed *should* favor the almost-identity—or a more general "many" solution to the problem of the many.

Conclusion

There seem to be serious difficulties for the so-called "supervaluationist" solution to the problem of the many. The almost-identity solution seems to be capable of offering a satisfactory explanation of the counting intuition that is ultimately rejected. And, anyway, the consideration offered by Lewis does not show that it is

need of assistance—but rather at most that it should be generalized, in a predictable way, to cover related cases such as the problem of the two; and that supervaluations are nonetheless required to deal with issues other than the problem of the many.

All in all, one might think that it is the second solution—not the first, and not with assistance from the first—that Lewis could have had in mind all along, his own presentation notwithstanding. But as I said, I would rather avoid entering into Lewisian exegesis here.[17]

Arché—The AHRC Research Centre for the Philosophy of Logic, Language, Mathematics and Mind
University of St Andrews

LOGOS—Grup de Recerca en Lògica, Llenguatge i Cognició
Universitat de Barcelona

dlds@st-andrews.ac.uk

References

Andjelković, M., & Williamson, T. (2000). Truth, falsity, and borderline cases. *Philosophical Topics*, *28*, 211–43.

Barnes, E. (MS). What is ontic vagueness?

Eklund, M. (2005). Deconstructing ontological vagueness. *Ms.* (Presented at the 2005 Central APA.)

Fine, K. (1975). Vagueness, truth and logic. *Synthèse*, *30*, 265–300.

Hudson, H. (2001). *A materialist metaphysics of the human person.* Ithaca: Cornell University Press.

Johnston, M. (1992). Constitution is not identity. *Mind*, *101*, 89–105.

Keefe, R. (2000). *Theories of vagueness.* Cambridge: Cambridge University Press.

Lewis, D. (1986). *On the plurality of worlds.* Oxford: Blackwell.

Lewis, D. (1993). Many, but almost one. In J. Bacon, K. Campbell, & L. Reinhardt (Eds.), *Ontology, causality and mind.* Cambridge University Press. (Reprinted in *Papers in Metaphysics and Epistemology*, Cambridge University Press, 1999 (q.v.).)

López de Sa, D. (MS). Bivalence and (Tarskian) truth and falsity.

Lowe, E. J. (1995). The problem of the many and the vagueness of constitution. *Analysis*, *55*, 179–182.

[17]Thanks to the members of the *Arché* research project on vagueness, to the participants of the *LOGOS* discussion group on vagueness, and to the audiences at the universities of Rio de Janeiro and Murcia. I owe special gratitude for criticisms and suggestions to Elizabeth Barnes, Eline Busck, Marta Campdelacreu, Ramiro Caso, Pablo Cobreros, Fabrice Correia, Richard Dietz, José A Díez, Manuel García-Carpintero, Patrick Greenough, Mario Gómez-Torrente, Andrea Iacona, Carrie Jenkins, Andrew Jorgensen, Max Kölbel, Sebastiano Moruzzi, Joan Pagès, Marco Ruffino, Pablo Rychter, Robbie Williams, Timothy Williamson, Crispin Wright, Elia Zardini, Ezequiel Zerbudis, and anonymous referees. Research has been funded by research projects HUM2004-05609-C02-01 (MEC) and BFF2002-10164 (ESF), and grant EX2004-1159 (MEC).

McGee, V. (1998). 'Kilimanjaro'. *23*, 141–63.
McGee, V., & McLaughlin, B. (2000). The lessons of the many. *Philosophical Topics*, *28*, 129–51.
McKinnon, N. (2002). Supervaluations and the problem of the many. *Philosophical Quarterly*, *52*, 320–39.
Morreau, M. (2002). What vague objects are like. *99*, 333–61.
Sider, T. (2001). Maximality and intrinsic properties. *Philosophy and Phenomenological Research*, *63*, 357–64.
Sider, T. (2003). Maximality and microphysical supervenience. *Philosophy and Phenomenological Research*, *66*, 139–49.
Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy*, *5*, 411–67.
Varzi, A. (2001). Vagueness in geography. *Philosophy & Geography*, *4*, 49–65.
Weatherson, B. (2003). Many many problems. *Philosophical Quarterly*, *53*, 481–501.
Weatherson, B. (2004). The problem of the many. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy.* (http://plato.stanford.edu)
Williams, J. R. G. (2006). An argument for the many. *Proceeding of the Aristotelian Society*, *106*.
Williamson, T. (1994). *Vagueness*. London: Routledge.