

[Forthcoming in *Philosophical Perspectives*]

The Hard Problem of the Many

Jonathan Simon

Abstract

A problem of the many Fs arises in cases where intuitively there is precisely one F, but when you look closely you find many candidates for being that F, each one apparently as well-qualified as the next. The problem arises for mundane things like rocks, houses, and coins. It also arises for entities that present special philosophical challenges, like persons and experiencers. In this essay, I present a new argument that the problem of the many experiencers is an especially hard problem of the many, and that property dualism — the view that properties that there is *something it is like* to instantiate are irreducible — may be the best way to solve it. The argument given here turns primarily on normative (i.e., moral) considerations, and is independent of existing arguments for property dualism such as the conceivability and knowledge arguments. It is also independent of existing arguments deriving metaphysical conclusions from the problem of the many experiencers (or related problems of the many), such as those in Unger (2004) and Zimmerman (2010).

Introduction

A problem of the many Fs arises in cases where intuitively there is precisely one F (in the region you are talking about), but when you look closely you find many candidates for being that F, each one apparently as well-qualified as the next. Imagine an apparently solitary cloud in an otherwise blue sky. Look closer, and you'll see lots of water vapor molecules, with no one collection of them more eligible than the others to count as *the* cloud.¹ And many things are like this when you look closely enough. The problem arises for mundane things like rocks, houses and coins. It also arises for entities that present special philosophical challenges, like persons and experiencers. In this essay, I present a new argument that the problem of the many experiencers is an especially hard problem of the many, and property dualism — the view that properties that there is *something it is like* to instantiate are irreducible — may be the best way to solve it.

There are a range of solutions to problems of the many. Here, I'll partition the salient options into three categories: uniform solutions, epistemic solutions, and magnetic solutions. According to uniform solutions to the problem of the many Fs, every candidate stands in the same relation to Fness. On some uniform solutions, every candidate is determinately an F; on other uniform solutions, every candidate is determinately not an F; and on yet other uniform solutions, for every candidate it is indeterminate whether it is an F. Many distinct solutions to problems of the many are uniform in this sense, including nihilism (Unger 1980), relative identity

¹The problem was introduced by Geach (1980) and Unger (1980). For comprehensive treatments see Hudson (2001) and Weatherston (2014).

theory (Geach 1980), manyism (Lewis 1993), partism (Hudson 2001), fractionalism (Bostrom 2006), and supervaluationism (McGee and McLaughlin 2000, Weatherston 2014).

According to epistemic solutions to the problem of the many Fs, precisely one of the candidates satisfies the predicate ‘F’, even though there is no metaphysically significant difference between the winner and its rivals. Epistemic solutions to problems of the many are related to epistemicist theories of vagueness like those in Williamson (1994) and Sorenson (2001), though one need not subscribe to an epistemicist theory of vagueness in order to endorse an epistemic solution to a problem of the many.

According to magnetic solutions to the problem of the many Fs, precisely one of the candidates satisfies the predicate ‘F’, and this winning candidate is inherently better qualified than the others, in light of some metaphysically significant difference between it and the others. This may be because some physical threshold singles out one candidate — roughly the way that an event horizon precisely delineates the boundary of a black hole — but it may also be because of some strictly *metaphysical* threshold, not accounted for by anything uncovered by natural science. For example, it may be that only one candidate *composes a sum*, or that only one candidate *constitutes an enduring individual*, or that only one candidate *is* an enduring individual.² It also may be because ‘F’ denotes some special, emergent property that only one of the candidates instantiates.³

²I include solutions that appeal to ontic vagueness here.

³Technically there are solutions according to which some proper subset of the number of candidates greater than one satisfy ‘F’. These solutions may be either epistemic or magnetic, or both. I note that while I mean for these categories to be more or less exhaustive, they are not exclusive. Hudson (2001)’s partist solution is arguably uniform, epistemic, *and* magnetic.

I will proceed as follows. First, in §§1-3, I will argue that no uniform or epistemic approach can solve the problem of the many experiencers. In §1 I consider uniform solutions according to which each candidate is (determinately) an experiencer, and uniform solutions according to which each candidate is (determinately) not an experiencer. In §2 I consider uniform solutions according to which for each candidate it is indeterminate whether it is an experiencer. In §3 I consider epistemic solutions.

Crucially, my arguments will not tell against the applicability of uniform or epistemic solutions to most other problems of the many. Rather, I will draw on special normative features of the experiential case to give special reasons why uniform or epistemic solutions cannot solve the problem of the many experiencers, even if those solutions can solve other problems of the many.⁴

I will then argue, in §4, that for empirical reasons no magnetic solution to the problem of the many experiencers is available which appeals to purely physical thresholds. Any magnetic solution to the problem of the many experiencers must appeal, somehow or other, to strictly metaphysical thresholds — thresholds not accounted for by anything uncovered by natural science.

Finally, in §5 I will argue that, since we are forced to appeal to metaphysical thresholds only because of special features of the problem of the many experiencers, there will be extra credit for solutions to the problem of the many experiencers which appeal to metaphysical thresholds that arise only in the experiential case — that is, to solutions which do not posit more metaphysical surprises than are necessary to solve this problem in particular. For this reason, I will suggest, the arguments of

⁴My arguments may extend also to the problem of the many persons. I will consider this as it arises.

this paper recommend property dualism — the view that properties that there is *something it is like* to instantiate are irreducible — or anyway, that is what the arguments recommend if you do not have independent reasons for thinking that some other more generic form of metaphysical threshold which does the job is called for.⁵

Thus, considered as a whole, this paper amounts to a novel argument for property dualism (though its individual components, I hope, are of interest on their own). As such, it invites comparison to the knowledge and conceivability arguments. But the premises of the argument here are wholly independent of the premises of the knowledge and conceivability arguments. Instead of appealing to the conceivability of zombies, and some kind of bridge between what is conceivable and what is possible, here I appeal to normative principles about the significance of phenomenal states (like pain), and empirical facts about where physical thresholds are to be found.

I note that I am not the first to recognize that the problem of the many experiencers raises distinctive challenges. Other notable entries in the genre include Merricks (1998), Unger (2004), Monton and Goldberg (2006), Zimmerman (2010), and Bynoe and Jones (2013), and related worries appear in Olson (1997) and Hudson (2001). But the arguments I will advance are wholly independent of those advanced by any of these authors, and in particular none make use of the normative considerations on which my argument hinges. The normative arguments I develop parallel arguments found in Briggs and Nolan (2015), Olson (2010) and Zimmerman (2002), in ways I discuss below, but those arguments are for very different conclusions.

⁵As, for example, Lowe (1995) and Markosian (1998) maintain.

I also note that my focus is on the problem of the many *experiencers* rather than the problem of the many *minds*. I take something to be an experiencer if and only if it instantiates a phenomenal property — a property that there is something it is like to instantiate. I do not assume that experiencers are *minds* — i.e., continuant substances whose experiences are unified.

1 They're Not All Experiencers:

Against Uniform “Manyist” Solutions

In this section I will advance an argument against uniform solutions to the problem of the many experiencers according to which every candidate is an experiencer, the sort of solution championed by Lewis (1993). I do not contest the applicability of analogous “manyist” solutions to other problems of the many; my arguments are specific to the problem of the many experiencers (though we arguably may adapt them to the problem of the many persons, a point to which I will return below).

Say that an *experiencer-cloud* is a maximal set of all candidates for being an experiencer that significantly overlap one another. I will use bracket notation ‘ $[x]$ ’ to denote the experiencer-cloud containing x .

Now, we do not ordinarily think of experiencer-clouds (or person-clouds) as the beneficiaries of our actions, as opposed to experiencers (or persons). But many of the actions we take will equally affect all members of a given experiencer-cloud. In tickling me, you also tickle everyone who massively overlaps me except for a few molecules.

For most intents and purposes, we can treat experiencer-clouds — at least person-shaped ones — as equal, while remaining neutral on how many experiencers (or persons) they contain. If John and Jane are distinct people who deserve equal moral treatment, then we should seek, in our action, to treat [John] and [Jane] equally. That is, we do not go wrong, morally, by ignoring the fact that where there is an experiencer there is an experiencer-cloud:

(EQUAL TREATMENT): : Fully distinct (person-shaped) experiencer-clouds deserve equal moral treatment, other things equal.

EQUAL TREATMENT can be thought of as a cancelling-out principle. So construed, it says that the extra members of [John] and the extra members of [Jane] cancel out, as far as the evaluation of action and the calculation of value is concerned.

This is not to assume that [John] and [Jane] will be equinumerous. To the contrary, biology is messy. Assuming that for every redundant neuron near John's brain — a neuron such that the loss or gain of it would not make a difference to John's phenomenology — there is some number of members of [John] that differ from one another only over whether they contain this neuron, then biological variation between John and Jane will mean that [John] and [Jane] may differ significantly in cardinality. For example, Jane may have a hyperactive hippocampal subgranular zone (a region responsible for adult neurogenesis) while John may not, in which case we may expect that $||\text{[Jane]}|| \gg ||\text{[John]}||$.⁶

⁶One proviso: the number of experiencers in an experiencer-cloud might be infinite, if we may differentiate candidates in some continuous manner. But we may then compare group size by *measure*, and in all likelihood some experiencer-clouds will take up more space than others.

But the facts of biological variation — that one individual may feature more redundant neurons, or greater brain volume than another — do not seem to be of any moral importance at all. In other words, it does not matter morally how many redundant neurons one has. Thus, EQUAL TREATMENT.

But now suppose that both John and Jane are experiencing an equally intense pain. And suppose that manyism is correct — that is, that each member of [John] is an experiencer, and so is each member of [Jane]. But then, assuming that experience is determined by brain state, it is highly plausible that each member of [John] and each member of [Jane] is experiencing an equally intense pain.

The trouble is that on most plausible moral theories, at least in situations where other things are equal, if you can either relieve the suffering of n experiencers or the suffering of m experiencers where m is greater than n , you are obliged to do the latter. This is apparently true for hedonic consequentialist theories.⁷ But non-consequentialists also have reason to accept it, since it falls out of the general duty of beneficence, at least in those cases where other things are equal, meaning at minimum that no one's suffering is more deserved than anyone else's, the agent stands in no special relation to any affected patient, and no side-constraints would be violated in assisting the greater number.⁸ Let us call this principle:

(HEDONIC BENEFICENCE): If option ϕ relieves the pain of n experiencers, option ψ equally relieves the pain of m experiencers, and $n < m$, then other things equal one ought to do ψ rather than ϕ .

⁷E.g. Mill (1862).

⁸For general discussion of the duty of beneficence see Beachamp (2013). For a non-consequentialist explanation of why you should save the many rather than the few see Dougherty (2013), for the opposing view see Taurek (1977).

EQUAL TREATMENT and HEDONIC BENEFICENCE are not inconsistent. But they make for an inconsistent triad with my target in this section: the claim that all candidate experiencers in an experiencer-cloud are indeed experiencers.

For suppose that the only options available amount to either giving [John] a painkiller or giving [Jane] a painkiller — you cannot do both. Without loss of generality, because biological variation is the rule, we may assume that $[[\text{Jane}]] > [[\text{John}]]$. So if manyism is true (if every member of [Jane] is an experiencer and so is every member of [John]), then HEDONIC BENEFICENCE implies that you ought to give the painkiller to [Jane], while EQUAL TREATMENT implies that you should be indifferent.

Where have we gone wrong? Given HEDONIC BENEFICENCE, if every candidate experiencer is an experiencer, then the normative fact about which experiencer-cloud should get the aspirin is settled by which has more spare neurons (even where this has no effect on the intensity or character of the experiences in question).

But as EQUAL TREATMENT is meant to codify, this is an abhorrent conclusion. A hyperactive hippocampal subgranular zone does not make you a utility monster. It does not matter morally how many redundant neurons happen to be in the vicinity of your brain.⁹

⁹Compare McKinnon (2002)'s case of the many coins. Suppose we accept that wherever there is one dollar coin there are many. It follows that if I have one dollar in my pocket, I have thousands of dollars in my pocket (and the exact number will hinge on the exact number of equally qualified candidates). Why aren't I rich? The answer is that money is a social object whose value is a function of a convention that we are all implicitly party to. If it is granted that there are swarms of overlapping coins wherever there is one, what follows is that we must ultimately understand that convention as being a convention to treat coin-shaped swarms of particles as the units of value rather than single coins, strictly speaking, at least in typical cases. But this is not an option in the case of interest to us, since HEDONIC BENEFICENCE is not merely a convention that we may amend at will.

But manyists must either accept the abhorrent conclusion and reject EQUAL TREATMENT, or reject HEDONIC BENEFICENCE. Which should it be?

In favor of rejecting EQUAL TREATMENT, we might argue that it derives its plausibility from our tacitly assuming that there is only one pain-experiencer per swarm. If manyism is correct and [Jane] really does contain twice as many pain experiencers as [John] contains, then perhaps EQUAL TREATMENT loses its plausibility.

There is something to this line of thought. Some metaphysical pictures support the intuition that there is a substantive moral difference between [John] and [Jane]. For example, if each eligible candidate experiencer corresponds to a distinct immortal soul, meaning that [Jane] corresponds to a legion of suffering souls twice the cardinality of [John]'s, then perhaps differential treatment is indeed warranted.

Note though that on such a picture, [John] and [Jane] differ over more than redundant neurons. There are further facts, about souls, which correlate with but do not reduce to the facts about redundant neurons. Plausibly these further facts do matter morally. But manyists who endorse reductive materialism can make no such appeal. According to reductive materialism there is nothing more to the difference between [John] and [Jane] than their differences in redundant neurons, brain volume, and so on. But then the principle that such redundancy does not matter morally entails that EQUAL TREATMENT cannot be abandoned. So if reductive materialism is true, anyway, then EQUAL TREATMENT cannot be abandoned.

The remaining option for manyists is to stand by EQUAL TREATMENT and reject HEDONIC BENEFICENCE. On this view, each experiencer-cloud matters the same, and thus, partisans of more densely packed experiencer-clouds each have less moral weight

than do partisans of sparsely packed ones (in roughly the way that a U.S. citizen’s vote counts for more if that citizen lives in a rural area).

We might try to support this view by claiming that all of the experiencers in a given experiencer-cloud share the same mental states, i.e. experiences. The suggestion, then, would be that we should count by *experiences*, rather than *experiencers*.

Alas, there is a problem of the many experiences that runs in perfect parallel with the problem of the many experiencers. A given experience, say a pain state, will in general consist of, or at least correlate with, a pattern manifested by some neural assembly. But any such assembly will have spare parts, at least if we look at a fine enough grain. Indeed even patterns within a single cell (e.g. one of Koch’s “grandmother cells”) will have spare parts if we look to the molecular level. So without loss of generality, we can assign a numerically distinct experience-candidate to each experiencer-candidate in a given experiencer-cloud.¹⁰

So that gloss may not help. Still, even without the gloss, rejecting HEDONIC BENEFICENCE has the advantage of allowing the manyist to keep calm and carry on distributing aspirin in the sane way. However, it has troubling implications. For HEDONIC BENEFICENCE cannot simply be rejected; it must be replaced with some other principle, stating what beneficence in fact demands. If experiencers matter less when they almost entirely overlap other experiencers, do they also matter less when they only overlap other experiencers a little bit? Where do we draw the line? Do Siamese

¹⁰Subtleties arise if we embrace a bundle theoretic conception, one according to which there are no ‘unified’ minds in any deep sense, but instead there are simply interrelated clusters of individual mental events realized by related brain areas. But this approach only multiplies our difficulties. If an ‘experiencer’ is a combination of experiences, but each experience in the combination overlaps a swarm of others that are qualitatively just like it, then we seem to have at least as many experiencers as there are combinations containing one experience from each swarm.

twins matter less before they are successfully separated than after?

It would help if we had some independent criterion of beneficence, one that makes no reference to experiencers. For example, maybe we should count by persons, rather than experiencers. This raises questions of course: do non-persons that suffer really have no moral status?

The real issue is that we are currently considering a challenge for those who defend the manyist solution to the problem of the many experiencers. There is also a problem of the many persons (on which more has indeed been written than on the problem of the many experiencers, see e.g. Hudson 2001). If we take the manyist approach to that problem, then changing the loci of value from experiencers to persons brings us nowhere, because biological variation also means variation in the cardinalities of person-clouds. On the other hand, if we pursue some non-manyist approach to the problem of the many persons, why not deploy that approach here?

Incidentally, Briggs and Nolan (2015) identify a challenge for some theories of personal identity which parallels the challenge for manyism about the problem of the many experiencers that I consider here.¹¹ Briggs and Nolan's challenge confronts theories according to which fission separates persons who were spatiotemporally coincident until the fission event, and the number of coincident persons is determined by the number of future fission events the person(s) present will undergo. On such theories, that is, the more fissions one is scheduled to undergo the more persons one is currently coincident with. Thus if we count moral status by persons, the more fissions one has scheduled, the more one matters morally (or anyway, the more the

¹¹See also Zimmerman (2002) and Olson (2010) for other related challenges.

collection-of-persons-that-coincide-with-one matters morally) — even setting aside the problem of the many persons.¹²

Unfortunately, the more promising solutions to Briggs and Nolan’s problem do not translate into promising solutions to ours. Their problem is resolved (bracketing problems of the many) if we hold that person-stages, rather than persons, are the loci of value — a position which is plausible insofar as persons are nothing over and above their stages (which many advocates of the affected views of personal identity hold). A related option would be to take experiencers to be the loci of value, and hold that experiencers are stage-bound entities, such that all of the perfectly coincident people have the same experiencer as a temporal part.

But clearly neither of these options is any use to us. We are looking for a rationale for revising HEDONIC BENEFICENCE. HEDONIC BENEFICENCE already says that experiencers are the loci of value, so that response to the Briggs - Nolan problem does not help. But neither does saying that person-stages are the loci of value, for there is a problem of the many person-stages that parallels the problem of the many persons. Perhaps we have some non-manyist solution to that problem, but then why not deploy it to solve the problem of the many experiencers? Our current focus is a challenge confronted by manyist solutions.

In sum, the problem is that there is no obvious replacement for experiencers as loci of value which is not afflicted by a comorable version of the problem. I conclude that neither option for the manyist is comfortable: it is no good giving up either on EQUAL TREATMENT or on HEDONIC BENEFICENCE. But this means we should

¹²Though see their footnote 5.

abandon manyism, that is, the supposition that each candidate experiencer really is an experiencer.

Moreover, the argument I have given also covers the possibility, suggested by Bostrom (2006), that the number of experiencers is *fractional*. I do not claim that Bostrom has succeeded in giving sense to the possibility that the number of experiencers in a collection of n individuals can be anything but a whole number less than or equal to n . But if Bostrom's suggestion does make sense, and if each candidate contributes the same fraction to the total tally, then the exact reasoning above suffices, provided we modify HEDONIC BENEFICENCE by replacing natural numbers ' $n < m$ ' with real numbers ' $r_1 < r_2$ '.

Similarly, the argument above covers Hudson (2001)'s *partist* approach according to which the many are different whole locations of one and the same thing. If one and the same thing can have different properties at different (whole) locations at the same time, then there is at least a case that we should count by manifestations-of-individuals rather than individuals. To illustrate, suppose we hold that in cases of self-visitation through time travel we have multiple manifestations of the numerically same individual. Suppose you are at a time where there exist three manifestations of Marty McFly and two manifestations of Doc. All five are in equally intense pain, and somehow you can either relieve the pain of the three Martys or the two Docs. Other things equal, you do more good by helping the Martys!

This also suggests a response to those who address my argument by appealing to something like Geach (1980)'s notion of relative identity. According to this approach, we hold that, even though there are many numerically distinct material composites

which are elements of [John], each of them is *the same experiencer* as all of the others. Crucially this is more than just the fairly innocuous claim that if manyism is true it is sometimes pragmatic to count loosely, and say there is only one cat on your rug, even if strictly (because manyism is true) there are 1001. HEDONIC BENEFICENCE is not a principle concerned with the pragmatics of communication, but with the underlying facts about the right and the good.

However if there is really a deep or joint-carving sense in which all of the experiencers contained in [John] are the same experiencer (despite being distinct material composites) then there is a case to be made that it is double-counting to take each of them to be a distinct locus of value, and so a case to be made that HEDONIC BENEFICENCE is false (if we count strictly).¹³

But what goes for partists may also apply here. For in the time travel scenario above, the three Martys are plausibly all the same experiencer, and likewise for the two Docs. If anything, finding out that each of the Martys and each of the Docs is a distinct material composite (though the same experiencer) only strengthens the intuition I am advocating here.

I conclude that rejecting manyism is the most attractive resolution of the conflict between HEDONIC BENEFICENCE and EQUAL TREATMENT.¹⁴ This covers many uniform solutions, but it leaves nihilism (Unger 1980), where we deny that there are any experiences or experiencers, and it leaves uniform solutions where we hold that for each candidate it is indeterminate whether it is an experiencer. I will finish this section with a discussion of nihilism, and address indeterminacy-based uniform solutions in

¹³My thanks to Kristie Miller and Kris McDaniel for (independently) pressing me on this point.

¹⁴At least supposing that reductive materialism is true, on which more below.

the next section.

Let me distinguish between weak nihilism and strong nihilism. Weak nihilism denies that any *single* thing instantiates phenomenal properties, but allows that pluralities can plurally instantiate them. Weak nihilism is compatible with a range of solutions to the problem of the many experiencers: weak nihilist manyists face the challenge I have just presented. Weak nihilist epistemicists face the challenge for epistemicists that I present below, and so on.

Strong nihilism denies that phenomenal properties are instantiated, by unities or pluralities. But Cartesian considerations address the strong nihilistic view. You have special insight into the fact that you are an experiencer. Thus you are in a position to know that there is at least one experience associated with your experiencer-cloud.

2 Some Are Definitely Experiencers:

Against Uniform “Indeterminist” Solutions

In this section I will argue against uniform solutions according to which for each candidate, it is indeterminate whether it is an experiencer. I will argue that a modified version of HEDONIC BENEFICENCE generates a conflict for these solutions much like the conflict I presented in the previous section. Note that some views in this category — supervenient views — maintain that (determinately) there is only one experiencer in each experiencer-cloud. As we will see, however, this makes no difference for our purposes.

As an aside, I note that in Simon (2017), I argue that ‘experiencer’ cannot be

vague, in the sense that there can be no such thing as a borderline case of ‘experiencer’. My arguments there do not rule out all forms of indeterminacy, but plausibly the sort of indeterminacy that is in play here is similar enough to vagueness to be covered by my argument. However, I mean here to rely neither on that argument, nor on the claim that the indeterminacy in question here amounts to vagueness.

Here, I assume only that the views in question (uniform indeterminist solutions to the problem of the many experiencers) construe ‘indeterminately’ as a status incompatible with truth, and incompatible with falsity. If some candidates that are indeterminately F are F, while others are not F, then not all candidates have the same status with respect to ‘F’, and the solution in question is non-uniform (epistemic or magnetic), rather than uniform.

With ‘indeterminacy’ so construed, what is to be said about the normative status of a being such that it is indeterminate whether it is feeling pain or nothing at all? As before, consider a case in which (as we would ordinarily describe it) both John and Jane are in intense pain, and there is only pain medication enough for one of them. On the kind of uniform solution we are currently considering, this translates into a case in which for every member x of [John], it is indeterminate whether x is experiencing anything, but determinate that if x is experiencing anything x is experiencing intense pain, and likewise for every member y of [Jane]. Suppose we give the medicine to Jane. It is potent, meaning that after “she” takes it, for each member y of [Jane], it is indeterminate whether y experiences anything, but it is determinate that if it is experiencing anything then it is experiencing pleasure.

As before, we are assuming that the cardinality of experiencer-clouds depends

on biological variation over brain volume and the like. So the choice of whom to give the medication to is again a matter of numbers: option ϕ converts n cases of indeterminate-whether-feels-nothing-or-feels-pain into n cases of indeterminate-whether-feels-nothing-or-feels-pleasure, while option ψ converts m cases in the same manner, and $n < m$. Here, as before, it seems that other things equal, you had ought to help the greater number rather than the lesser. In other words:

(HEDONIC BENEFICENCE*): If option ϕ brings us from a situation in which for each of n things, it is indeterminate whether they are in agonizing pain or feel nothing, to a situation in which it is indeterminate whether they feel great or feel nothing, and option ψ brings us from a situation in which for each of m things, it is indeterminate whether they are in agonizing pain or feel nothing to a situation in which it is indeterminate whether they feel great or feel nothing, and $n < m$, then other things equal one ought to do ψ rather than ϕ .

So we have that if it is indeterminate, for each candidate, whether it is an experiencer, standard biological variation will lead to conflicts between HEDONIC BENEFICENCE* and EQUAL TREATMENT.

We have already discussed the costs of rejecting EQUAL TREATMENT. Is it easier to reject HEDONIC BENEFICENCE* than it is to reject HEDONIC BENEFICENCE? Recall that the challenge with rejecting HEDONIC BENEFICENCE is that one ends up treating some experiencers as of lesser value than others, even if they are qualitative duplicates, solely as a function of how many other experiencers they overlap.

Here too, the intuition is one of equality: that every candidate experiencer should have the same moral status as every other — or at least, every other that is the same

in all mental respects. As before, there is no obvious replacement locus of value.

I note that the problem here is a problem even on supervaluation-style semantics for indeterminacy, according to which ‘There is exactly one mind for each experiencer-cloud’ comes out determinately true. That semantic result does not offer us any rationale for rejecting EQUAL TREATMENT OR HEDONIC BENEFICENCE*.

I stress that nothing I have said above challenges indeterminist (or manyist) solutions to most problems of the many. Principles of beneficence do not govern our behavior towards clouds. Principles of beneficence do arise in some other cases: for example, the problem of the many persons. But it is debatable whether there are really distinct problems here. Would principles of beneficence apply to persons even if they were phenomenal zombies? Do zombie lives matter? If not, then the problem of the many experiencers is the unique source of the challenge, else, the problem of the many persons turns out to independently pose an analogous challenge, but only because of features that clearly do not generalize to most other cases.

3 It’s Not What You Don’t Know: Against Epistemic Solutions

Generally speaking, an epistemic approach to vagueness or a problem of the many is an approach according to which there is a sharp cut-off, and a single candidate in the “cloud” which satisfies the relevant predicate, but we are inescapably ignorant about which one it is.¹⁵

¹⁵Prominent defenders include Williamson (1994) and Sorenson (2001). See also Hudson (2001).

Now, my interest here is with approaches that are epistemic in this sense, but which also deny that there is any metaphysically significant difference between the winning candidate and the losers, in light of which the winner is inherently more qualified than its fellow candidates. I will not offer a precise account here of *metaphysically significant difference*, but I take it to be a sufficient condition that one candidate is somehow *more natural* than the others in the sense of Lewis (1983), or singled out by a joint in nature. Another sufficient condition for there to be a metaphysically significant difference between two things is for them to be sufficiently ‘far’ from one another in a suitably joint-carving quality space (meaning that sufficiently many insignificant differences can compound into a significant difference).

Not all epistemic approaches to problems of the many deny that there are metaphysically significant differences between the winning candidate and the losers. For example, Hawthorne (2006) sketches a “metaphysically inflationary” version of epistemicism according to which the reference relation is a highly natural relation, though we are ignorant of some of the details about its extension, much as we are ignorant of some mathematical truths (according to the Platonist).

Insofar as there is an instance of a highly natural (semantic) relation involving one of the candidate experiencers, and there are no such instances involving any of the others, there is for my purposes a metaphysically significant difference between that candidate and the others. Hawthorne’s inflationary epistemicism will count for my purposes as a magnetic solution. From here on, when I speak of epistemic solutions, I have non-magnetic epistemic solutions in mind — ones according to which being singled out as the referent of some vague term ‘F’ is not metaphysically significant

in and of itself.

My argument against uniform (non-magnetic) epistemic solutions to the problem of the many experiencers turns on the following two principles:

(NORMATIVE SUPERVENIENCE): Normatively significant differences supervene on metaphysically significant differences.

(PAIN MATTERS): There is a normatively significant difference between an entity which is experiencing intense pain, and an entity that is not experiencing anything at all.

By a normatively significant difference I mean a major difference in moral status, like the difference between something being of no moral relevance at all, and something being very bad. I stress that **NORMATIVE SUPERVENIENCE** does not deny that metaphysically insignificant differences can *cause* metaphysically (and normatively) significant ones. Nuclear bombs go off because of microscopic alterations to individual atoms, after all.

It is also consistent with **NORMATIVE SUPERVENIENCE** that there is a normatively significant difference between two cases that are *intrinsically* more or less the same. If I desire A but not B, then even if A and B are intrinsically more or less the same, there is a metaphysically significant difference between them, thanks to the asymmetry of my desire. Metaphysically significant differences (between A and B) need not be intrinsic to A and B.

Now, it might seem as though **NORMATIVE SUPERVENIENCE** would be anathema to any thorough-going epistemicist, at least to a non-magnetic epistemicist. For if terms

like ‘obligation’ are vague, an epistemicist treatment would likely entail sharp cut-offs in whether something is obliged or not, that do not correspond to any metaphysically significant differences. However, it is not a foregone conclusion that terms like ‘obligation’ are vague. One might follow Dworkin (2013) in endorsing a “Right Answer Thesis” according to which normative questions all admit of (knowable) answers, and so deny the vagueness of the relevant terms.

In any event, there is real cost to denying NORMATIVE SUPERVENIENCE: it is unclear how to deliberate without it. When we seek to justify our actions to ourselves and others, we seek to find salient descriptive differences between cases we treat one way and cases we treat another. But it is a necessary condition on such a difference being morally salient that it is metaphysically significant (in the relatively weak sense that I have in mind).

Consider the line of reasoning that begins when Peter Singer asks you to justify throwing away the letter from UNICEF. You attempt to find descriptive differences that stand a chance of being morally salient, between that case and the case of the drowning child. Plenty of differences come to mind that are metaphysically significant — but that nevertheless are clearly morally irrelevant. One example is the difference between your distance to the child drowning in the park, and your distance to the child starving in Bangladesh.

However, it is even easier to see the moral irrelevance of metaphysically insignificant differences. Imagine that someone claims that whether or not it is permissible to ignore a letter from UNICEF hinges on the precise margins of the envelope. This is what NORMATIVE SUPERVENIENCE denies. The metaphysical significance of a difference

is certainly not sufficient for its moral salience, but it is necessary.¹⁶

What about PAIN MATTERS? May it be rejected? One intriguing development of this idea (in a somewhat different context) is what Lee (2012), (2014) calls the deflationary response. On this response we deflate the normative distinction between conscious beings and nearby *quasi-conscious* beings (where a being is *quasi-conscious* just in case it is not conscious but is very similar to beings that are conscious). The deflationary response motivates the rejection of PAIN MATTERS, holding that if the pain of conscious beings matters, then the quasi-pain of quasi-conscious beings matters just as much, or almost as much.

Unfortunately that generous extension of moral significance to non-conscious beings will not help us here: it brings us back to the dilemma confronted by manyists. For the deflationary position suggests that we must consider the standing of the quasi-conscious along with the conscious when we assess the impact of our actions. But that motivates:

(QUASI-HEDONIC BENEFACTENCE): If option ϕ relieves the pain or quasi-pain of n experiencers or quasi-experiencers, option ψ equally relieves the pain or quasi-pain of m experiencers or quasi-experiencers, and $n < m$, then other things equal one should do ψ rather than ϕ .

And this brings us back into conflict with EQUAL TREATMENT. I conclude that there is no good epistemic solution to the problem of the many experiencers, even if there may be epistemic solutions to other problems of the many. Once again, the same may

¹⁶See the related criticisms of epistemic accounts of moral vagueness in Sider (2002), Dougherty (2013) and Schoenfeld (2016).

go for the problem of the many persons. As before, this depends on whether we think that personhood is what matters, even divorced from phenomenal consciousness.¹⁷

4 No Physical Thresholds: Against Reductive Materialist Magnetic Accounts

So a few related normative considerations show us that there will neither be uniform solutions nor (non-magnetic) epistemic solutions to the problem of the many experiencers. This means that, if we hope not to overturn any normative precepts, we must opt for a magnetic solution. Magnetic solutions, again, are those according to which the winning candidate wins because it is inherently better qualified than the others: there is a metaphysically significant difference between it and the others.

Magnetic solutions fall into two categories: those that invoke metaphysically significant distinctions that the underlying physics does not suggest (strictly metaphysical thresholds), and those that only rely on distinctions that the underlying physics does suggest (physical thresholds). In this section I will argue on empirical grounds that physical thresholds are unavailable to solve the problem of the many experiencers.

Physical thresholds generally are not defined with infinitesimal precision. Nothing

¹⁷The case against nihilism about persons may also be weaker than the case against nihilism about experiencers. And it is easier to see that personal identity might not matter morally (Parfit 1984) than it is to see how intense pain might not matter morally. But if there are principles of beneficence that apply to persons irrespective of their phenomenal lives, arguments analogous to those given here could certainly be constructed.

like an event horizon exists to carve out one collection of water vapor molecules as best-suited to the role of being a cloud. At a fine enough level of resolution the boundary between the cloud and the sky becomes murky. A single straw may break the camel's back, but which molecules exactly are the ones that compose that straw?

It is also worth noting that appealing to *maximality* does not lead us to the sought for physical thresholds. It seems that many of the terms and concepts that give rise to problems of the many have an in-built maximality condition requiring that if something is an F, it is not a large proper part of another F.¹⁸ This principle is plausible, and should not be confused with the claim the real F is the fusion of all of the other candidate Fs. We have in general no guarantee that the fusion of all of the candidates is even a candidate, let alone the most eligible one (it might be too big). Why favor the fusion of all candidates rather than their intersection? Moreover, in the general case no sufficiently sharp threshold partitions candidate Fs from non-candidates.

Most theories of the physical basis of consciousness imply that there are no perfectly precise physical thresholds delineating the conscious from the non-conscious. On a global workspace theory, for example, a representation is conscious insofar as it is widely available for consumption by a variety of cognitive systems. Obviously 'wide availability' is not a precise threshold. Even on specific-wavelength theories like the 40 Hz theory, that identify conscious processing with processing in some specific frequency band, there is some tolerance: frequency is measured in the aggregate and a range of variability is inevitable. Of course, one might think of such a theory as

¹⁸See relevant discussions in Merricks (1998), Hudson (2001) and Sider (2003).

being like a qualitative description of a natural kind: we take the actual contour of the final theory to be given by the contour of the natural kind that best fits it. But here we have no good reason to suppose that any relevant natural kind has perfectly precise contours.

Quantum and integrated information (IIT) theories of consciousness come closest. There are many ways that quantum phenomena might play a role in conscious processing. Perhaps the final theory of consciousness must be given partly at the quantum level because quantum interference or collapse phenomena are somehow directly relevant to the makeup of consciousness.¹⁹

According to Penrose and Hameroff (2014)'s Orchestrated Objective Reduction (Orch-Or) theory, coherent quantum processes take place within neuronal microtubules that give rise to consciousness.²⁰ But even this does not yield absolutely precise thresholds that single out one candidate experiencer from each experiencer-cloud, since if there is such computation, there will be a swarm of overlapping systems that equally make use of it.

These authors and others have also posited that quantum mechanics involves a collapse mechanism, and that this mechanism is somehow essentially driven by consciousness.²¹ But to be plausible, this approach requires that consciousness antecedently demarcates a joint in nature. Otherwise the claim would be that there is some intrinsically unremarkable macroscopic property that just happens to play a

¹⁹Antony (2006) argues that, because 'consciousness' is not vague a quantum theory of consciousness may be correct. But Antony's claim is only that *if* there is a correct physical theory of consciousness it must be expressible in precise, quantum language: he does not establish that there exist suitable joints in nature.

²⁰See Tegmark (2000) *pro contra*.

²¹See Atmanspacher (2015) for review.

role in fundamental physics, and this is implausible. A possible alternative might be to *identify* collapse events with conscious events, but this path is not for the faint of heart. Note that, coupled with Tumulka (2004)'s GRW_f theory, the leading relativistic collapse theory, this yields an exotic form of panpsychism. I conclude that delving into the quantum depths is unlikely to reveal the kind of physical threshold we are searching for.

Tononi (2016)'s Integrated Information Theory (IIT) proposes that there is a quantity, Φ , which is a measure of a system's *integrated information*, and such that conscious systems are those which take on local maximum values of Φ . Insofar as the quantity Φ is physically reducible, unambiguously defined, and such that for any swarm of experiencer-candidates there is one with a well-defined minimum Φ value, then IIT may resolve the problem of the many experiencers in a selective but reduction-friendly way.

However, Φ is not unambiguously defined, and it is not such that for any swarm of mind-candidates there is one with a well-defined minimum Φ value. Tegmark (2015) identifies 420 distinct characterizations of Φ . Of these, he argues that six stand out as most attractive. But six is still five too many. Worse, Φ -measures are generally not deployed at the level of granularity that our current task requires. Φ -measures measure integration or non-separability in systems construed, roughly, as causal networks. But assuming that the level of causal efficacy relevant to consciousness is the level of neurons or neuronal assemblies, Φ -measures will be blind to distinctions at the chemical or molecular level. Thus if our favorite Φ -measure assigns apparently minimal Φ to some candidate, the result of subtracting one peripheral molecule and

adding another may yield the same Φ , leaving us with a tie.²²

5 From Metaphysical Thresholds to Property Dualism

So, as an empirical matter of fact, it turns out that no physical thresholds are available. I conclude that the only acceptable solutions to the problem of the many experiencers are magnetic solutions that involve strictly metaphysical thresholds: thresholds not accounted for by anything uncovered by natural science. What are the options?

Many have been canvassed in the literature. To name four, we can restrict composition so that only one set corresponding to a candidate composes a sum (Markosian 1998), or ration out composition so that each set composes the same sum (Hudson 2001). We can restrict constitution so that only one sum constitutes a member of the relevant sortal, though it may be vague which one (Lowe 1995 and Van Inwagen 1990), or we can ration out constitution, so that each sum constitutes the same member of the relevant sortal (Jones 2015).

But if, like me, you hold that uniform or epistemic or physical threshold solutions may be available to solve most problems of the many, and the problem of the many experiencers is more or less alone in requiring a metaphysical threshold solution, then you have reason to eschew solutions that introduce the relevant kind of threshold at the generic level of formal ontology, in contrast to those which introduce it specially

²²See also Aaronson (2015) for an independently troubling critique of IIT: on some reasonable measures of Φ it implies that CD players are more conscious than we are.

for the mental domain.

Moreover, domain-independent solutions that involve rationing out composition or constitution or the like are incomplete, for they still owe an account of the mind-body nexus. After all, one can allow that only one plurality from the cloud composes a sum, while still maintaining that every plurality from the cloud (plurally) instantiates phenomenal properties.²³

In contrast, a property dualist solution faces no such further challenges, and invokes no more metaphysical surprises than are necessary to the case. On this view, there are phenomenal properties that are irreducible to physical properties, and one experiencer-candidate per cloud instantiates them. This candidate is *ipso facto* the only experiencer in the cloud.

Some authors (notably Unger 2004 and Zimmerman 2010) have suggested (for reasons different from those given here) that the problem of the many minds (which includes the problem of the many experiencers) calls for *substance* dualism. I do not make this claim here. Here, I claim only that property dualism is a fitting solution: substance dualism entails property dualism, but property dualism suffices for the thresholds we may need.

There is a further challenge for many metaphysical threshold solutions, however, including the property dualist solution. As Hudson (2001) argues, if one endorses a solution like Markosian (1998)'s brute composition, one has no good reason for maintaining that there is exactly one candidate that composes a sum. This difficulty is a consequence of the claim that composition is brute, i.e., not governed by any

²³My thanks to Maya Eddon for pointing this out to me.

deeper principle. Likewise, if phenomenal properties are irreducible, this suggests that their instantiation is likewise not governed by any deeper principle. What assurance does the turn to property dualism give us that there is indeed exactly one experiencer per experiencer-cloud?

A first thing to say here is that, as I note above, the pressure to accept EQUAL TREATMENT is mitigated if reductive materialism is rejected. That principle is motivated by the thought that redundant neuronal parts do not matter morally. But if something over and above redundant neuronal parts is at stake then arguably that principle loses its plausibility. Thus, even if we have no reason to expect precisely one experiencer per cloud, we have a natural way of resolving the tension faced by uniform solutions.

A second thing to say is that if phenomenal properties are irreducible their instantiation is likely governed by psychophysical law (as opposed to being anomalous). But it is then open that the best laws we can find turn out to involve principles of attraction and repulsion, that secure that phenomenal subjects are both unified, and not too close to one another. I shall not offer arguments that there are indeed such laws here; I am content to note that nothing rules them out.

Of course, for those who are already persuaded that some other more general form of metaphysical threshold is called for, suitable for solving many problems of the many, my argument merely serves as auxiliary support, rather than a case for property dualism. But it is worth noting an analogy in the dialectic of the conceivability argument: at least as Chalmers (2009) develops it, the conceivability argument hinges on a reductive account of metaphysical modality which Chalmers

calls *modal rationalism*. Those who embrace primitivist or emergentist accounts of metaphysical modality are in a position to resist Chalmers' property dualist (or Russellian monist) conclusions, much as those who embrace emergentist solutions to all problems of the many are in a position to resist mine.

A final loose end concerns the connection of the problem of the many experiencers and the problem of the many persons. If the move to property dualism cannot resolve the latter, then more generic machinery may be called for after all. Happily, some have already embarjed on the project of deploying an emergentist account of consciousness in the service of an emergentist account of persons.²⁴

Acknowledgments

My thanks to David Chalmers, Maya Eddon, Dana Goswick, Kris McDaniel, Kristie Miller, Michael Raven and an audience in Graz for comments on a recent ancestor and to Ned Block, Paul Boghossian, Patrick Greenough, Geoffrey Lee, Colin Marshall, Angela Mendelovici, Daniel Nolan, Jeff Sebo, Daniel Stoljar, Peter Unger, and audiences in Canberra, Geneva and New York for comments on more distant ancestors. Finally my thanks to the editors of *Philosophical Perspectives*.

²⁴See Gilmore (2015) for an example of how this might be done, and see McDaniel (forthcoming) p.178-185 for some worries.

References

- Aaronson, S. (2015) “Why I Am Not An Integrated Information Theorist (or, The Unconscious Expander)”. Shtetl-Optimized. (Weblog) url: www.scottaaronson.com. Retrieved 2015-11-23.
- Antony, M. (2008). “Are Our Concepts *Conscious State* and *Conscious Creature* Vague?” *Synthese*.
- Antony, M. (2006). “Vagueness and the Metaphysics of Consciousness.” *Philosophical Studies* 128: 515-538.
- Barnes, Elizabeth, 2010, “Ontic Vagueness: A Guide for the Perplexed,” *Noûs*, 44: 601-627.
- Beauchamp, Tom, “The Principle of Beneficence in Applied Ethics”, *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), forthcoming URL = <https://plato.stanford.edu/archives/win2016/entries/principle-beneficence/>.
- Bostrom, N. (2006) “Quantity of Experience: Brain Duplication and Degrees of Consciousness.” *Mind* 116:185-200
- Briggs, R. & Daniel Nolan, D. (2015). “Utility Monsters for the Fission Age.” *Pacific Philosophical Quarterly* 96 (2):392-407.
- Bynoe, W. & Jones, N.K. (2013). “Solitude Without Souls: Why Peter Unger Hasn’t Established Substance Dualism”. *Philosophia* 41 (1):109-125.
- Chalmers, D.J. (2009). “The Two-Dimensional Argument Against Materialism.” In Brian P. McLaughlin & Sven Walter (eds.), *Oxford Handbook to the Philosophy of Mind*. Oxford University Press.
- Chalmers, D.J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Dougherty, T. (2014). “Vague Value.” *Philosophy and Phenomenological Research* 89(2):352-372.
- Dougherty, T. (2013). “Rational Numbers: A Non-Consequentialist Explanation of Why You Should Save the Many and Not the Few.” *The Philosophical Quarterly*. Vol. 63, No. 252, pp.413-427.
- Dworkin, R. (2013). *Justice for Hedgehogs*. Harvard University Press.
- Evans, G. (1978). “Can There Be Vague Objects?” *Analysis* 38 (4):208.
- Fine, K. (1975) “Vagueness, Truth and Logic.” *Synthese*, vol. 30, pp. 265-300.
- Geach, P. T., (1980), *Reference and Generality*. 3rd edn., Ithaca: Cornell University Press
- Gilmore, C. (2015). “Personal Identity, Consciousness, and Joints in Nature.” *Journal of Ethics* 19 (3-4):443-466.
- Hameroff, S; Penrose, R. (2014). “Consciousness in the universe: A review of the ‘Orch OR’ theory”. *Physics of Life Reviews*. Elsevier. 11 (1): 39-78.

- Hawthorne, J. (2006). "Epistemicism and Semantic Plasticity." *Oxford Studies in Metaphysics* 2:289-322.
- Hudson, H. (2001). *A Materialist Metaphysics of the Human Person*, Ithaca: Cornell University Press.
- Jackson, F. (1982). "Epiphenomenal Qualia." *Philosophical Quarterly* 32:127-136.
- Johnston, M. (1992), "Constitution is Not Identity," *Mind*, 101: 89-105.
- Jones, N. (2015), "Multiple Constitution." In Karen Bennett & Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics* 9. Oxford: OUP. pp. 217-261
- Lee, G. (2014) "Materialism and the Epistemic Significance of Consciousness." in *Current Controversies in Philosophy of Mind*, Kriegel, U. (ed.) Routledge.
- Lee, G. (2012). "Alien Subjectivity and the Importance of Consciousness". In Adam Pautz & Daniel Stoljar (eds.), *Themes from Block*. MIT Press.
- Lewis, D. (1993), "Many, but Almost One," in *Ontology, Causality and Mind: Essays in Honour of D M Armstrong, John Bacon* (ed.), New York: Cambridge University Press.
- Lewis, D. (1983). *New Work for a Theory of Universals*. *Australasian Journal of Philosophy* 61:343-377.
- Lowe, E. J., (1995), "The Problem of the Many and the Vagueness of Constitution," *Analysis*, 55: 179-82.
- Mackie, P. (2011). "Property Dualism and Substance Dualism." *Proceedings of the Aristotelian Society* 111:181-199.
- Markosian, N. (1998), "Brutal Composition," *Philosophical Studies*, 92: 211-49.
- McDaniel, K. (2017). *The Fragmentation of Being*. Oxford University Press.
- McGee, V. and 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.
- Merricks, T. (1998), "On Whether Being Conscious is Intrinsic." *Mind*, 107 (428):845-846.
- Mill, J. S., (1861). *Utilitarianism*, edited with an introduction by Roger Crisp. New York: Oxford University Press, 1998.
- Monton, B. and Goldberg, S. (2006). "The Problem of the Many Minds." *Mind* 115:463-470.
- Olson, E. (2010) "Ethics and the Generous Ontology." *Theor Med Bioeth.* 31:259-270.
- Olson, E. (2002). "Thinking Animals and the Reference of I." *Philosophical Topics* 30 (1):189-207.
- Olson, E. (1997). *The Human Animal: Personal Identity Without Psychology*. Oxford University Press.

- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Sattig, T. (2013), "Vague Objects and the Problem of the Many" *Metaphysica*, 14: 211-223.
- Schoenfeld, M. (2016). *Ethics* 126 (2):257-282.
- Sider, T. (2003). "Maximality and Microphysical Supervenience." *Philosophy and Phenomenological Research* 66: 139-149.
- Sider, T. (2002), "Hell and Vagueness," *Faith and Philosophy* 19 (2002): 58-68.
- Simon, J. (2017). *Vagueness and Zombies*. *Philosophical Studies*. Vol. 174 (8), 2105-2123.
- Sorensen, Roy, (2001). *Vagueness and Contradiction*, Oxford: Oxford University Press.
- Taurek, J.(1977). "Should the Numbers Count?", *Philosophy & Public Affairs*, 6, 293-16.
- Tegmark, M. (2016). Improved Measures of Integrated Information. *PLoS Comput Biol.*12(11).
- Tegmark, M. (2000). "Importance of quantum decoherence in brain processes". *Phys. Rev. E*. 61 (4): 4194-4206.
- Tononi, Giulio; Boly, Melanie; Massimini, Marcello; Koch, Christof. (2016) "Integrated information theory: from consciousness to its physical substrate". *Nature Reviews Neuroscience*. 17 (7): 450-461.
- Unger, P. (2008). *All The Power in the World*. Oxford University Press.
- Unger, P. (2004) "The Mental Problems of the Many." In D. Zimmerman (ed.), *Oxford Studies in Metaphysics*, Vol. 1. Oxford: Clarendon Press.
- Unger, P. (2002). "Free Will and Scientifiphicalism." *Philosophy and Phenomenological Research* 65 (1):1-25.
- Unger,P. (1988). *Conscious Beings in a Gradual World*. *Midwest Studies in Philosophy* 12 (1):287-333.
- Unger, P. (1980). "The Problem of the Many," *Midwest Studies in Philosophy*, 5: 411-67.
- Van Inwagen, P. (1990). *Material Beings*. Ithaca. Cornell University Press.
- Weatherson, B. (2014) "The Problem of the Many", *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2015/entries/problem-of-many/>.
- Weatherson, B. (2003). "Many Many Problems." *Philosophical Quarterly* 53.
- Williams, J. Robert G., (2006), "An Argument for the Many," *Proceedings of the Aristotelian Society*, 106: 411-419.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Zimmerman, D. (2011). "From Experience to Experienter." in Baker, M.C. and Goetz, S. (eds.)*The Soul Hypothesis*, London: Continuum Publishing.

Zimmerman, D. (2010). "From Property Dualism to Substance Dualism." *Aristotelian Society Supplementary Volume* 84 (1):119-150.

Zimmerman, D. (2002). "Persons and Bodies: Constitution Without Mereology?" *Philosophy and Phenomenological Research*, Vol. LXIV, No. 3, 599-606.