Building Thoughts from Dust: a Cantorian Puzzle

Forthcoming in *Synthese*

Abstract

I bring to light a set-theoretic reason to think that there are more (identifiable) mental properties than (identifiable) shapes, sizes, masses, and other characteristically "physical" properties. I make use of a couple counting principles. One principle, backed by a Cantorian-style argument, is that *pluralities* outnumber *particulars*: that is, there is a distinct plurality of particulars for each particular, but not *vice versa*. The other is a principle by which we may coherently identify distinct mental properties in terms of arbitrary pluralities of physical properties. I motivate these principles and explain how they together imply that there are more mental properties than physical properties. I then argue that certain *parody* arguments fail for various instructive reasons. The purpose of my argument is to identify an unforeseen "counting" cost of a certain *reductive* materialist view of the mind.

Building Thoughts from Dust: a Cantorian Puzzle

1. The Question

What sort of thing am I? Perhaps the cleanest and simplest answer is that I am entirely physical in every respect. This answer is a kin to the *type-identity* theory, which implies that all my properties are analyzable in terms of "physical" properties, such as shape, size, mass, and other properties informed by the hard sciences.[1] Or is there more to me than that?

I shall introduce a new strategy for thinking about the nature of persons using principles of counting. The usual objections to reductive physicalist views focus on mental properties, like *happiness* or *thinking carefully*. I will argue instead that the mental properties *outnumber* the physical ones. The conclusion is that one current reductive theory of minds is false.[2] My goal is to present a certain tool for investigation, rather than to give an unassailable argument. Identity theorists may take from my argument a reason to think they must pay a certain "counting" cost, which is the cost of having to posit a breach in certain intuitive counting principles. I will also explain how my argument strategy may be useful for knocking away various *restricted* identity theses, where there is no independent reason to pay the "counting" cost.

[1] The term 'physical property' is a term of art. I have in mind the sort of properties that can be sensed from a "third-person" perspective, such as shape, size, motion, etcetera, or properties that can be analyzed wholly in terms of such properties. Physical properties include, or are analyzable in terms of, the sorts of properties physicists *qua* physicists study. By contrast, candidate examples of *non*-physical properties include properties of mathematical entities, such as *being divisible by two* or *being congruent*.

Perhaps I should be more modest: by 'physical property,' I intend to mean *whatever it is* that "type-identity" theorists mean when they use the term. I am assuming (for the sake of argument, at least) that "type-identity" theory is itself intelligible. I should add that we can give counting arguments with respect to restricted classes of "physical" properties, such as shape, mass, size, etcetera, as I will explain in a later section.

[2] Recent discussions and defenses of the identity theory include, for example, Perry 2001, Polger 2006 and Block 2009.

2. The Counting Argument

The opening outline of the argument is as follows:

1. For any class (or plurality[3]) of physical properties, the *p*s, there is a mental property of *thinking that the ps are physical*.
2. There are more classes of physical properties than there are physical properties.
3. Therefore, there are more mental properties than physical properties. (1, 2)
4. Therefore, not every mental property is a physical property.

Let us have a closer look at the premises. Premise 1 asserts that there is a mental property for any plurality of physical properties. Take *shapes*, for example. Given any class of shapes, we can coherently describe a unique mental property in terms of those shapes. There is a procedure for doing so: for any class of shapes, the *q*s, there is a distinct mental property of *thinking that the qs are shapes*. (There are other procedures, too.) Notice that there is no procedure to construct a new *shape* from any arbitrary class of shapes: for example, a shape constructed from *all* shapes would be partly constructed out of *itself*, which does not seem to be coherent. (Or if self-construction *is* coherent, then let S be a shape that is constructed from just those shapes that are *not* constructed from themselves: S is constructed from itself just if it isn't, and so S is not coherent.)

We can likewise identify mental properties in terms of various other physical properties, including masses, velocities, changes in shape, spins, and so on. We can do this using the same procedure, which allows us to identify a unique type of thought for each class of physical

[3] The argument does not require that there be such things as sets or classes. Readers are welcome to translate "class" talk into plural reference talk.

properties.[4] To be clear, when I say that the mental properties are unique, I mean they are distinct from *each other*, since they feature different pluralities of physical properties.

Alternatively, we may build complex thoughts from arbitrary "lower-order" thoughts about physical states. To illustrate, take any physical property *p*, and let mental property *m* be the property of *being a thought about p*. I'll call *m* "logically simple" because *m* isn't a logical construction (conjunction, disjunction, etc.) of other mental properties (even if *p* is). Then for any plurality of logically simple mental properties, there is a logical construction of them: for example, there is their *conjunction*. In this way, we identify a unique mental property for any plurality of physical properties. (We will consider possible objections in a moment.[5])

Premise 2 is Cantor's theorem applied to classes (or pluralities): for any particular items, the classes (taken as pluralities) of those items outnumber the items themselves. So, the classes (or pluralities) of physical properties outnumber the physical properties themselves. That sounds reasonable. Plus, there's a Cantorian argument for it, which I'll relegate to a footnote.[6]

---

[4] To be clear, when I say that the mental properties are distinct, I do not mean they are each distinct from a physical property. Rather, I mean they are distinct from *each other*, since they feature different pluralities of physical properties.

[5] In particular, we'll consider attempts to show that premise 1 overly generalizes. I will explain why those attempts each include a premise that is importantly different from premise 1.

[6] What follows is Cantor's proof of Cantor's theorem applied to the class (or plurality) of physical properties. Suppose there are no more classes of physical properties than physical properties. Then for each class C of physical properties, there is a distinct physical property, which we may call 'C's partner'. Now consider that some properties may be in the classes they are partnered with; for example, {p1, p2} might be partnered with p1. Nothing rules that out. But not every class can be partnered with one of its own members. For if every class were partnered with one of its members, then every *singleton* class (or "singleton plurality") would be partnered with its member, leaving no other properties for the other, non-singleton classes to be partnered with. So, we cannot suppose that every class is partnered with one of its members. It follows, then, that there are at least some classes that are partnered with a property that is *not* one of its members.

We have just seen that some properties cannot be members of the classes they are partnered with. Now let 'C' be the class of just those properties that are not members of the classes they are partnered with. The existence of C follows from the Axiom of Separation, if the physical properties are containable in a *set*. And even if they are not containable in a set, we may use what Pruss & Rasmussen (Forthcoming) call plural comprehension: for any formula ψ, if there are some things that satisfy ψ, then there are the things that satisfy ψ (where, in this case, ψ = '*x* is not a member of its partner'). Next, let P_C be a property that is partnered with C. There is such a property as P_C since every class is partnered with a distinct property, given our starting assumption. A contradiction is now two steps away. Step one: P_C cannot be a member of C, because C, by definition, only contains properties that are *not* in

From premises 1 and 2, it follows that there are more mental properties than physical properties, assuming that *being about the xs* is distinct from *being about the ys* when the *x*s and *y*s don't comprise the same plurality. From here, we may infer that some mental properties are not physical.

Perhaps we can go further. Consider that the divide between *physical* properties and *non-physical* properties may seem to be far greater than the divide between any two mental properties. Mental properties fall together under the same basic category. So, if some are non-physical, it's perhaps plausible that they all are.

Even if we don't take this last step, we still have the conclusion that not every mental property is a physical property. (We also have the conclusion that the determinable property, *having a thought*, can have non-physical determinates, which implies that *having a thought* is not itself wholly analyzable in terms of physical properties.) This is a significant result because it implies that there is more to mentality than is captured by a complete physical description of brains and bodies. The upshot: the *type-identity* theory carries a cardinal cost. (The argument does not target every so-called "type-identity" theory but just those that analyze mental properties as third-person physical properties.[7] The argument also leaves open non-reductive physicalist alternatives.)

---

the classes they are partnered with. Step two: P_C *must* be a member of C, because if P_C is not a member of C, then since C contains all those properties that are not in the classes they are partnered with, it follows that C contains P_C, which contradicts the antecedent. So, P_C cannot be in C, and P_C must be in C, which is a contradiction. To avoid the contradiction, we must deny the starting assumption—that are no more classes of physical properties than physical properties. (The "plurality" version of this proof replaces talk of classes with talk of pluralities, and the pairing is understood as a plurality of pairs.)

[7] The conclusion of the counting argument is compatible with Searle's proposal (2004) that mental properties, though *irreducibly* subjective in nature, are "physical" by virtue of being *causally* reducible to basic physical properties even if they aren't *ontologically* reducible to such properties. But it is not compatible with typical formulations of type identity theory, such as the theory that Polger (2004) defends.

I should emphasize that the Cantorian counting argument also says nothing about whether mental properties are physically “realizeable.” You might think, for example, that mental properties are second-order properties that have only physical realizers in our world.

On the other hand, the argument does challenge a certain metaphysical interpretation of the psycho-physical correlations we witness in neuroscientific research. Suppose the counting argument is sound. Then the mental property, *being a thought*, has non-physical determinates, because there are more ways to be a thought than there are physical properties. By contrast, all determinates of (ways to be) a neurophysical kind are, it would seem, neurophysical. If that’s right, then *being a thought* isn’t identical to a neurophysical kind. In other words, some mental kinds, including actually instantiated ones, are *not* neurophysical (even if they have a neuro-physical base or ground). Thus if the counting argument is sound, we have greater reason to be cautious when moving from *a posteriori correlation* to *identity*.[8]

3. Objections and Replies

The argument is disarmingly simple. I believe its strength becomes apparent, however, when we test it against objections. What follows are what I take to be the most serious or instructive objections I have encountered.

[8] That isn’t to say that the counting argument rules out all psycho-physical identity theories. Some identity theorists may be happy to suppose that certain psychological properties are reducible to physical properties without supposing that *all* mental properties are reducible. For example, one might think that each property of the form *being a thought that P* has two components: (i) a psychological component, and (ii) a content component, where the content component is not itself a psychological in nature. Then identity theorists who reduce *psychological* properties may grant that non-psychological components of mental properties are themselves non-physical.

Those identity theorists may still appreciate the counting argument because the argument supplies a new reason to think that thoughts have a non-physical content component—and thus that the physical facts of reality do not exhaust the actual mental facts. I am grateful to an anonymous referee for emphasizing those identity theories that are unchallenged by my argument.

Objection 1: The argument counts *uninstantiated* properties, such as being a thought about each and every type of particle in the Andromeda Galaxy. But many type-identity theorists deny the Platonist doctrine that there are uninstantiated properties. Therefore, the argument will be unappealing to many philosophers in the intended audience.

Reply: Note first that the thesis that Platonism causes trouble for the type-identity is itself a substantial (and surprising) philosophical thesis. After all, not every type-identity theorist rejects Platonism. So, I think the counting argument would still be valuable even if Platonism is presupposed.

But more importantly, the counting argument doesn't require Platonism. Even if there are no uninstantiated properties, counting arguments of this sort are still instructive. Consider, for example, a counting argument against the thesis that shapes are integers. That argument will make use of the following premise: for every integer, a distinct shape can be defined, but not *vice versa*.[9] Now suppose there happens to be no object with more than one thousand sides. Would the counting argument then be defeated? Surely not. We could still argue that shapes aren't integers by comparing the number of *definable* shapes with definable integers. The counting argument still tells us what sort of thing a non-actual shape *would* be like if it *were* to actually exist (that is, if its definition were to actually pick out something). For example, we can estimate what a thousand-sided shape would be like even if there is no such a shape. The suggestion, here, is that our investigation of the nature of shapes need not depend upon which shapes happen to actually exist.

[9] Distinct shapes are constructed from lines and angles, and the number of lines and angles is the same as the number of points on a continuum, which is $\aleph_1$ (assuming the Continuum Hypothesis). The cardinality of the set of integers, by contrast, is only $\aleph_0$.

Consider the famous diagonal argument used to show that there are more decimal numbers than integers. That argument gives us insight into decimals vis-à-viz integers, *even if* the only "existing" numbers are ones that happen to be instantiated by pluralities of things in the world (or if mathematical language is merely fictional or definitional). Whether or not there can actually *be* uninstantiated numbers is beside the point. Similarly, a counting argument may give us insight into mental properties vis-à-viz physical properties, *even if* the only "existing" properties are ones that happen to be instantiated. I suggest, then, that this first objection fails to target the heart of the counting argument.

Objection 2: Some types of thoughts will be too complicated to be physically realizable. Take, for example, a thought about each and every possible physical property. That type of thought cannot be physically realized, it might be argued, because it would be too complicated to fit within a space of any size. Therefore, it is a mistake to infer that there are more mental properties than physical properties.

Reply: Fortunately, counting arguments need not take sides on whether the counted items can be physically realized. Recall the counting argument concerning shapes and integers. That argument does not require that all definable shapes are physically realizable. Perhaps some shapes cannot be physically realized. The conclusion that there are more definable shapes than integers still holds up. Similarly, even if some mental properties cannot be physically realized, the conclusion that there are more (definable) mental properties than (definable) physical properties still holds up. What matters for the argument is not that there can be some brain or body that can instantiate the properties in question; what matters is merely that such properties are intelligible.

A reason to think that the mental properties on the table are intelligible is that we have a systematic way to coherently define them: for any physical properties, the *p*s, let *M* be the mental property of thinking about each of those *p*s. Moreover, mental properties of the form *thinking about the ps* constitute a categorically unified class of properties. That is to say, there is no non-arbitrary division among properties of that form, with seemingly coherent properties on one side and seemingly incoherent ones on the other. Such properties seem to differ merely in their complexity and in the physical properties they are about, but those differences do not seem to affect intelligibility. So, even if certain mental properties cannot be physically realized, no premise in the counting argument is called into question. Definability and realizability are different matters.

Objection 3: The following parallel argument shows that there are more *physical* properties than physical properties:

1. For any class (or plurality) of physical properties, the *p*s, there is a distinct physical property that is identical to the conjunction of the *p*s.
2. There are more classes of physical properties than there are physical properties.
3. Therefore, there are more physical properties than physical properties. (1, 2)

This parallel argument is clearly unsound, for its conclusion is contradictory. But if the parallel argument is unsound, then so is the original counting argument.

Reply: I reply that premise 1 of the parallel argument has the following defect: it strictly entails a contradiction. Here is how. Premise 1 entails that there is a conjunction for any plurality of physical properties. Therefore, there is a conjunction of the plurality of *all* physical properties. That conjunction would be a conjunction of itself, since it is itself a physical property (assuming premise 1). Now let the *r*s be just those physical properties that are *not* conjunctions of themselves. Premise 1 entails that there is the conjunction R of the *r*s. But a classical problem arises: R is a conjunction of itself just if it is not. So, premise 1 lands us in a Russellian contradiction.[10]

The same cannot be said about the original counting argument: no contradiction results from supposing that for any physical properties, there is a unique type of *thought* about those properties. Moreover, the original counting argument is backed up by procedures for coherently describing mental properties in terms of physical properties. We cannot give a similar procedure for coherently describing physical properties in terms of arbitrary physical properties. We have seen, for instance, that we cannot coherently describe a physical property as a conjunction of just those physical properties that are not conjunctions of themselves. The parallel argument, then, unlike the original counting argument, is not supported by a procedure for coherently describing properties. The arguments are importantly different. (Of course, premise 1 of the original argument *together with type-identity theory* may entail a contradiction. But it would be question-begging to object to premise 1 just *on that basis*. I will say more about the issue of "question-begging" when I consider Objection 6.)

Upon further reflection, the objection also delivers an interesting lesson about the limits of recombining physical properties. We just saw that there cannot be a physical *conjunction* for every plurality of physical properties (else the Russellian contradiction). By the same reasoning, we find that there is no way of packaging physical properties so that each plurality is packable into a distinct physical property.[11] To see why, suppose there is a package, like "arrangement" or

[10] A further difficulty is that not all conjunctions of distinct pluralities are themselves distinct. For example, the conjunction of p_1 with p_1 & p_2 is plausibly not distinct from the conjunction of p_1 with p_2.

[11] For more on the problems with packaging pluralities, see Pruss & Rasmussen Forthcoming.

"interactions among" or "realized by" or "is a function of," such that for any physical properties, the *p*s, those *p*s are packaged into a distinct physical property.[12] For example, suppose that for any physical *p*s, there is a physical *q* identical to an arrangement of the *p*s. Now either there are packages of *packages*, or there are not. If there are no packages of packages, then it's not the case that there is a package for *every* plurality of physical properties, since the packages are themselves physical properties (per hypothesis). In that case, there is this limit on physical properties: not every plurality of them has a package. If we suppose instead that every plurality has a package, then we fall into the Russellian contradiction: the package of all packages includes itself, whereas the package of all non-self-including packages includes itself if and only if it doesn't. So, there are limits—bounded by a contradiction—to the packaging of physical properties.[13]

Fortunately, we may free ourselves from those same limits if we allow packages to fall under a different category. For example, no contradiction arises from supposing that every plurality of physical properties is packable into a *set* or *class* or … *thought*. There are ways, then, to package every plurality of physical properties. In each case, the packages have a nature, whether abstract or mental, that is irreducible to the nature of the things they package. That's just what the counting argument implies.[14]

---

[12] I'm expressing the argument in terms of "physical properties," but we could equivalently give the argument in terms of "physical state types."

[13] We can make sense of the limitations in terms of *repeats*: for example, if *P* is a package of *p1* and *p2*, it's unclear how to construct a distinct, coherent physical package consisting of *P*, p_1, and p_2. See note 8.

By contrast, packages formed by the "thinking about" operator don't plausibly produce repeats: for example, *thinking about p1 and p2* is plausibly distinct from *thinking about p1 and p2 and thinking about p1 and p2*. (In the next objection, we'll consider what happens when the "thinking about" operator applies to all *mental* pluralities; then, we probably get repeats.)

[14] I am grateful to two anonymous referees who independently brought up the issue of packaging physical properties in ways other than the way of conjunction.

Objection 4: The following parallel argument shows that there are more *mental* properties than mental properties:

1. For any class (or plurality) of mental properties, the *m*s, there is a distinct mental property of *thinking that the ms are mental*.
2. There are more classes of mental properties than there are mental properties.
3. Therefore, there are more mental properties than mental properties. (1, 2)

This parallel argument is clearly unsound, for its conclusion is contradictory. But if the parallel argument is unsound, then so is the original counting argument.

Reply: I reply that premise 1 of this parallel argument, like the premise 1 of the previous parallel argument, strictly entails a contradiction. The contradiction is produced in a familiar way. Let the *r*s be just those mental properties that are *not* identified partly in terms of themselves. Then premise 1 entails that there is a mental property R that is identified in terms of just the *r*s.[15] But that's impossible because R must be identified in terms of itself just if it is not.

The original argument is not like that. No contradiction is deducible from any of its premises. The corresponding "premise 1" is this: for every plurality of physical properties, there is a unique mental property. But there is no way to deduce, from that premise, the existence of self-identified mental properties or the existence of a mental property that is identified in terms of all and only non-self-identified mental properties. The premises of the original counting argument may jointly entail that a version of reductive physicalism is false, but there is no contradiction in *that*.

Moreover, there is no *procedure* for coherently identifying mental properties in terms of arbitrary pluralities of mental properties. Recall, for example, the *r*s—the mental properties that

[15] Assuming plural comprehension—see note 6.

are not self-identified. We have just seen that any property identified in terms of the *r*s is self-identified just if it is not. Thus, it seems that no property can be coherently identified in terms of all and only the *r*s. By contrast, recall that there is a procedure for coherently identifying a mental property in terms of any arbitrary class of shapes, masses, velocities, spins, and other physical properties without generating any evident incoherencies. An important difference, then, between "workable" counting arguments and "unworkable" ones is the presence of a procedure for identifying coherent properties. (I will say more about workable procedures in my response to the next objection.)

It is worth noting that there are indeed "workable" ways to generalize the counting argument. The *most* generalized counting argument would be one that implies a conclusion of the following form: mental properties are not all *G*, for any *G* meeting the conditions necessary for the existence of a coherent procedure for defining distinct mental properties in terms of classes of *G*s. I consider it an open question what properties, in general, might meet those conditions. For all I know, it might be that any property that isn't strictly or conceptually equivalent to *being mental* meets the conditions. In that case, the most generalized counting argument would imply that there is no complete, substantial 'reductive' analysis of the mental in terms of properties falling under a type (such as *being physical*) that is not itself strictly or conceptually equivalent to *being mental*. That is a significant result because it implies that mental properties (at least certain of them) are irreducible to anything; they are, one might say, just what they seem to be in the mind's eye.

This last point reveals that the counting argument is more exactly about *mental reduction* than about *physicalism*. The argument in its most general form supplies a reason to treat mental properties (certain of them) as irreducible to any class of properties, where for any subclass, there is at least one, distinct corresponding mental property. To be clear, some reductions may be unproblematic. For example, it may be that a mental property of the form *being a thought that* $\Phi \vee \psi$ is reducible to a property of the form *being a thought that* $\sim(\sim\Phi \wedge \sim\psi)$. In general, reducing mental properties is okay when the base properties have a sufficiently rich logical form or structure. That's because the logical structure places the Russellian boundary on arbitrary packaging of base properties into a property of the same kind. If, on the other hand, the base properties lack the requisite logical form or structure, *then* the reduction is a problem—per the counting argument. This result fits nicely with the non-reductive theory that mental properties, including their basic logical structure, are irreducible.

Objection 5: The counting argument relies upon intuitions concerning what propositions can be *about*. But those intuitions lead to intolerable paradoxes and so cannot be trusted. For instance, it seems we can think about just those thoughts that *aren't about themselves*: we can think, for example, that the thoughts that aren't about themselves are about something else. However, such a thought is about *itself* just if it is not. So, there cannot be any such thought.[16] And if *that* thought cannot exist, then how can we be sure that there can be thoughts about any arbitrary physical properties?

Reply: Notice, first, that the above paradox requires an inference from

(i) There exists the thought that the thoughts that aren't about themselves are about something else.

to

[16] Cf. Grim and Plantinga 1993.

(ii) There exists a thought that is about all and *only* those thoughts that are not about themselves.

Consider that this inference cannot be valid if either of the following hypotheses are true: (a) 'the thoughts that aren't about themselves are about something else' is about *all* thoughts, because it expands to 'for *every thought* T, if T is not about itself, then T is about something else'; or (b) 'the thoughts that aren't about themselves are about something else' is about the property *being a thought that isn't about itself*.[17] On either option, the inference fails. Moreover, the inference must fail if there can be universal thoughts, such as the thought expressed by (i). Surely, it's plausible that there can be universal thoughts. So, I doubt (ii) is deducible from (i).

Second, even if other paradoxes can be found in the neighborhood, there is a principled way to protect ourselves from paradoxical definitions of the sort we have encountered. The paradoxical definitions are produced by procedures that have this in common: they attempt to define a thought T when there is no way to individually *specify* each of the things T is supposed to be about. For instance, if we try to define a thought that's specifically about each and every individual thing, including my shoelaces and your mom, then there is no way to specify each of the things this thought should be about without *already* defining the thought itself—since the thought itself is one of the very things it is about. The same problem arises for defining a thought that's about *just* those thoughts that are not about themselves. We cannot individually specify each of the things this thought is supposed to be about without *first* defining the thought itself. There is a circularity problem here, which leads to a problem of specification. I believe the problem of specification reveals a principled way to separate potentially good counting arguments from the bad, paradoxical-generating ones: the bad ones face the problem of specification, whereas the good ones do not.

[17] Cf. Rasmussen 2014, p. 167.

Consider, by contrast, the diagonal argument concerning integers and decimals. There is no hint of paradox when it comes to defining the integers and the decimals because there is a procedure that specifies the integers and decimals without circularity. The same is so for the counting argument applied to shapes. Shapes are independently specifiable (definable): there are well-defined mathematical principles for specifying shapes and classes of them. Therefore, there is no hint of paradox when it comes to defining mental properties in terms of thinking about shapes. Same goes for masses, velocities, spins, and other physical properties, it seems: each is individually specifiable. So, the original counting argument, just like Cantor's diagonal argument, is miles away from the unsound paradoxical counting arguments.

Objection 6. The argument is question-begging. Its first premise—the premise that for any class of physical properties, there is a corresponding mental property—cannot be true unless type-identity theory is false. For suppose type-identity is true. Then there cannot be a unique mental property for each class of physical properties because mental properties just *are* physical properties, and there cannot be a unique *physical* property for each class of physical properties (by Cantor's theorem). So premise 1 cannot be true unless some mental properties are not physical, which is the very thesis at issue.

Reply: The question of when an argument is "question-begging" can be tricky. Consider, first, another argument that contains a premise that is true only if its conclusion is true:

(1) All rectangles have four sides.

(2) A square is rectangle.

(3) Therefore, a square has four sides.

In this case, each premise is only true if the conclusion is true because the premises and the conclusion alike are necessarily true (if they are true at all). Yet, presumably we would not consider this argument to be question-begging. The argument seems good.

On the other hand, if it is fairly easy to see that a particular premise is true only if a certain conclusion is true, then someone who doubts the conclusion might, on that basis, resist the premise. And even if the connection is not so easy to see, once someone sees the connection, one might become less sure of the premise.

Recall the *parallel* argument about shapes and integers. That argument purports to show that there are more (definable) shapes than (definable) integers, and therefore, not every shape is an integer. Suppose someone is unconvinced by that argument on the grounds that the first premise—the premise that for every set of integers, one can construct a distinct shape—contradicts the thesis that every shape is an integer. Skepticism here might be motivated by a sufficiently good reason to believe that every shape is indeed an integer.

But clearly not everyone must be skeptical. There is a considerable cost of denying the first premise in the counting argument concerning shapes and integers. The cost arises from the fact that we lack a principled (non-question-begging) way to say *which* of the seemingly definable shapes count as coherent shapes and which ones do not. The differences between the definitions of shapes don't *seem* to account for a difference in coherence.

A similar cost seems to arise if we deny premise 1 of the counting argument concerning physical properties. Suppose we deny premise 1—the premise that for every class of physical properties, there is the mental property of thinking about that class. Then we do not have a principled (non-question-begging) way to say *which* seemingly definable mental properties count as coherent (intelligible) properties and which ones do not. This result is especially puzzling because we have a procedure for defining seemingly coherent mental properties in terms of any given physical properties.

Of course, *if* type identity theory is true, then the general procedure for defining mental properties in terms of physical properties must fail, somehow, somewhere. The problem, though, is that we have no principled account of why it should fail where it does. The procedure seems to generate perfectly coherent definitions: no particular mental property of the form *being a thought about the ps*, where the *p*s are all well-defined shapes, masses, velocities, or any other definable physical properties, can be shown to entail a contradiction, as far as I know, even *assuming* type-identity theory. (That is, no contradiction is deducible from any particular case.) So, there is a cost here to pay if the principle fails.

On the other hand, one might think that physicalists are already committed to explaining away various intuitions about the mental. Furthermore, a type-identity theorist may think the arguments for the identity theory provide a good reason to explain away any intuitive appeal of the counting argument. So, for example, although *being a thought about the ps* might *appear* coherent for any physical *p*s, the arguments for type-identity theory imply that appearances here are misleading, just as zombies might appear to be possible, even though they are not. Nevertheless, unless the arguments for type-identity theory are infallible, any new costs discovered by, say, the counting argument could, in principle, tip the scales. The counting argument serves to offer identity theorists a reason to *reassess* the scales, since the argument brings to light a new cost to consider.

I would like to add a final, concessive comment. Suppose you reject premise 1 of the counting argument. Still, you may accept the premises of a more restricted counting argument. So, for example, you might focus instead on just *shapes* to argue that some mental properties are not a shape. Or you could include masses, sizes, changes in shape, and so on, to target particular *kinds* of type-identity theories. I bring this up for the sake of those who either are unclear about what the broad term 'physical property' should mean or who might otherwise resist the unrestricted version of premise 1. Such a person might find certain restricted versions of 1 more plausible. In this way, even a committed type-identity theorist could make use of a counting argument to support certain theses about the nature of mental properties. Counting arguments can provide a useful strategy for probing the nature of mentality on a variety of views.

I'll stop here. Further investigation may reap additional objections (such is the nature of philosophy), but I believe I have said enough to get this new "counting" argument strategy on the table for discussion.

4. Wrap Up

I have argued that the number of mental properties outstrips the number of physical properties. Central to the argument is the observation that thoughts have the feature of *aboutness*; it is this feature that allows us to construct (identify) any number of mental properties in terms of arbitrary (and identifiable) classes of properties. People often report a pre-philosophical intuition that aboutness is not anything like the physical properties studied by the hard sciences. But why trust this intuition? The counting argument may add teeth to the intuition by backing it up with certain principles of counting.

I will close by emphasizing a certain advantage of the counting argument over previous arguments against the type-identity theory. Previous arguments, such as the Multiple Realizability Argument,[18] the Zombie Argument,[19] the Knowledge Argument,[20] and the

[18] Putnam 1967.

[19] Kirk 2005.

Replacement Argument,[21] depend upon intuitions about what situations are metaphysically possible. The counting argument, by contrast, does not require us to figure out whether certain exotic situations are possible. Instead, the argument makes use of a mathematical counting principle together with a topic-neutral procedure for defining certain kinds of mental properties.

Of course, there are reasons to like type-identity theory—such as that it provides what is surely the simplest explanation of psycho-physical correlations. My goal has only been to bring to the table a certain theoretical cost of the theory. I leave it to readers to decide if the price is right.[22]

---

[20] Jackson 1986.

[21] Plantinga 2006.

[22] I owe thanks to many, more than I remember. In this case, I will simply thank my wife, Rachel, who bore the burden of reviewing all previous drafts and listening to me talk about nearly all the objections and comments I've received from others.

Works Cited

Block, N. 2009. “Comparing the Major Theories of Consciousness,” in Michael Gazzaniga (ed.) *The Cognitive Neurosciences IV*. MIT Press

Grim, P., and Plantinga, A., 1993, “Truth, Omniscience, and Cantorian Arguments,” *Philosophical Studies* 71: 267–306.

Halmos, P., 1960. *Naive set theory*. Princeton, NJ: D. Van Nostrand Company.

Jackson, F., 1986. “What Mary Didn't Know,” *Journal of Philosophy* 83: 291–295.

Jech, T. 2003. *Set Theory*. Berlin, New York: Springer Monographs in Mathematics.

Kirk, R., 2005. *Zombies and Consciousness,* Oxford: Clarendon Press.

Perry, J. 2003. *Knowledge, Possibility, and Consciousness*. MIT Press.

Plantinga, A. 2006. “Against Materialism,” *Faith and Philosophy* 23: 3-32.

Polger, T. 2004. *Natural Minds*. MIT Press.

Pruss. A. & Rasmussen, J. Forthcoming. “Problems with Plurals,” *Oxford Studies in Metaphysics*.

Putnam, H. 1967. “Psychological Predicates.” in W.H. Capitan and D.D. Merrill (eds.), *Art, Mind, and Religion*, Pittsburgh: University of Pittsburgh Press, 37-48.

Rasmussen, J. 2014. *Defending the Correspondence Theory of Truth*, Cambridge University Press

Searle, J. 2004. *Mind: a Brief Introduction*. Oxford: Oxford University Press.