Identity and Indiscernibility

Katherine Hawley

Abstract

Putative counterexamples to the Principle of Identity of Indiscernibles (PII) are notoriously inconclusive. I establish ground rules for debate in this area, offer a new response to such counterexamples for friends of the PII, but then argue that no response is entirely satisfactory. Finally, I undermine some positive arguments for PII.

1. Introduction

According to a Principle of the Identity of Indiscernibles (PII), entities which are indiscernible in some specified respect are identical. We can generate different PIIs by varying the entities concerned — concrete or abstract particulars, universals too? — varying the degree of indiscernibility — intrinsic properties, purely qualitative properties, all properties? — or varying the modal strength — metaphysical or physical necessity, mere contingency? (Swinburne 1995 and Forrest 2002 offer alternative taxonomies.) Criteria of identity have the same form, ranging over entities of a given kind (e.g. sets), and linking their identity to indiscernibility in a specific respect (co-extensionality for sets).

Putative counterexamples have been central to debate about such principles. The best-known is Max Black's world which contains 'nothing but two exactly similar spheres … [each] made of chemically pure iron … every quality and relational characteristic of the one would also be a property of the other' (Black 1952, p. 156); some see counterexamples amongst quantum particles (see French and Krause 2006 for detailed discussion). A PII is threatened if there is a situation in which objects of the relevant kind seem to be indiscernible in the relevant respect without being identical. To generate such a threat we need to establish first that there is a situation involving a certain arrangement of qualitative features, and moreover that the qualitative arrangement is due to a pair of distinct but indiscernible objects. The first of these points is established empirically in the actual-world quantum cases, and via more controversial conceivability arguments in cases like Black’s. The second step must involve the familiar, mysterious, philosophical method of comparing different explanations, weighing theoretical virtues and vices, and making tentative inferences.

The same two-step process is required in debate about the Special Composition Question. Here, some think that no pluralities have sums, others that only certain pluralities have sums, still others that every plurality has a sum (for discussion see, e.g., van Inwagen 1990). We cannot defeat the no-sums view just by pointing to a counterexample — a butterfly! — for we can agree on the qualitative arrangement of the world yet reasonably disagree about the distribution of individuals. Debates about which objects have proper parts, about whether persisting objects have temporal parts, and about the nature of properties (is this instance identical to that, or are they distinct tropes?) have the same nature. Facts about identity and parthood are not directly deducible from the agreed empirical facts.

So building a counterexample to a PII is a two-step process. We must show that a certain qualitative arrangement is actual or possible (actual, if we are to defeat modest PIIs which claim only contingent truth), and we must also show that two distinct indiscernibles are responsible for that arrangement. Characterising the debate in this abstract fashion is useful for several reasons. First, it reminds us that whether quantum physics conflicts with PII is not a purely empirical question. Second, it clarifies the dual role of conceivability in Black’s case: even if conceivability is a good guide to the possibility of qualitative facts, it may nevertheless be a poor guide to the possibility of identity facts. Third, it prompts us to think about what sorts of qualitative arrangement might be best explained by PII-defying identity facts.

As I shall use the phrase, the qualitative arrangement consists of those facts about the world which do not immediately settle questions about identity and parthood, and which can be accepted in common by those who disagree about persistence, composition and the like; ‘settle’ here is an imprecise epistemic notion, not a matter of metaphysical determination.

‘Qualitative’ more usually picks out those properties and relations whose instantiation does not require the existence of any specific object: thus *composing something* usually counts as qualitative, while *composing the Eiffel Tower* is non-qualitative. (This distinction is set out in Adams 1979, Sect. I.) But my usage is narrower: even *composing something* counts as non-qualitative in my sense.

2. What sort of qualitative arrangement?

In this section I discuss what sort of qualitative arrangement is required for a putative counterexample to a PII, what we need in order to take first of the two steps towards building a counterexample. This project is of interest in its own right; moreover it is significant because Michael Della Rocca has recently attempted to use the fact that certain qualitative arrangements obviously fail to establish a counterexample to PII as leverage against all possible counterexamples (Della Rocca 2005).

Suppose someone claims, bizarrely, that the desk in my study is co-located with an indiscernible desk: my desk and its duplicate do not differ even modally or sortally. To make this story fit the empirical facts, we must suppose that the duplication is qualitatively insignificant: the two desks do not collectively have double the mass of an ordinary wooden desk. After all, the supposed two desks are made of the same molecules at the same time: this explains both their indiscernibility, and the fact that together they have only as much mass as either would alone, though this explanation would not account for allegedly undetectable duplicate simples. (Wasserman (2002) develops a line like this on behalf of those who suppose, somewhat less bizarrely, that my desk is co-located with a distinct lump of wood.) Now, if my desk really does have a co-located exact duplicate, then we have a counterexample to PII. But there is, of course, no reason to think that my desk has such a duplicate.

Why is there no reason to think that my desk has a duplicate? According to Della Rocca (2005), this is a puzzle confronting anyone who rejects PII. If PII is true, then pairs of distinct, indiscernible objects are impossible. But, asks Della Rocca, if you reject PII, and thus accept the possibility of distinct indiscernibles, what grounds do you have for rejecting scenarios like the duplicate desks? Why think that the qualitative arrangement in my office is best explained by the existence of a single desk, rather than by the existence of two (or two thousand) indiscernible desks?

One possible response is to concede that distinct objects must be discernible in some respect, even if only in their spatial relations to other objects. This concession is compatible with maintaining that distinct objects can be intrinsically indiscernible; here, we are reminded of the vast variety of PIIs.

But Della Rocca's argument does not force this concession. Suppose that distinct objects may be entirely indiscernible, even co-located. Then it is indeed metaphysically possible that there be two or twenty 50kg desks occupying a single desk-shaped location in my study, compatibly with the actual qualitative arrangement of my study. But the fact that such qualitatively-undetectable duplicates are metaphysically possible does not mean that we should believe that they are actual, nor even that we should suspend belief about their existence. Ordinary methodological principles direct us towards quantitative parsimony: other things being equal, posit as few objects as are necessary to explain the phenomena.

Quantitative parsimony rationalises, e.g., scientists' decision to account for the missing mass in beta decay by claiming that a single particle — a neutrino — is emitted in each decay. *Qualitative* parsimony is neutral between this claim and the hypothesis that twenty-seven even smaller particles are emitted during each decay: each hypothesis posits the existence of one new kind of entity. But quantitative parsimony counts in favour of the single-neutrino theory (Nolan 1997, Baker 2003). It is metaphysically possible that there be particles emitted in groups of twenty-seven during beta decay, with the same results as we actually attribute to the single neutrino. Nevertheless, it is reasonable to suppose that this is not what actually happens. Similarly, we can acknowledge the metaphysical possibility that my desk has a co-located duplicate without thereby acquiring reason to doubt that there is only one desk in my study; conversely, my legitimate confidence that I have only one desk does not entail that such duplication is metaphysically impossible.

Too easy? Della Rocca considers the related suggestion that the ‘one desk’ hypothesis is simpler than the ‘two desks’ hypothesis. He argues that simplicity directs us to identify the ‘two’ spheres in Black’s case: once we admit that simplicity (or quantitative parsimony) is a guide to reasonable belief, we will always be prompted to identify indiscernible objects, and thus to endorse PII.

But quantitative parsimony must be considered alongside other types of economy, and simplicity, explanatory power and so on. Indeed, unchecked quantitative parsimony would lead us straight to ontological monism — the view that there is but a single object — and while admiration for quantitative parsimony is widespread, ontological monism is not. So the principle needs to be applied with care. In cases like the duplicate desks, there is no countervailing pressure to violate quantitative parsimony and thus violate PII. But Della Rocca’s attempt to generalise this result fails, because other cases are different.

Crucially, the desk case involves qualitatively insignificant duplication: it is part of the story that whether there are one, two or twenty such desks, the qualitative arrangement in my office is constant. Putative counterexamples to PII must involve qualitatively *significant* duplication, so that it makes a qualitative difference how many of the supposedly indiscernible objects there are. That is, the qualitative arrangement must be one for which the existence of multiple objects is a promising explanation.

The quantum ‘counterexamples’ involve qualitatively significant duplication: in such cases, apparently indiscernible particles are mereologically disjoint, and have additive properties like mass and charge, so it makes an empirically detectable difference how many such particles there are in a system. As Black describes it, his universe also involves qualitatively significant duplication: the sum of the ‘two’ spheres has twice the mass and twice the spatial extension of either individual sphere, and moreover it is not itself spherical.

Cast in these terms, much of the debate about Black's case has focused on whether he has genuinely established the possibility of a world with qualitatively significant duplication of this sort. Ian Hacking writes

> … there is no possible universe that must be described in a manner incompatible with [PII]. Yet this fact is no proof of [PII], for some imagined universes may be described in a way that violates the principle. In short, it is vain to contemplate possible spatiotemporal worlds to refute or establish the identity of indiscernibles. (Hacking 1975, p. 249)

Most readers interpret Hacking as highlighting the empirical equivalence of two indiscernible spheres in a Euclidean spacetime, on the one hand, with a single sphere in a cylindrical spacetime on the other (e.g. Landini and Foster 1991, French 1995, Forrest 2002; Adams 1979, Sect. III has a detailed discussion). Hacking's point may be that there is no genuine difference between these 'two' possibilities, or perhaps that evidence for the possibility of the Euclidean-two-sphere world can be explained away by the genuine possibility of the one-sphere-in-a-cylindrical-spacetime world. Hacking's critique raises important issues, but from here on I will assume that Black has established the possibility of a world with the qualitative arrangement one would expect of two indiscernible spheres two miles apart in a Euclidean spacetime; I will focus on the gap between establishing this and establishing a counterexample to PII.

3. How to defend PII?

As the literature shows, even those putative counterexamples to PII which involve a suitable qualitative arrangement are not universally persuasive. There are strategies which accept the problematic qualitative arrangement broadly at face value, but attempt to reconcile it with PII.

The first such strategy is to identify the 'two' objects (I will call this the 'identity defence' of PII). The second is to find a relevant respect in which the objects are, after all, discernible (the ‘discerning defence’). But there is a neglected third option: accept the qualitative description of the case, but deny that either of the alleged indiscernibles exist. Instead, posit a single object, with the qualitative features which would have been possessed by the sum of the two indiscernibles, had they both existed and had a sum. In Black’s case, this would involve claiming that his universe contains only a simple, partless object, which extends through a disconnected spatial region. This object is not a sphere, and doesn’t contain any sphere as a part. I will call this third option the ‘summing defence’.

In cases where duplication is qualitatively insignificant, the summing defence is equivalent to the identity defence. The putative duplicate desks are together indiscernible from a single desk, which is why principles of quantitative parsimony count against the ‘two desks’ hypothesis. So the summing defence — posit a single object qualitatively indiscernible from a sum of the ‘two indiscernibles’ — simply collapses into the obvious identity defence — posit a single object qualitatively indiscernible from each of the ‘two indiscernibles’. Where duplication is qualitatively significant, however, the summing defence is distinctive, and, I shall argue, it is sometimes superior to both the identity defence and the discerning defence.

3.1 The identity defence

First consider the identity defence, as a response to cases which apparently involve qualitatively significant duplication. *Prima facie*, it might seem hopeless to claim that although a case is qualitatively equivalent to one containing two Fs (e.g. two spheres), it in fact contains only a single F. Nevertheless, John O’Leary-Hawthorne offers such a strategy to those who think that concrete particulars are bundles of universals (O’Leary-Hawthorne 1995; for more recent discussion, see Hawthorne and Sider 2002). Bundle theorists may take Black’s universe to contain a single bi-located sphere, one which is located both two miles from itself and in the same place as itself.

As O’Leary-Hawthorne points out, if particulars are entirely composed of universals, then it is plausible that they, like universals, are multiply-locatable. But oddities emerge as we consider the mereology of this situation. Suppose the sphere has mass 1kg, and volume $1m^3$, which is to say that the corresponding universals are elements of the multiply-located bundle. Now, what should we say about the sum of the 'two' spheres?

On the one hand, since the 'two' spheres are one and the same object, presumably the sum of those 'two' spheres is simply the sphere itself, which thus has mass 1kg and volume $1m^3$, and, indeed, a spherical shape. On the other hand, we might expect the sum of this 1kg sphere and that 1kg sphere to have mass 2kg and occupy a $2m^3$ non-spherical scattered region; at least, it looks as if something in that universe has those features. The bundle theorist might try denying that the 'two' spheres have a sum. After all, they are spatially separated, and apparently don't have much causal interaction. But the sum of *o* plus *o* is guaranteed to exist if *o* does, for it is just *o* itself.

Perhaps the single sphere includes *having volume exactly* $2m^3$ in its bundle, as well as *having volume exactly* $1m^3$. After all, it occupies two distinct $1m^3$ regions simultaneously. And maybe this is no worse than including *being located at some distance from itself* as well as *being co-located with itself*. But then the bundle also includes both *being spherical* and *being two-sphere-shaped*, and perhaps *having mass exactly 1kg* and *having mass exactly 2kg*. I, at least, am beginning to lose my grip on how property-exclusion is supposed to work here.

Perhaps we should relativised the single sphere's properties to various regions.[1] That is, it has mass 1kg relative to one of the spherical regions it occupies, has mass 1kg relative to the other spherical region, and has mass 2kg relative to the scattered two-sphere-shaped region. After all, we are contemplating objects which are wholly present in more than one place, and those who believe that persisting objects are wholly present at more than one time need to find a way of relativising property attributions to times, in order to account for change.

[1] I thank an anonymous referee both for this suggestion and for the observation that it does not fit the quantum cases.

There are several problems with this suggestion, given the bundle-of-universals theory of particulars. First, it drastically reduces the number of intrinsic properties in the bundle, leaving only those which the object has relative both to the spherical regions separately and to the larger region. Second, it imports machinery designed to account for variation into a case premised on lack of variation, i.e. indiscernibility. Third, as we will shortly see, it cannot apply in the quantum cases, where the putative indiscernibles are co-located.

I think the best option for the bundle theorist who wants to pursue the identity defence is simply to bite the bullet: even though Black's universe looks and would behave as if it contains two distinct spherical objects which collectively instantiate *having volume 2m*^{3} and *having mass 2kg*, this is an illusion. No object or even plurality of objects in that universe has mass 2kg. This is not a *reductio* of the identify defence. But these considerations do show that taking the defence seriously will require us to completely rethink the relationships between what properties are instantiated by the objects in a world and the way in which regions of that world appear to be filled and qualified.

Moreover, the identity defence is even less appealing when we consider co-located qualitatively significant duplicates (as certain quantum particles are alleged to be). Here, the suggestion would have to be that a single bundle is instantiated several times in the same location, with qualitatively significant consequences. It is hard to see what the force of this claim is, what it is, e.g., for *having charge +1* to be instantiated n times at the same spacetime point, unless this is for *having charge +n* to be instantiated (but by what?) or for *having charge +1* to be instantiated by each of n distinct objects (in violation of the identity defence).

3.2 The discerning defence

The second option for handling putative counterexamples to PII is the discerning defence: find some overlooked but relevant respect of discernibility. Simon Saunders has recently shown how to hunt for such respects, directing our attention to Quine's 'third grade of discriminability' — weak discernibility (Saunders 2006, Quine 1976). Two objects are weakly discernible if they stand in some symmetric but irreflexive relation. If Black's universe contains two distinct spheres, two miles apart, then they are weakly discernible, since each is two miles from the other, but not two miles from itself.

The discerning defence is not always available: not every putative counterexample is one in which the putatively distinct objects are weakly discernible. Saunders himself notes the distinction between certain entangled fermions, which are suitable for this treatment since each has spin opposite to the other, but not to itself, and certain entangled bosons, which are not even weakly discernible since each has spin the same as the others.

What are the merits of the discerning defence where it is available? I will focus upon mere weak discernibility, because stronger respects of discernibility are unlikely ever to have been overlooked. In particular, I will discuss whether a PII based on mere weak discernibility is worth defending.

A PII based on mere weak discernibility claims that, necessarily, any two distinct objects are at least weakly discernible (stand in some irreflexive relation). According to Steven French

> … there is the worry that the appeal to irreflexive relations in order to ground the individuality of the objects which bear such relations involves a circularity: in order to appeal to such relations, one has to already individuate the particles which are so related and the numerical diversity of the particles has been presupposed by the relation which hence cannot account for it. (French 2006, p. 5)

What is the alleged circularity here, and why is it a worry? As French illustrates, one significant motive for adopting a PII is the desire to ground facts of identity and distinctness in non-identity facts. Seen in this light, a bold PII claims that differences in intrinsic properties are the only grounds for distinctness, and a PII based on mere weak discernibility tells us that the fact that otherwise-indiscernible objects stand in some irreflexive relation can ground the fact that they are two distinct objects. Some may be sceptical of notions such as ‘grounding’, ‘dependence’, and the like. But in order to explore the options open to those who do want to ground identity facts, I will ignore such doubts: granted the assumption that some facts ground others, can facts about the weak discernibility of objects ground facts about their distinctness?

In Black’s scenario, each of the two spheres *a* and *b* is two miles from an iron sphere, and they differ only in that *a*, unlike *b*, is two miles from *b* and *b*, unlike *a*, is two miles from *a*. (Likewise, a suitable pair of fermions, *f* and *g*, differ only in that *f*, unlike *g*, has opposite spin to *g*, and *g*, unlike *f*, has opposite spin to *f*; there is no determinate fact of the matter about the spin direction of either considered separately.) Consider the properties in which *a* and *b* differ: *a* has *being two miles from b*, whilst *b* has *being two miles from a*. If *a* and *b* are distinct, these are evidently distinct properties: a third object might instantiate one but not the other property.

But what grounds the fact that they are distinct properties? There seem to be two relevant options. The first is that the distinction between the properties is grounded in the distinction between *a* and *b*; that is, that the monadic property *being two miles from b* depends for its identity upon the two-place relation *being two miles from* and the object *b* (and similarly for the property *being two miles from a*). The second option is that the distinction between the two monadic properties is somehow more fundamental than the distinction between *a* and *b* themselves. (See Lowe 2005 for related discussion of property identity in the context of Evans’s argument against ontically vague identity, and Lowe 1998 for discussion of identity-dependence.)

Suppose that some third object *c* instantiates *being two miles from b*, and moreover that *c* stands in the two-place *being two miles from* relation to *b*. These do not look like independent facts: which of them grounds the other? It is hard to be sure of the rules here, but a couple of considerations suggest that the instantiation of the two-place relation grounds the instantiation of the monadic property. First, the two-place relation is a better candidate for naturalness; it could feature in lawlike generalisations, and its obtaining is an intrinsic feature of the sum of its relata, whereas the obtaining of *being two miles from* *b* is not an intrinsic feature of anything. Second, the fact that *c* stands in *being two miles from* to *b* and the fact that *b* stands in *being two miles from* to *c* look as if they are both grounded in the fact that *b* and *c* collectively instantiate *being two miles apart*. In contrast, if we resist the reduction of the monadic *being two miles from b* to the combination of *being two miles from* and *b*, then we struggle to understand why the fact that *c* is two miles from *b* is conjoined with the fact that *b* is two miles from *c*. (See Dorr 2004, Fine 2000 and MacBride 2007 for more sophisticated discussion of related issues.)

I conclude, that, given the assumption that some facts ground others, the fact that *c* instantiates *being two miles from b* is grounded by the fact that *c* stands in the two-place *being two miles from* relation to *b*. So the identity of the property *being two miles from b* is grounded in the identity of the *being 2 miles from* relation, and in the identity of *b*. Similarly, the identity of the property *being 2 miles from a* is grounded in the identity of that relation, and the identity of *a*. This means that the distinctness of the two monadic properties is grounded in the distinctness of the two objects, *a* and *b*. So we cannot take the fact that *a* and *b* are discernible in respect of these properties to ground the fact that *a* and *b* are distinct.

French's concerns on behalf of those who seek to ground identity facts in facts about indiscernibility are vindicated. When two objects are weakly discernible, this fact is grounded in the fact that the objects in question are distinct; weak discernibility cannot itself be the ground of distinctness. But highlighting overlooked respects of mere weak discernibility is the only plausible option for the discerning defence, and so the discerning defence is not available to those whose principal motive for defending a PII is the desire to ground facts about the identity and distinctness of objects. So the appeal of the discerning defence, even where it is available, is limited.

3.3 The summing defence

It is useful to compare the discerning defence with the summing defence. Following the latter, we can take Black's universe to contain exactly one object, with just those properties — mass 2kg, volume $2m^3$ — we would expect from the sum of the two spheres. According to the identity defence, the property *being spherical* is instantiated in Black’s universe by exactly one thing; according to the discerning defence the property is instantiated by exactly two things; according to the summing defence, it is not instantiated at all, since the universe contains only a scattered two-sphere-shaped object. In the quantum cases, an advocate of the summing defence would deny the existence of the individual particles, and accept only the existence of the larger system.

When considering a scenario in which the discerning defence is available, and choosing between the discerning and the summing defences, our question is whether the scenario contains two objects which differ in an irreflexive relation, or whether it contains only a single, larger object. I will assume that those who adopt the discerning defence will also accept that the two distinct objects have a sum, so that our question is whether the large object has two distinct proper parts. This assumption is for ease of exposition only, since the point at issue is not the existence of the ‘sum’ but the existence of the two smaller objects.

Issues about the circumstances under which objects have proper parts are usually discussed in connection with the Doctrine of Arbitrary Undetached Parts, the Inverse Special Composition Question or the Simple Question (van Inwagen 1981, van Inwagen 1990, Markosian 1998, respectively). Two sorts of consideration seem relevant. The first is whether there are features of the larger object which are best explained by the existence and features of the smaller objects. The second is whether, considered separately, each of the smaller objects is a good candidate for existence. These considerations may overlap: after all, having ‘parts’ which are good candidates for existence is a feature of the larger object which is best explained by the existence and features of those smaller objects. (I am writing as if the existing objects are a mere sub-group of the objects; my aim is to avoid narrowing the focus to regions, trope-bundles, portions of stuff, or other object-surrogates.)

So, should we distinguish objects which are only weakly discernible, or should we follow the summing defence and posit a single, simple, larger object? First, consider whether the nature of the larger object needs explaining by that of the smaller ones. If *a* and *b* are only weakly discernible, then *a* bears some relation to *b* which *a* does not bear to *a*, even though *a* and *b* are intrinsically indiscernible. So the instantiation of this relation fails to supervene upon the intrinsic natures of the relata: the discerning defence involves commitment to nonsupervenient relations. This is a familiar point in unfamiliar guise: the relevant relation between Black's spheres is spatiotemporal, and spatiotemporal relations are of course nonsupervenient, whilst the relevant relation between fermions arises from their entanglement, and it is well-known that quantum entanglement poses a problem for Humean Supervenience (Teller 1986, Hawley 1999).

But a qualitative arrangement which would involve a nonsupervenient relation between two objects if they existed gives us little reason to posit such objects in addition to the larger object 'they' compose, i.e. little reason to opt for the discerning rather than the summing defence. Such a situation is paradigmatically one in which the properties of the 'whole' need no explanation in terms of putative parts.

What about other aspects of the larger object? In Black's scenario the two-sphere object has mass 2kg (let us suppose); in the quantum cases too the system as a whole has mass, charge and other additive properties. Is the mass property of the whole best explained by positing proper parts with mass properties of their own? Perhaps, perhaps not. But however strong your reductive urge here, it will not prompt you to differentiate cases in which the putative parts are mutually discernible from cases in which they are not. Either the properties of the whole require explanation in terms of those of the parts — in which case PII will be violated where no respect of discernibility can be uncovered — or else they do not, and we should favour the summing defence over the discerning defence even where the discerning defence is available.

Consideration of the larger object seems to favour the summing defence over the discerning defence. What about the putative parts themselves — are they (via the corresponding filled regions, trope-clusters, or portions of stuff) good candidates for existence? Here, we need to distinguish cases like Black's, in which the putative parts are not colocated, from cases like the quantum systems in which the individual particles, if they exist, have the same location properties. I will treat these in turn.

When applied to cases like Black's the summing defence of PII commits us to the existence of scattered simples. Scattered objects are not objectionable *per se*: most of us believe in these. But the problem with scattered simples is that it is hard to see what more could be required for the existence of an object than the existence of a maximally connected portion of matter; that is, it is hard to see what prevents each of the spherical regions from exactly containing an object. It is very natural to suppose that existence conditions must be intrinsic, or, if they are extrinsic, only in a negative, maximal way, that what it takes for an object to exist can only be a question of getting the local conditions right, including appropriate isolation where necessary (Sider 2001, Markosian 1998). The summing defence of PII violates this intuition: what prevents the spherical region from containing an object is the existence of another spherical matter-filled region, together with the PII.

This problem arises for both the identity defence and the summing defence of PII in the face of noncolocated indiscernibles. Suppose we have two distinct regions, each containing one of two almost-indiscernible distinct objects (compare Adams 1979). It follows both from the identity defence and from the summing defence that neither of these regions would have contained an entire object, if either had been marginally different in a way which ensured the indiscernibility of their occupants. Advocates of the summing defence might retreat to arguing that there is only a single, scattered object even in the world which appears to contain two almost-indiscernible objects, but this looks like the beginning of the road to ontological monism. On the other hand, advocates of the identity defence can hang onto the idea that whether a region contains an object is locally determined even though the entire location of that object is not locally fixed. (The interaction of multiply-locatable objects with extended simples creates many complications for our talk of locations; see Hudson 2006, Parsons 2007.)

But both the identity and the summing defences of PII in such cases push us to accept that facts about the number of distinct objects entirely located in a region do not depend upon the intrinsic features, or even the maximality-related features of that region, that they are instead determined by the features of the entire universe, and in particular by whether there is an exactly-matching region elsewhere. It looks as if the discerning strategy is the best option here, at least for those who do not care about grounding identity; those who do care must pay the high price of adopting either the identity or the summing defence.

What about the quantum cases, in which the putative parts of the system share all their location properties? Here at least the summing defence — denying the existence of the separate parts — does not commit us to the existence of scattered extended simples. It is worth comparing Saunders' treatment of certain systems of entangled fermions with his treatment of certain systems of bosons. As we have seen, in the former case he points out that the fermions are weakly discernible (though this will not satisfy those who want a PII to *ground* identity facts). In the latter case, there is no irreflexive relation between the bosons, and so Saunders argues that bosons are not real objects.

> The only cases in which the status of quantum particles as objects is seriously in question are therefore elementary bosons … we went wrong in thinking the excitation numbers of the mode [of the quantum field], because differing by integers, represented a count of things; the real things are the modes … Similar conclusions follow in the classical case … Problems only arise if relative distances and velocities are zero [i.e. if particles are not even weakly discernible], in which case, if no more refined description is available, they will remain structureless and forever combined, and we would do better to say there is only a single particle present (with proportionately greater mass). This [is] a classical counterpart to elementary bosons … (Saunders 2006, p. 60)

Whatever type of entity a mode is, the classical analogue suggests that Saunders adopts the summing defence of PII in cases where the discerning defence is unavailable, where the putative objects in question are not even weakly discernible. Yet he opts for the discerning defence where it is available.

Why discriminate in this way? Why not adopt the summing defence in both cases? We are discussing how to choose between the summing and discerning defences, in cases where both are available, and in particular we are now considering whether there is any aspect of the putative parts which presses us to recognise their existence in addition to the existence of the system as a whole. Again, there seems to be nothing which presses us to acknowledge the existence of putative parts if and only if they would be weakly discernible if they existed. When thinking about whether an object has parts, it just seems irrelevant whether, if they existed, those parts would be indiscernible one from another. We may consider the object (system) as a whole, and ask whether its properties are mereologically reducible, and we may consider the integrity of the putative parts separately. Anything else seems superfluous. (I was reaching for this point in Hawley 2006.)

4. So why PII?

Those who want to defend PII against putative counterexamples have a number of options. They may adopt the identity defence, though this is hopeless where the putative indiscernible objects are co-located, and problematic even where they are not. They may adopt the discerning defence where this is available, but this will not satisfy the urge to ground identity facts in facts about (in)discernibility, and it is unmotivated in cases where the putative objects are co-located. They may adopt the summing strategy, with some plausibility in co-location cases. Or they may resist the conceivability arguments that seek to establish scenarios like Black's.

But perhaps the most interesting question is this: why bother? That is, what motivation is there for accepting PII as a constraint either in establishing which qualitative arrangements are genuinely possible, or in inferring facts about identity and parthood from facts about the qualitative arrangement?

One reason for advocating a PII is that it is a consequence of some other metaphysical view. For example, a PII for properties follows from the identification of properties with sets, plus the co-extensionality criterion of identity for sets. More famously, a PII for concrete particulars follows from the identification of such particulars with bundles of universals. I cannot here assess all the various metaphysical theories which entail PIIs for entities of one sort or another. But in the face of putative counterexamples, we must count the cost of the various defences of PII in the balance against whatever benefits are promised by the metaphysical view in question.

A different kind of reason for advocating a PII is broadly empiricist, even anti-metaphysical. This begins from the idea that all we can know of objects, all that is physically relevant about objects, is the properties they instantiate. After all, objects' properties determine how they behave, and thus how we can interact with them. The substitution of one object for a 'distinct' yet indiscernible object would be just that — indiscernible — and could make no practical difference; facts about objects which go beyond facts about their qualitative features are inaccessible to us, and perhaps talk about such facts is just empty.

These are attractive thoughts, but only so long as we think about substituting indiscernibles one for another. It is surely true that it make no practical difference which of two indiscernible objects we are dealing with. But it may be of great practical importance whether we have one such object present or two (or two thousand). The two-object case is discernible from the one-object case, at least where duplication is qualitatively significant. As we have seen, it is possible to salvage PII in the face of such cases. But concern for empirical significance provides us with no good reason to do so.

A related thought is this: quantitative parsimony is a theoretical virtue, so it is a good methodological principle to assume that objects are identical, unless there is reason to think that they are distinct. True enough, but we should recall that discernible differences are not the only reason to think that objects are distinct; their adding up to more than either taken separately is also a reason to recognise their multiplicity.

A more thorough-going empiricism might reject any attempt to infer facts about individuation, identity and parthood from what I have been calling the ‘qualitative arrangement’. Perhaps such facts are epistemically inaccessible, or dependent upon convention, or nonexistent: one can imagine the full range of anti-realist positions. This view is a serious challenge to debate about PII as it is usually conducted. It does not, however, support PII, but invites scepticism about both PII and its negation.

So metaphysical grounds for a PII must stand or fall on their merits, relative to the costs of handling counterexamples. And moderate empiricism does not lead to PII, whilst more extreme empiricism undermines the whole debate. Finally, I consider a motive for PII which combines something of the metaphysical and the anti-metaphysical. This is the thought that the cost of rejecting PII is an unwelcome commitment to ungrounded brute facts about identity, to transcendental individuation, haecceity, or similarly dank metaphysics. (See French and Krause 2006 for detailed discussion; Della Rocca (2005) seems motivated by opposition to brute identity facts.)

What are ‘identity facts’ and in what sense are they ungrounded if PII is false? If PII is false, then among the pairs of indiscernible objects there are some pairs whose members stand in the relation of identity and some pairs whose members do not. But this distinction may have a qualitative ground. Where the members of a pair are identical as well as indiscernible, their sum has the same features as each member. Where the members of a pair are indiscernible but not identical, their sum has unmysterious features which are different from those of either member (greater mass, for example), at least where qualitatively significant duplication is concerned. And this amounts to a qualitative difference between the pairs. So the falsity of PII is compatible with the truth of the claim that whether ‘two’ objects are identical or distinct must have a qualitative ground.

One might object that the distinctness of the two objects in question must come before their having a sum which differs from either considered separately. It’s *because* they are two that they add up significantly, not vice versa. I don’t know whether this claim is true, but it seems to be on a par with the claim that the distinctness of two objects must come before their qualitative difference, that it is *because* they are two that they can differ qualitatively, not vice versa. But if this second claim is true, then advocates of PII must abandon their project of grounding ‘identity facts’ in qualitative discernibility and again we lose the motivation for PII.

However sometimes ‘identity facts’ seems to mean not facts about whether given objects are identical or distinct, but rather facts about which object is which. If *a* and *b* are distinct yet indiscernible, as the falsity of PII would permit, what fixes which is *a*, and which *b*? By assumption, no qualitative feature of either *a* or *b* is doing this work, yet *a* is *a*, and *b* is not, so we seem to have an ungrounded fact, a fact which can reasonably be called an ‘identity fact’. The objects differ in which objects they are, without differing in any other respect.

There is a genuine, difficult problem in this vicinity, one concerning modal (transworld) identity conditions. Given that *a* and *b* are actual, we might wonder what could possibly determine which object in another possible world was *a*, and which *b*. But denying PII doesn’t create extra difficulties here: even if PII is true, and *a* is discernible from all other actual objects, we can still ask which object is *a* in another possible world in which no object is an exact duplicate of the actual *a*. PII’s guarantee that every object has an intraworld-unique collection of properties would provide neat universal modal identity conditions only if objects had all their properties essentially. (Mackie 2006 is an excellent, book-length discussion of modal identity conditions; it contains only three brief references to PII.)

So modal worries do not motivate PII. Is there a nonmodal ‘which is which’ fact, a fact which goes beyond the fact that the objects are distinct, a fact which goes ungrounded if PII is false? Maybe there is, but I, for one, cannot see what it is. Whether or not identity facts are always grounded is independent of whether PII is true.[2]

---

[2] Many people have helped me with this paper. In particular I thank audiences in St Andrews and at a Rutgers conference on *Mereology, Topology and Location*, Michael Della Rocca, Steven French, James

KATHERINE HAWLEY

Departments of Philosophy

University of St Andrews

Fife, KY16 9AL, UK

kjh5@st-and.ac.uk

References

Adams, Robert Merrihew 1979: ‘Primitive Thisness and Primitive Identity’, *Journal of Philosophy*, 76, pp. 5–26 .

Black, Max 1952: ‘The Identity of Indiscernibles’, *Mind*, 61, pp. 153–64.

Baker, Alan 2003: ‘Quantitative Parsimony and Explanatory Power’, *British Journal for the Philosophy of Science*, 54, pp. 245–259.

Bermudez, José (ed.) 2005: *Thought, Reference and Experience*. Oxford: Oxford University Press.

Della Rocca, Michael 2005: ‘Two Spheres, Twenty Spheres, and the Identity of Indiscernibles’, *Pacific Philosophical Quarterly*, 86, pp. 480–92.

Dorr, Cian 2004: ‘Non-Symmetric Relations’, *Oxford Studies in Metaphysics*, 1, pp. 155–92.

Fine, Kit 2000: ‘Neutral Relations’, *The Philosophical Review*, 109, pp. 1–33.

French, Steven 1995: ‘Hacking Away at the Identity of Indiscernibles’, *Journal of Philosophy*, 92, pp. 455–66.

- 2006: ‘Identity and Individuality in Quantum Theory’, *The Stanford Encyclopedia of Philosophy* (Spring 2006 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2006/entries/qt-idind/>.

French, Steven and Krause, Decio 2006: *Identity and Individuality in Modern Physics*. Oxford: Oxford University Press.

Forrest, Peter 2002: ‘The Identity of Indiscernibles’, *The Stanford Encyclopedia of Philosophy* (Summer 2002 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2002/entries/identity-indiscernible/>.

Ladyman, two anonymous referees for *Mind*, and the Editor, Tom Baldwin. My work on this project was supported by a Philip Leverhulme Prize, which I gratefully acknowledge.

Hacking, Ian 1975: ‘The Identity of Indiscernibles’, *Journal of Philosophy*, 72, pp. 249–56.
Hawley, Katherine 1999: ‘Persistence and Non-Supervenient Relations’, *Mind*, 108, pp. 53–67.
- 2006: ‘Weak Discernibility’, *Analysis*, 66, pp. 300–3 .
Hawthorne, John and Sider, Theodore 2002: ‘Locations’, *Philosophical Topics*, 30, pp. 53–76.
Hudson, Hud 2006: *The Metaphysics of Hyperspace*. Oxford: Oxford University Press.
Landini, Gregory and Foster, Thomas 1991: ‘The Persistence of Counterexample: Re-examining the Debate over Leibniz’ Law’, *Noûs*, 25, pp. 43–61.
Lowe, E.J. 1998: *The Possibility of Metaphysics*, Oxford: Oxford University Press.
- 2005: ‘Identity, Vagueness and Modality’, in Bermudez 2005, pp. 290–310.
MacBride, Fraser 2007: ‘Neutral Relations Revisited’, *Dialectica*, 61, pp. 25-56.
Mackie, Penelope 2006: *How Things Might Have Been*. Oxford: Oxford University Press.
Markosian, Ned 1998: ‘Simples’, *The Australasian Journal of Philosophy*, 76: 213–29.
Nolan, Daniel 1997: ‘Quantitative Parsimony’, *British Journal for the Philosophy of Science*, 48, pp. 329–43.
O’Leary-Hawthorne, John 1995: ‘The Bundle Theory of Substance and the Identity of Indiscernibles’, *Analysis*, 55, pp. 191–6.
Parsons, Josh 2007: ‘Theories of Location’, *Oxford Studies in Metaphysics*, 3, pp. 201–32.
Quine, W.V.O. 1976: ‘Grades of Discriminability’, *Journal of Philosophy*, 73, pp. 113 –6.
Saunders, Simon 2006: ‘Are Quantum Particles Objects?’, *Analysis*, 66, pp. 52–63.
Sider, Theodore 2001: ‘Maximality and Intrinsic Properties’, *Philosophy and Phenomenological Research*, 63, pp. 357–64.
Swinburne, Richard 1995: ‘Thisness’, *Australasian Journal of Philosophy*, 73, pp. 389–400.
Teller, Paul 1986: ‘Relational Holism and Quantum Mechanics’, *British Journal for the Philosophy of Science*, 37, pp. 71–81.

Van Inwagen, Peter 1981: ‘The Doctrine of Arbitrary Undetached Parts’, *Pacific Philosophical Quarterly*, 62, pp. 123–37.

- 1990: *Material Beings*. Ithaca, NY: Cornell University Press.

Wasserman, Ryan 2002: ‘The Standard Objection to the Standard Account’, *Philosophical Studies*, 111, pp. 197–216.