The Identity of Indiscernibles

MICHAEL WREEN

Abstract
A number of distinctions are needed to assess The Identity of Indiscernibles. The first concerns its purported status as true or, more strongly, necessarily true. The second concerns the nature of the properties the principle ranges over. One way to divide up properties is in terms of those that entail the existence of a particular object, those that entail the existence of an object but no particular object, and those that don't entail the existence of an object at all. A second way is divide up properties is in terms of 'indicative' and 'counterfactual' properties. Combining these three distinctions yields twelve different versions of the Identity of Indiscernibles. In the long run, however, four more are needed to do the principle justice. Sixteen different versions of the principle, then, need to be, and in this paper are, assessed. Particular attention is paid to methodology and to constraints on the generation of properties needed to support various versions of the principle. Bottom-up reasoning is found sufficient to vindicate the great majority of them, but there does not appear to be a good bottom-up or top-down argument for the two most interesting and important versions of the Identity of Indiscernibles, and there is some reason to think both false. An open mind needs to be kept, however, for definitively establishing or refuting them probably requires a relatively complete and plausible metaphysics.

Informally, the Identity of Indiscernibles is:

If 'two' objects share all their properties, they are really one.[1]

[1] Written as a stand-alone piece, this paper is intended to be comprehensible independently of acquaintance with the vast literature on the Identity of Indiscernibles. To ensure such self-containment and accessibility, references to the literature are virtually non-existent, and a number of issues that could be critically explored – especially (a) quantum physics and the Identity of Indiscernibles, (b) interpretations of Leibniz on the principle, and (c) the Identity of Indiscernibles and bundle theories of physical objects (all three of which are heavily represented in the literature) – are bypassed or no more than very lightly touched upon. Space limitations, clarity, and coherence would dictate such limitations, even if the desire for ready accessibility didn't. Available from the author upon request, however, is an extensive bibliography covering a fair amount, even if not quite all, of the recent literature on the principle, as well as on a number of topics in the near vicinity.

doi:10.1017/S0031819114000461 © The Royal Institute of Philosophy, 2014

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

Two points need to be made at the start. First, the 'objects' the principle speaks of can be any things at all, any items included in an ontology. That means that if things besides physical objects exist – events, numbers, sets, immaterial souls, properties, relations, possibilia, propositions, or functions, for example – then the principle ranges over them. That much is accepted by everyone who discusses the principle in any form. However, discussions of the Identity of Indiscernibles have concentrated almost exclusively on physical objects. There may any number of reasons for this. Two that come to mind are, first, there's widespread agreement that physical objects exist, but considerable disagreement about whether anything else does, and, second, there are difficulties enough with physical objects, so other objects, if there are such, can be considered later. This second reason is reinforced by the facts that physical objects may be easier to deal with than anything else, and an understanding of them may yield insights respecting other objects. That said, and in following tradition, only physical objects will be considered here.

Second, since the properties the principle ranges over might well determine whether or not it is acceptable, different ranges of properties need to be considered independently of each other. If the properties the principle ranges over are unrestricted, for example, a weak version of it will result, while if restricted in certain respects, a much stronger version will result. The Identity of Indiscernibles, in other words, comes in stronger and weaker forms, and the forms shouldn't be conflated. The need to be on the lookout for just such a conflation will become evident in what follows.

1.

In second-order logic, the principle – call it IOI – is:

$$(x)(y)(F)[(Fx \equiv Fy) \supset x = y]$$

The two questions to ask about it are: (i) Is it necessarily true? and (ii) Is it at least contingently true, by which I mean (using the term non-standardly), true in this world?

If 'F' is unrestricted, if it ranges over every single property, relational as well as non-relational, in other words, if the Identity of Indiscernibles is taken to be

(A) If 'two' objects share *all* their properties, then they are really one,

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

then it's almost certainly necessarily true. Consider a world in which it doesn't seem to be true, a world consisting of two iron spheres, a and b, which have exactly the same shape, size, chemical composition, and so on, and stand to each other as mirror images. Even in this world, there is one property that a has but b lacks: the property of being identical with a. As (1) this is a property not just that a has and b lacks, but that every other object, c, d, e, and so on, would also lack regardless of how populated or diverse the world; and as (2) there is a comparable property, *being identical with x*, for every other object x, about which (1) holds mutatis mutandis, regardless of how populated or diverse the world – as that is the case, there are no possible circumstances in which two numerically distinct objects share all their properties. If 'F' is unrestricted, then, if 'F' ranges over all properties, IOI is necessarily true. And if it's necessarily true, then of course it's true – that is, true in this world.

2.

Suppose, though, that we distinguish between properties that are purely qualitative and properties that aren't, with the former being properties that don't entail the existence of any particular object, and the latter being properties that have such an entailment. On this definition, *being square* is a purely qualitative property, but so is *being on top of a refrigerator*, since an object's having it doesn't entail the existence of any particular object, any particular refrigerator. *Being on top of Mount Everest*, however, as well as *being identical with a*, aren't purely qualitative properties, since an object's having either does entail the existence of a particular object, Mount Everest or a. Some purely qualitative properties are thus non-relational, in that they don't entail the existence of any other object at all – these are roughly the 'intrinsic properties' of objects – while some are relational, in that they do entail the existence of some object, although no object in particular. Slightly qualified, and slightly inaccurate even so, this last statement can be taken to define relational and non-relational properties in general: relational properties entail the existence of an object, non-relational properties don't. As the terms have been defined, then, there are purely qualitative non-relational properties (e.g., *being square*), purely qualitative relational properties (e.g., *being on top of a refrigerator*), and non-purely qualitative relational properties (*being on top of Mount Everest*). There are no non-purely qualitative non-relational

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

properties, because by definition, non-purely qualitative properties entail the existence of a particular object, and thus the existence of an object, and so are relational.

Suppose that 'F' ranges over purely qualitative non-relational properties (intrinsic properties). In other words, suppose we consider:

(B) If 'two' objects share all their purely qualitative non-relational properties, then they are really one.

Basically, (B) invites us to think of the 'two' objects only in relation to themselves, that is, only in relation to those properties, like *being square*, that they would have even if nothing else existed. So considered, (B) certainly doesn't look like a necessary truth. The world consisting of two iron spheres mentioned above is one in which it would appear to be false, for a and b are stipulated to 'look exactly alike', that is, to share all their properties that make no reference to anything outside of themselves.

Even if that's so, (B) still might be contingently true. It could well be that no two objects in this world share all their purely qualitative non-relational properties. That can be, and has been, disputed, however, for it's been argued that some quantum phenomena, for example, share all such properties, yet are numerically distinct.

But things are more complicated than appears here, and, as the next section will show, (B) will need to be revisited. Later still, it will need to be revisited again.

3.

What if IOI ranges over all purely qualitative properties, including relational ones? In other words, what about

(C) If 'two' objects share all their purely qualitative properties, then they are really one?

Here it might be argued that IOI survives as a necessary truth, for, in the two-spheres world, the property of being such that a tired philosopher would sit on it but on nothing else if he were to enter the world is a property of one sphere but not the other. Given any world, with any number and any number of kinds of objects, this property would be had by one object only, thus singling it out from the rest.

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

A number of doubts could be raised about this argument, however. The first concerns the legitimacy of subjunctive or counterfactual properties. If there are no such properties – and some philosophers have claimed just that – the argument certainly fails. This, however, doesn't seem to be a particularly strong objection, as counterfactual properties, though perhaps strange and puzzling in certain respects, seem as legitimate as their indicative brethren. Properties ride in tandem with the possible assignment of truth-values to propositions, and as we do assign truth-values to counterfactual propositions, and would be hard pressed to get along if we didn't, counterfactual properties seem respectable enough.

The second objection is that the property in question isn't one that doesn't entail the existence of a particular object. In effect, the property is that of being such that a tired philosopher would sit on it and it alone of all the objects in the world if, etc., and 'it and it alone' does pick out a particular object – the object that the tired would sit of if, etc. In other words, the argument fails because the property in question isn't purely qualitative.

The strength of this objection depends on the meaning of 'particular object'. The term is somewhat indefinite, and whether a definite description embedded in a property specification entails the existence of a *particular* object isn't settled by the term as it stands, that is, without clarification, qualification, or stipulation. Does, for example, 'the first man to come within 10,000,000 miles of the surface of a star' denote a particular object, in the requisite sense of the term, or just whatever object it may be that satisfies the description? In the absence of clarification or qualification of the key concept here, or perhaps stipulation regarding it, the force of the objection is also, at best, somewhat indefinite.

But that's 'at best'. Actually, if 'entails the existence of an object', particular or not, is taken strictly, there's no problem from the start. The definite description is buried behind the counterfactual supposition, and thus an object's having the property in question doesn't entail the existence of any object at all, except the object that has it.

The problem with this reply, however, is that it proves too much. If correct, it shows that the property in question isn't a relational property at all, *a fortiori*, not a purely qualitative relational property. Worse still, if the reply is correct, it really vindicates (B) as well as (C), despite what was said in the previous section. This isn't just counterintuitive in and of itself, but methodologically suspicious: a reply that proves so much so easily just doesn't seem right.

Although I will ultimately argue that this response is on the right track, and the property in question is indeed irrelevant to a

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

vindication of both (B) and (C) – this will emerge in section 6 – in the present context it could be replied that all that has been shown is that the argument for (C) needs to be muted in some way, not thrown out altogether.

The best way to do so, in order to avoid the untoward and counter-intuitive implications just noted (i.e., the property isn't a relational one; being non-relational, it thereby too quickly and easily proves (B)), is to adjust the definition of a purely qualitative relational property. One quick and simple way to do so is this: A property is a purely qualitative relational one if and only if (a) it's purely qualitative, and either (b) its instantiation entails the existence of at least one (possibly other) thing, or (c) the satisfaction of any counterfactual component (or components) it has would entail the existence of at least one (possibly other) thing. If needed or desired, a similar qualification could be added to the definition of a purely qualitative non-relational property. The adjustment made here is blatantly ad hoc and not nearly as precise as it should be, but for two reasons, I'd like to temporarily table further discussion of the definition, and indeed of the second objection in general. My reasons for doing so are two. First, even if the new definition blocks a proof of (B) while allowing a proof of (C) to go through, (B) still needs to be reconsidered in light of the above, and as far as it's concerned, the second objection, and any replies and counter-replies to it, are beside the point. Second and more important, what needs to be said about the definition can best be said after certain other matters are attended to. I'll make good on this claim in section 4.

4.

Consider once again

(B) If 'two' objects share all their purely qualitative non-relational properties, then they are really one,

but with counterfactual properties in mind, for example, the property of turning red if at least one thing but not everything were to turn red. (For purposes of illustration, I'll assume here and in what follows that redness is an intrinsic property of objects. Nothing of substance hangs on the particular property chosen in this case.) Being non-relational, this property doesn't entail the existence of any object other than the one that has it. In the world of two iron spheres, one object would have it and the other lack it, even if we couldn't tell

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

which. (If we like, we can imagine the change, if it were to occur, not to be subsumable under universal laws, just as some quantum phenomena aren't.) The property singles out one object from all the rest, and thus it looks as if, contrary to first impression, such as that of section 3 above, and despite the qualms aired at the end of the previous section, (B) is just as necessarily true as (A).

But we shouldn't start celebrating just yet. In the case of (A), coming up with the property of being identical with a was enough to meet the challenge to the claim that (A) is necessarily true, was enough to distinguish the two spheres even though they seemed to share all their properties and yet be numerically distinct. That, however, was because, and only because, the property in question was generalizable to all objects, no matter how many objects there happen to be. In other words, there are analogue properties to *being identical with a* – namely, *being identical with b, being identical with c*, and so on – for every object in a world, no matter how many objects the world contains and no matter what properties those objects have. With analogue properties of this sort, each object has a property that every other object lacks, and so no two distinct objects share all their properties or could share all their properties.

The properties invoked to distinguish the spheres when (B) and (C) were in question, however, can't be generalized in this way to yield the result that no matter what, an object will have a property that every other object lacks. (This, incidentally, is an objection to (C) not considered in the immediately preceding section, and is a much stronger objection than the two considered there.) Consider, for instance, a four iron-sphere world. The property of being such that a tired philosopher would sit on it, etc., will distinguish one sphere from the other three, but it won't distinguish the other three from each other, and, more importantly, it isn't generalizable, in terms of analogue properties, in a way that will allow us to distinguish each of the other three from each other and also from the first sphere. It thus won't let us conclude, on the basis of the iterative considerations that figured in the discussion of (A) – basically, point (2) back in section 1 – that for each object in a world, no matter how populated the world or what the properties of the objects in the world, there's a guarantee that every object has at least one property that every other object lacks. That's what's needed, and not merely a property that singles out one object from all the rest, a property that only one object has. Even an infinity of such 'singling out' properties won't do. An analogue property like *being such that a tired philosopher would lean on it*, etc., as opposed to *sit on it*, will distinguish whichever sphere has it from all others, and every other such analogue property

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

will do the same. That, however, doesn't come close to guaranteeing that for every object, there is at least one property that it has and every other object lacks (symbolically: $(x)(\exists F)\{Fx\cdot(y)[\text{-}Ixy \supset -Fy]\}$). All we get with each analogue property is a guarantee that one and only one object has whatever, in this case, purely qualitative relational property is in question. For all that, two numerically distinct objects may still share all their purely qualitative relational properties.

The situation is even worse for a purely qualitative non-relational property like *turning red if at least one thing but not everything were to turn red*. In a four-sphere world, three spheres might have that property. Better would be to focus on *turning red if at least one thing and nothing else were to turn red*. That's a property that can be had by only one object. It, however, faces exactly the same problem that purely qualitative relational properties do, the problem detailed in the preceding paragraph. Neither (B) nor (C) has been vindicated, then, and the initial judgment of (B), in particular, that it isn't a necessary truth, seems to stand.

5.

But that also may be a little too quick. Even if no good argument for the necessity of (B) and (C) has been provided, the situation hasn't been shown to be hopeless. So far, all that's been shown is that the properties mentioned above (and, by parity of reasoning, all their kin) are of no avail. Others, for all that's been said, may well do the trick. But what would these properties be, and how are they to be found?

The second question here is the crucial one. A method for finding the properties is needed. Since (A) is necessarily true, the natural place to begin looking for such a method is in the argument for its necessity. And an examination of that argument is revealing, for it shows that attention needs to be paid to the constraints on the properties needed, that is, the properties that the properties must have. The constraints are essentially those indicated by (1) and (2) back in section 1. It must be possible, first, to find a property that an object, picked at random, alone has, and second, on the basis of the nature of that property, to generate, in some iterative or recursive manner, like properties for every single object in the world, even if the world contains an infinity of objects. When (A) was under consideration, that was easy. Each object, x, was matched up with itself, via the property-frame *being identical with x*. With (B) and (C), however, non-purely qualitative property-frames, such as *being identical with x*,

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

are no longer available, so something other than iteratively generated self-identity properties is needed for the matching. But what?

Whenever a potentially infinite matching of the sort in question is needed, numbers and functions suggest themselves. A numerical function, say $f(x) \rightarrow x$, maps a number into another number, and only one number. It does this even if there's an infinity of numbers to map. In addition, some functions – once again, $f(x) \rightarrow x$ is an example – are such that no number is mapped into more than once. Thus, since there is an infinity of positive whole numbers, if they could be iteratively or recursively nestled into purely qualitative relational and non-relational properties in the right way, a way that guarantees that for each object, it and it alone has one of those properties, a vindication of (B) and (C) would be possible.

In light of this, consider the property-frame *being sat upon by exactly N people, were x people to come into existence and each person to sit on at least one object and each object to be sat upon by a different number of people*. 'x' here holds a place for the number of objects in whatever world is in question (assuming, for the sake of argument, that we can count objects as such; if not, x can simply be taken to be a countable infinity). Its value would remain constant within a world, but vary across worlds (again assuming that we can count objects), depending on the world's population. N, however, varies within a world, and, starting at one, is increased by one as objects are run through. This generates property after property, and, the content of each property being what it is, every object in a world would have exactly one of these properties, and thus no two objects would share all their properties. Put differently: in any world, only one object would have the property when 'one' is filled in for 'N'; only one, and a different one from the first, would have the property when 'two' is filled in for 'N'; only one, and a different one from the first two, would have it when 'three' is filled in for 'N'; and so on for every object in the world. Since this guarantees that no matter what the world, every object in it would have a property, a purely qualitative relational property, that that object and no other object has, (C) would be vindicated as a necessary truth.

One doubt that could be raised about this defense of (C) concerns content as opposed to form. That is, although the structure of property generation is secure, and it does result in object-property matching of the desired sort – every object being matched with a property that it uniquely has – the content of the property-frame, *being sat upon by... persons...*, may give pause in at least some worlds, namely, those in which being sat upon by even one person is a difficult business. In a world consisting of only clouds, for example, it would be

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

hard to rest one's feet, even counterfactually, and it would be even harder to do so in a world consisting of only sub-atomic particles. This suggests that even if there's a property-frame structure that will work across all possible worlds, there is no one property-frame content that will do the trick. Moreover, given the vast variety of physical objects and their properties, it would seem that there's no algorithmic way to generate the needed property-frame contents.

Although I can't definitively prove that this is only a minor problem, one that erects no very high barrier to accepting (C) as a necessary truth, I do think it's just that. Property–frame contents could be fadged up from world to world on an ad hoc basis, and tailored so as to guarantee that what holds for an iron-sphere world holds for every other world as well. There's no proof that we could come up with the goods in every case, of course, but given a specification of the objects in a world, the task seems, in no instance, a very difficult one. Certainly the imagination isn't taxed by cloud world and sub-atomic-particle world.

Much the same set of considerations holds for (B). A property-frame like *turning to its complementary color(s) and then reverting to its original color(s) exactly N times, were every object to turn and revert at least once and no two objects to do so the same number of times* would ensure every object in a world, no matter what the world, has a purely qualitative non-relational property that every other object in that world lacks. As before, 'N' would range over the positive integers, and would, at a minimum, mark the number of objects in the world. Factors respecting structure and content discussed above in relation to (C) also apply mutatis mutandis to (B).

6.

Now to pull the rug out from under much of what was said above – but in a sense to leave everything in place even so.

In section 2 non-purely qualitative properties were distinguished from purely qualitative properties, and the latter further divided into relational and non-relational properties. It's evident, however, that another distinction has also been at work in the last three sections, that between counterfactual properties and non-counterfactual or (for want of a better term) indicative properties. I take this last distinction to be roughly that between properties that include a counterfactual element in their specification and those that don't. Thus instead of having two questions – Is it necessarily true? Is it contingently true? – to ask about each of three principles, (A), (B), and

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

(C), the two questions have to twice asked about each of the three principles. First, six questions have to be asked when the properties ranged over include both indicative and counterfactual properties, and second, the same six questions have to be asked if the properties ranged over don't include counterfactual properties.

That, in fact, is why the revised definition of 'purely qualitative relational property' tendered back in section 3 won't do: its '(b) or (c)' clause effectively blurs the distinction between two forms of IOI, two forms which differ markedly in their strength and utility. (And that is also why the last few sections make good on the claim that ends the introductory section of this paper.) Section 3 could not possibly show that (C), as it is supposedly being considered there, as ranging over indicative properties alone, is correct. Counterfactual people barging into worlds are utterly irrelevant to the truth of such a claim. Two different kinds of purely qualitative relational property thus need to be distinguished, one as defined by the '(b)' clause of the definition back in section 3 – these are indicative purely qualitative relational properties – and another as defined by the '(c)' clause of that definition – these are counterfactual purely qualitative relational properties. The same distinction between the indicative and the counterfactual also holds, of course, for purely qualitative non-relational properties.

Put into tabular form, then, the results obtained so far are really these:

	(A)		(B)		(C)	
	Nec.	Cont.	Nec.	Cont.	Nec.	Cont.
Indic. Prop.	True	True		Maybe		
& Count. Prop	True	True	True	True	True	True

The second row indicates that with counterfactual properties included, IOI comes out true no matter what – though showing as much requires the argumentation of the previous sections. The result here is of little interest or utility, however, as the counterfactual properties that ensure the necessity of (A)–(C) are wild and wooly philosophers' toys, of small metaphysical significance and little use in any possible world, or even the actual world, for solving the sorts of practical and philosophical problems that typically arise. Certainly they don't help in the epistemic process of actually distinguishing objects from each other – of actually *discerning* differences and diversity, in other words, despite the fact that the principle is

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

called the Identity of *Indiscernibles*. Nor do they help with any other philosophical or non-philosophical task, including any metaphysical task, that I can think of. That, in fact, also holds for the necessary truth of (A) when only indicative properties are ranged over. The attribution of *being identical with a* and such like by an epistemic agent would require already determining that a and b are two, presumably on the basis of spatio-temporal placement relative to an actual or a possible – that is, a counterfactual – agent or object. Thus anything needed to solve a more substantial philosophical or extra-philosophical problem would probably be available independently of and prior to use of an (A)-form of IOI, or else wouldn't be available at all. Having a necessary truth can be like having millions of dollars in the bank – great funds to draw upon for important projects – but it can also be like having millions of shares of worthless mining stock – no cash value and no dividends.

Of the remaining four boxes in the table, one, the Maybe, notes that the debate over whether every object in this world has some indicative 'intrinsic' property not had by any other object, is still unresolved. This was mentioned briefly at the end of section 2. But the empty box at the end of the row, it seems obvious, should be marked True. Given the wide diversity of objects about us and the fact that our world isn't point, line, or plane symmetric, every object is virtually guaranteed to have at least one purely qualitative relational property not had by any other object.

Or at least seemingly so – that, too, can be disputed. Even if the world isn't spatially symmetric, it is, according to some philosophers, temporally symmetric. Nietzsche's doctrine of The Eternal Recurrence, at least interpreted one way, entails as much, as do similar views held by a number of other philosophers, including Plotinus, Origen, the Stoics, and some pre-Socratics. Temporal symmetry is also held and by a number of contemporary physicists, though for reasons very different from those that have attracted philosophers. If Nietzsche and friends are right, (C) may not even be contingently true. For the moment, let's put a Maybe in the box.

The remaining two boxes are also blank. Nothing said above offers the least reason for thinking that (B) or (C), confined to indicative properties, is necessarily true.

7.

Worse still, (B) and (C) don't look necessarily true, and the two-sphere world considered above seems to show that appearances

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

aren't deceptive in this case. With indicative purely qualitative relational and non-relational properties –hereafter, simply relational and intrinsic properties – as the only resources to draw upon, and seemingly nothing there of any use, falsehood seems inescapable.

But it's still not time to hoist the white flag. Spatio-temporal properties may come to the rescue. One iron sphere has the property of being located at m,n,o,p, the other of being located at q,r,s,t. The spheres are thus discernible, and, better still, the property-frame *being located at x,y,z,w* is easily generalizable in the way indicated in section 5: no matter what the objects in a world – no matter how many uniform spheres a world may contain, for instance – there's a guarantee that every object in that world has a spatio-temporal property that every other object lacks. Whether this vindicates (B) and (C) or (C) alone depends on whether space and time are absolute or relative. Clearly, though, (B) and (C) need not both surrender.

While I think that, strictly speaking, this does show that one or both principles are necessarily true (and hence that the first row of the table should either read True all the way across, or that one blank needs to be filled in with True and one of the Maybes changed to True), there are at least two possible rejoinders, one of little moment, the other of much greater significance.

The first is that there are no temporal or spatial properties. This, like the view that there are no counterfactual properties, is little more than special pleading to avoid the conclusion, and highly counterintuitive in light of our handling of temporal and spatial predicates in our descriptive and explanatory practices.

The more powerful reply is to accept the conclusion, but argue that more than anything else what it shows is the inadequacy of the notions of an intrinsic property (a purely qualitative non-relational property) and of a relational property (a purely qualitative relational property) as defined above. To focus just on the former, the claim would be that although spatial-temporal properties satisfy the definition of an intrinsic property, intuitively they don't seem to be intrinsic properties at all. The intrinsic properties of a car, for example, don't seem to include where it's parked on a certain night. Related to this is the fact that our pre-theoretical understanding of intrinsic properties is that they are characteristics of things as they are in and of themselves, and provide answers to the question, What's it like? Where is it, and where was it or will it be on such-and-such a date, are very different questions. The same holds if space and time are relational, not absolute.

Also worth noting in this regard is that our concept of a physical object as such is tightly tied to that of spatio-temporal locale. It's

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

arguable that the most fundamental principle of individuation for physical objects invokes spatio-temporal properties: x is identical with y if and only if the spatio-temporal co-ordinates of x and y are the same. In effect this reinforces the point that, strictly speaking, spatio-temporal properties demonstrate the necessity of (C) and possibly (B). Since the principle is universal in scope, however, it also demonstrates that the mere use of it in establishing identity or diversity tells us nothing about what x and y are or are like, nothing about their intrinsic properties.

In short, the reply here is that refined definitions of 'intrinsic property' and 'relational property' are needed to prevent the two from salvaging (C) and possibly (B) in a facile way, and that this isn't special pleading but based on the anomalous status of temporal and spatial properties. In slightly different terms, a further distinction is needed – the table in the previous section is incomplete as it stands – to capture all that's wanted in an examination of (B) and (C).

I think this objection is correct: the table of section 6 should read True all the way across, or at most one blank and one Maybe should remain, but a new table needs to be drawn up, one with a modified (B) and (C), a (B') and (C'). (B') and (C') should incorporate an understanding of intrinsic and relational properties that captures our fundamental intuitions about both, while excluding the sort of spatio-temporal properties in question. No definition of 'intrinsic property' or 'relational property' that divides up the terrain in the way just indicated will be proposed and defended here, however, as a careful and rigorous treatment of the issue would require a paper at least as long as this one. In order to continue to pursue the issue at hand, namely, the acceptability of various forms of IOI, I'm simply going to avail myself of new, refined notions of relationality and intrinsicality, whatever they may be, and march onward.

The new table is:

	(B')		(C')	
	Nec.	Cont.	Nec.	Cont.
Indic.		Maybe		Maybe

8.

Is anything left? Are there any properties which argue for filling in the blanks with Trues? Haecceities, primitive thisnesess, suggest

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

themselves, but their status is unclear. If they aren't intrinsic or at least purely qualitative properties, then it's at least arguable that their existence tells *against* The Identity of Indiscernibles; and not only did Scotus, who is most closely associated with the notion, explicitly deny that an haecceity is a universal or form – in today's parlance, a qualitative property – but it's hard to make sense of the notion as such. A very rough explication of an haecceity, considered as a qualitative property, might go something like this. The words 'being this particular thing and no other' and 'having this particular existence and no other' apply to each individual thing, but what is indicated by the words is, in each case, a different property, its haecceity. The haecceity of each thing has something to do with the property of being that particular thing and something to do with the property of existence, or having existence, or having that particular existence, but each haecceity is different from those properties, strictly speaking, for they're properties had by every thing, while each haecceity is had by only one thing.

The notion thus explicated isn't an impossible one, but it is obscure and hard to work with – linguistically, an intrinsic specification of even one haecceity seems impossible. A great deal of work is being done by the indexicals 'this' and 'that', which are taken to cover a content that is, in effect, ineffable. Apart from that, the notion is a shadowy kin to that of being identical with x, even if an haecceity is non-relational property. Each haecceity would be the basis of a relational property on the order of *being identical with a, being identical with b*, and so on. As such, each object having a different haecceity would amount to little more than each object being the object it is and not some other object. True enough, but not very enlightening or useful. Given the lack of clarity and utility of haecceities, then, they're probably best put to the side.

9.

Other than haecceities, try–try–try–try–to think o' something different as I might, nothing comes to mind.

Even so, that doesn't automatically mean that Falses should be filled in the empty boxes. For one thing, the failure to come up with anything may speak only of the limits of my imagination. Perhaps a King Lear will emerge to teach us differences. Still and all, the fact that no one, and not just myself, has come up with anything over the years is discouraging, and gives us some reason, though not a knockdown reason, or perhaps even a particularly strong one, for

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

thinking (B') and (C') false (here and in what follows, I'll be speaking of (B') and (C') only as possible necessary truths).

I say the reason here may not be a particularly strong one, and I want to emphasize and argue for the point at some length. With putative necessary truths which aren't axiomatic or self-evident, what's needed to establish them, that is, to provide adequate justification for believing them, is a proof – a mathematical, logical, or metaphysical demonstration of some kind, such as those offered for the infinity of primes, for 'if G or H, then either if G then H or if H then G,' and for (A). On the other hand, what's needed to establish that such a proposition is false, that is, to provide truly adequate justification for believing their denials, is again some sort of demonstration – a proof in mathematics, logic, or metaphysics. A counterexample is one sort of disproof – a proof of the denial of a putatively necessarily true proposition – in mathematics and logic because counterexamples and counterarguments in such fields involve limited and clearly defined objects, facts, and properties. Because of that, very little is left to guesswork, no epistemic question of great or even middling significance remains that something might be there and relevant, and yet have been overlooked. The claim that all integers greater than 1,002 are divisible by three (without remainder) can be shown not to be a necessary truth by citing 1,003 and dividing. No very strong doubt remains that something might have been missed that would overturn the counterexample and disproof – except, of course, the fact that human beings are fallible and can make mistakes about absolutely anything, including the division of small numbers – because what needs to be considered is minimal and well defined. The claim that 'if M then N' is a necessary truth (where neither 'M' nor 'N' is a necessary truth or falsehood and the two are logically independent), can be shown to be false by citing a world in which M obtains but not N. The confined, limited, and clearly defined facts that need to be considered provides a virtual guarantee that nothing relevant has been overlooked which would overturn the counterexample, and show it not to be such. Falsity is thus established in such cases, at least if we once again put to the side the ever-present possibility of human error.

But it's just such a virtual guarantee that's missing in the case of the two-iron-spheres-world counterexample to (B') and (C'). After all, many philosophers have been or still are tempted to declare IOI false in *all* forms when presented with the two-spheres world. They miss (or dispute the relevance of) *being identical with a* and the iterative function it suggests, as well as the other properties and functions discussed above, and thus fail to see that the counterexample doesn't

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

take relevant facts into account. The same may hold for the two-spheres world and (B') and (C'): it seems to be a counterexample to (B') and (C'), but given the infinite, open-ended, and not clearly defined set of properties that are relevant to their truth or falsity, there's no very great assurance that something relevant hasn't been missed, something that would upset the counterexample.

In slightly different terms, the point is this. The two-spheres world looks like a counterexample, and thus a demonstration of (B')'s and (C')'s falsity, but the backing needed to be as epistemically assured as we should be, that is, to have a fully adequate justification that it is a counterexample, isn't there. All we – myself and others – have done is look for a property and iterative function to salvage (B') and (C') in such a world and come up empty handed. That's not too different from looking for a counterexample to a supposed necessary truth and coming up empty handed, which is no very good evidence that the proposition is necessarily true. The Goldbach Conjecture, for example, certainly seems to be true – it holds for every even number greater than two that people have examined – but absent what's needed for adequate epistemic assurance that it is true – a proof – we can't be sure that something hasn't been overlooked, that there isn't a counterexample that no one has come across yet, or even a counterexample that's so large that humanity, even aided by computers, will never be able to detect. With Goldbach's Conjecture, we should either be agnostic, or at best, given what's required for truly adequate justification, tailor our confidence to the weak, basically inductive support that it has. Although the case of (B') and (C') concerns a disproof, based on a counterexample (the two-spheres world) to two supposed necessary truths, while that of the Goldbach's Conjecture concerns a proof, based on positive instances of a supposed necessary truth, in neither case has a demonstration been provided, for in neither case is there the required assurance that nothing has been overlooked. Nor, in the case of (B') and (C'), is this mere human fallibility at work, the mere fact that we might have missed something. Base-line fallibility of that sort is ubiquitous, a factor present in the case of every proof or disproof, including that of the falsity of all integers greater than 1,002 being divisible by three. Rather, it's due to what's required in the nature of the case of a proof or disproof of a putative necessary truth of the sort at hand: epistemic assurance that the nothing in the inherently open-ended, infinite, extremely various, and ill-defined set of objects and properties that are relevant and that would vindicate (B') and/or (C') has been neglected. Again, the comparison with Goldbach's Conjecture is a propos.

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

10.

The epistemic issue here is connected with the assessment of *argumenta ad ignorantia*. If I argue

> there's no reason to think that there is milk in the refrigerator;
> therefore, there isn't milk in the refrigerator,

the argument may be very strong, very weak, or anything in-between, depending on how extensive and careful my search of the refrigerator has been, and how extensive and careful a search must be to warrant the inference. If the refrigerator is small and not crammed with items, two minutes of attention is plenty; if it's the size of a barn and stuffed to the gills, an hour may not be enough to warrant an inference of more than minimal strength. And if I don't look at all, then even if the premise is true the inference stinks; but if I search for an hour in a typical refrigerator and do so with great conscientiousness, the inference is terrific even if the premise is false (and all the food spoiled). It's the nature of the case that tells us what's required for how strong an inference. If I argue that

> there's no reason to think that the Goldbach Conjecture is false;
> therefore, it's true,

and the only backing I have for the premise is that other people and I have run through a lot of numbers, perhaps millions, in search of a counterexample and found nothing, the inference is extremely weak. A lot is required in this case, because so many – an infinity – of factors are potentially relevant, and there is no way to tell whether an unexamined number won't upset the apple cart. On the other hand, if I argue that

> there's no reason to think that anything relevant has been missed in the counterexample-counter-argument to the claim that has been every integer greater than 1,002 is divisible by three;
> therefore, nothing relevant has been missed;
> therefore, it's false that every integer greater than 1,002 is divisible by three,

the argument is very good if the usual and not very extensive sort of care has taken, just because not many factors needed to be looked at. Not much is needed, in other words, because of the nature of the case, while much more is needed, because of the nature of the case, when it comes to the Goldbach Conjecture. It is, in fact, arguable that no finite, inductive search can provide anything at all in the way of a justification for believing the Conjecture. (Distinct and much more

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

implausible is the claim that an extensive search provides no justification for acting in any circumstances as if the Conjecture is true.) If I argue that

> there's no reason to think that any relevant factors have been overlooked in the counterexample-counter-argument to (B') and (C');
> therefore, no relevant factors have been overlooked;
> therefore, (B') and (C') aren't true,

the strength of the inference depends not just on the truth of the premise but on what's required for a sufficiently thorough search for relevant factors, and how extensive and careful a search has been conducted. Given the nature of the case and what we've learned from an examination of other, weaker forms of IOI, a lot is required: a property that overturns the inference may be just around the corner, and there's no settled or simple method (as there is in the case of the divisibility-by-three example) of telling how many corners there are, or where the corners are. The two-iron-spheres argument is thus closer to the Goldbach Conjecture argument than to the divisibility-by-three and refrigerator arguments (the last of which doesn't concern a purported necessary truth, of course, though that's irrelevant in this context). The failure to provide a property and function that would vindicate (B') and/or (C') in such a world is therefore a relatively weak reason, though not (as may be the case with the Goldbach Conjecture) no reason at all, to think them false.

11.

Independent of all that, there's another reason why the mere failure to come up with something doesn't provide sufficient assurance that the blanks in the table back in section 7 should be fill in with Falses. So far, all attempts to vindicate principles like (A), (B), (B'), (C'), and so on have proceeded from below. That is, the argumentative strategy employed has been to rummage around in an attempt to find a property that, plugged into the proper sort of iterative function, will yield the desired result. From the property the ascent is made to the principle.

Browsing and sluicing in this way, panning for gold in an a priori manner, is fine as far as it goes, but not every proof in mathematics, logic, or metaphysics is or needs to be bottom-up. A top-down proof will do just as well, and actually better if it can be pulled

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

off. To be clearer and more explicit about what's meant here: some proofs in mathematics, logic, and metaphysics start from other principles and work their way down, or sometimes across (if the spatial metaphor isn't being stretched too thin here), to get to what's wanted. Historically speaking, in fact, that's been the norm as far as IOI is concerned: arguments for the Identity of Indiscernibles, usually in a very strong form, have proceeded by way of purely metaphysical, highly theoretical considerations. If any such argument succeeds, failure to be able to cite or specify a property difference between objects would be of no moment, for there would be a guarantee that there's a difference there, a guarantee that no matter what, every object has an intrinsic or relational property that no other object has. The search for a specific property could be abandoned, and the no-reasons argument for the falsity of (B') and (C') declared to be devoid of strength altogether. But that's only if there's a good top-down metaphysical argument to be had.

12.

For brevity's sake, I'll consider just two such arguments here.

Strangely enough, the first is due to A.J. Ayer. After airing a two-iron-spheres-like argument against IOI, apparently in the form of (B') or (C') or something very like them, Ayer goes on to say that

> in framing such [two-iron-spheres] examples a use is freely made of the distinction between subjects and predicates. It is assumed, for instance, that [we are] entitled to refer to a plurality of [iron spheres] in order to raise the question whether they are. . . discernible. . . . And this being so it is not surprising that [such an] example should appear to refute the principle of the identity of indiscernibles. . . . But it is just here, it may be argued, that the mistake arises. Being accustomed to speak of objects in the plural without further qualification, we come to think that this alone is enough to differentiate them. We are led to treat numerical difference as a relation among others, which may obtain independently of anything else. But this, the argument continues, is an illegitimate extension of the concept of number. It is only where there is discernibility that it makes any sense to talk of a plurality of objects. In a wholly predicative language, these

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

> fanciful universes, in which things are supposed to differ only numerically, could not even be described.[2]

Although he doesn't find this argument altogether convincing, Ayer is

> disturbed by. . . the consequences of rejecting it. . . . Suspicious of the category of substance [as he is, he] cannot see how asserting that an individual exists can be to assert anything more than that some predicate, or set of predicates, is instantiated. [He is thus] still inclined to hold that the principle of the identity of indiscernibles is necessarily true.[3] (page 35)

The talk of subjects and predicates here, and of number requiring discernibility, is more superstructure than basic argument. As I see it, the basic arguments are essentially two. The first is:

(1) The two-spheres world is a counterexample to (B') and (C') if and only if the notion of substance is viable.
(2) But the notion of substance isn't viable.
(3) Therefore, the two-spheres world isn't a counterexample to (B') or (C').

That's to show that the two-spheres counterargument doesn't work. More positively, and more directly in favor of (B') and (C'), is a second, more implicit argument:

(4) An object is a collection of intrinsic and/or relational properties.
(5) So if and only if the collections of such properties are the same, the objects that the collections are, are the same.
(6) Hence, (B') and (C') are true (or at least (C') is – though this qualification will henceforth be ignored for simplicity's sake).
(7) Moreover, since the key premise here, (4), is a metaphysical one that concerns the nature of objects, (B') and (C') aren't just true, but necessarily true.

It's thus Ayer's overarching metaphysical commitments that lead him not just to reject the counterexample, but to argue that no two objects share, or can share, all their intrinsic and relational properties. If he's right, we have, in effect, an a priori guarantee that even if we can't find a difference between two objects, there still is a difference

[2] A.J. Ayer, 'The Identity of Indiscernibles', in A. J. Ayer, *Philosophical Essays* (Macmillan Publishing, London (1954)), 34–35.
[3] *Ibid*. 35.

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

between them, and it's our own limitations that prevent us from finding it.

A detailed discussion of arguments (1)–(3) and (4)–(7) isn't possible here, but a few brief comments can be made. First, premises (1) and (4) depend on (a) how the notions of substance and property are understood, (b) why Ayer thinks the former an illegitimate concept, and (c) the strength of his positive arguments for objects as bundles of properties. At the very least, these are highly controversial matters.

Second and less programmatically, if Ayer were right, a two-spheres world would be an impossibility. It certainly doesn't seem to be such, though, a fact that can be underscored by the following thought experiment.

Imagine that there are two very similar iron spheres in our world, and that at a certain time objects other than the spheres start to pass out of existence one by one, though very rapidly. Nothing comes into existence in their place. While this is happening the differences between the spheres are themselves disappearing – as time goes by, they share more and more properties. After a while no object in addition to the spheres remains. At that point the spheres differ in only a small number of intrinsic properties. Then the spheres change one last time and share all their intrinsic properties. If Ayer were right, the spheres would have to pass out of existence at that point as well, for the world would suddenly become an impossible one. That, however, is not only highly counterintuitive, but leaves why the spheres would have to vanish, why the universe suddenly becomes an impossible one, as a complete mystery.

What an unsympathetic critic might conclude from this is that Ayer's argument needs to be stood on its head: what has been proved is that objects aren't collections of intrinsic and relational properties – premise (4) is wrong. Premise (2) might also be denied and the notion of substance embraced, but the conclusion of Ayer's first argument can be avoided without taking on such a, so to speak, substantial metaphysical commitment. If spatio-temporal properties are brought into the picture and used to individuate objects, (1) can be jettisoned and the viability of the notion of substance left as an open question.

A caveat needs to be appended to all this. The objections just aired depend on the thought experiment's being coherent. If, however, contrary to appearances, it harbors a hidden contradiction, the last two paragraphs come to naught. Criticism of Ayer's argument would then have to focus on points (a), (b), and (c) of four paragraphs

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

back. The burden of proof, however, rests on the claim of incoherence.

13.

A second argument is inspired by Leibniz, a great defender of even the strongest forms of IOI. Without going into textual details, and without even claiming that the argument that follows is Leibniz's – this, to disarm the zealots known as historical scholars of their wonted and often fully justified fury – the argument is

(1) The spatio-temporal co-ordinates of any object are essentially linked to, and a function of, the nature of the object.
(2) The nature of an object comprises all its intrinsic properties.
(3) Hence, since two objects necessarily have different spatio-temporal co-ordinates, they necessarily have different natures, different intrinsic properties.
(4) Thus, necessarily if 'two' objects shared all their intrinsic properties, they would really be one.

Leibniz makes use of The Principle of Sufficient Reason in his full argument for IOI, but that point aside, as this argument makes plain, he rejects premises (1) and (2) of Ayer's argument. In opposition to the suggestion made near the end of the last section, however, he also rejects spatio-temporal locale as a principle of individuation – more accurately, he rejects it as the most metaphysically fundamental principle of individuation. The nature of the object itself is a more basic principle, since it determines spatio-temporal locale. Spatio-temporal locale is an extrinsic denomination, and thus inherently a function of objects as they are in themselves. In effect, this amounts to holding that it's the complete concept of an object that individuates it, that the complete set of properties that an object has is necessary and sufficient for its being the object it is.

The two places at which this argument might be questioned are the first and second premises. If they pass muster the rest follows.

Consider a universe consisting of only an iron sphere. (Although the matter is too complicated and far afield to discuss here, I take this to be a real possibility for Leibniz, even though he holds that space and time are relational.) On the first premise of the argument, its position in space and time would be determined by its being red, round, smooth, and so on. Prima facie, that's quite implausible. Perhaps some of the implausibility can be dissipated if relational

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

properties respecting God are also taken into account, but Leibniz himself doesn't go that route, and it seems a refuge of desperation anyway, and one that needs to be explained and justified. If we table the point then, and move on to the two-spheres world considered earlier, the situation doesn't improve. The fact that the spheres are two would be a function of an unknown, and actually unknowable, difference between them. The difference would be unknowable, because however extensive our stipulative specification that the spheres have properties A, B, C, D, and so on in common, there being two would necessitate there being some further property, somewhere out at the limit of infinite specification, that they differ in, a property that accounts for their different locales. That's not impossible, but without a good argument in its favor, it strains credulity.

The second premise, a linchpin of Leibniz's metaphysics, is well known and needn't be discussed at length. The idea that the concept of an individual contains all that's true of it, that every property an individual has it has essentially – *baying at the moon one fine night*, for example, thus being an essential property of Walter L. Weber if he did, does, or will bay at the moon one fine night – is odd and counterintuitive, but, as with the first premise, not impossible. For Leibniz, its oddity speaks only of our inability to fully grasp the concept of Mr. Weber and our misplaced confidence that we do. Far from thinking the second premise counterintuitive, God would find it obvious.

The thought experiment discussed in the last section is a problem for any argument that purports to establish a strong form of IOI, Leibniz's included. Here, it bears on both the first and second premises, showing from another direction that in our intuitive understanding of things, spatio-temporal locale isn't a function of an object's intrinsic properties, and objects don't have all of their properties essentially.

The objections to Ayer are of necessity somewhat unfair, and those to Leibniz much more so, because both are taken very much out of context. Leibniz's adherence to strong forms of IOI is embedded in a rich, complicated, and profound metaphysics, and in the absence of a full and sympathetic exposition of that metaphysics, dismissal of any arguments of his should be met with a great deal of caution.

14.

The main findings of this paper can be briefly summarized. The Identity of Indiscernibles can be considered as a necessary or a contingent truth, but in either case, different versions of it need to be

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press

distinguished, in accordance with the type of property or properties ranged over. The great majority of appropriately distinguished versions of the principle, including those indicated in the table of section 6, are true. Not as clear is the status of its two strongest and most interesting forms, those in which what alone figure are indicative intrinsic and relational properties, restricted so as to exclude certain spatio-temporal properties. In neither case could a good bottom-up or top-down argument for the principle, considered as a necessary truth, be found. That provides only some reason for thinking them false, however, and not as much as might be suspected. A thought experiment adds to the case against them, but although things look bad, an open mind needs to be kept. Perhaps the celebrated Leibniz is even now looking down on us, wryly smiling at our ignorance, error, and presumption.

MICHAEL WREEN (michael.wreen@marquette.edu) *is Professor of Philosophy at Marquette University. He is the author of about 70 articles in logic, medical ethics, aesthetics, argumentation theory, the philosophy of language, and other areas of philosophy. Recent articles include 'A P.S. on B.S.' in* Metaphilosophy *(which is indeed on, but hopefully not, b.s.) and 'The Ontology of Intellectual Property' in* The Monist. *His 'The Standing Is Slippery' was published in* Philosophy *in 2004.*

https://doi.org/10.1017/S0031819114000461 Published online by Cambridge University Press