THE TRIVIALITY OF THE IDENTITY OF INDISCERNIBLES

Samuel Z. Elgin[1]

Abstract

The Identity of Indiscernibles is the principle that objects cannot differ only numerically. It is widely held that one interpretation of this principle is trivially true: the claim that objects bearing all of the same properties are identical. This triviality ostensibly arises from haecceities (properties like *is identical to a*). I argue that this is not the case; we do not express a triviality with haecceities, because it is impossible to express the haecceities of indiscernible objects. I then argue that this inexpressibility generalizes to all of their trivializing properties. If there is a trivial version of the Identity of Indiscernibles, it is a version that we cannot express.

Introduction

The Principle of the Identity of Indiscernibles (hereafter, the PII) is the principle that objects cannot differ *solo numero*; there can be none that are numerically distinct, yet are the same in all other respects.[2] This principle has troubled philosophers at least since Leibniz (1991). For it simultaneously seems to be a principle that *must* be false—while it also seems to be one that *must* be true.

A canonical counterexample was introduced by Black (1952). Consider a world that contains only two perfectly homogenous and qualitatively identical spheres. These spheres share all of the same properties; they have the same composition, mass, size, electric charge, temperature, color (and so on) as one another. As far as I know, no one holds that this is how the world actually is—but it seems perfectly clear that this situation is metaphysically possible. What's more, this world violates the PII. By stipulation, the spheres differ numerically, yet are the same in all other respects. So it *is* possible for objects to differ only numerically, and the PII is therefore false.

But another argument apparently establishes that the PII is true—and trivially true at that. Indeed, nearly every contemporary discussion of the Principle begins by acknowl-

[1]My thanks to Michael Della Rocca and Catherine Elgin for their helpful feedback on earlier drafts of this paper.

[2]Here, 'can' and 'cannot' have modal force; the PII is the principle that it is metaphysically impossible for objects to differ only numerically. Some (e.g., Casullo (1984)) defend a contingent version of the PII—according to which there happen to be no objects which differ *solo numero*, but there could have been. See French (1989) for an objection to this version.

edging this triviality.[3] On one interpretation, the PII amounts to the claim that objects that bear all of the same properties are identical. The PII is thus the converse of Leibniz's Law—according to which identical objects bear all of the same properties. This interpretation is held to be trivial due to the existence of haecceities: properties like *is identical to a*. Any objects that bear of the same properties (in general) bear the same haecceities (in particular). And, quite clearly, all objects that bear the property *is identical to a* are identical to one another.

We can formalize this triviality using the resources of higher-order logic: a language in which we represent properties by binding open sentences with λ-abstracts.[4] For example, we represent the property *is red* with '$\lambda x.(x \textit{ is red})$' and the property *is identical to a* with '$\lambda x.(x = a)$'. Higher order systems are strictly needed to quantify over properties—quantification used to formalize '*a* and *b* bear all of the same properties' (which we do with '$\forall \lambda X.(Xa \leftrightarrow Xb)$'). Within this language, we can prove that objects bearing the same properties are identical as follows:

i.	$\forall \lambda X.(Xa \leftrightarrow Xb)$	Supposition
ii.	$\lambda x.(x = b)(a) \leftrightarrow \lambda x.(x = b)(b)$	*i*, $\forall$-Elim
iii.	$\lambda x.(x = b)(b)$	Classical Logic
iv.	$\lambda x.(x = b)(a)$	*ii, iii,* $\leftrightarrow$ *Elim*
v.	$a = b$	*iv*, β-conversion

The truth of the PII thus becomes a simple matter of logic: trivial in every sense of the word. As a result, a central interpretive puzzle—perhaps even *the* central interpretive puzzle—of the PII concerns what a nontrivial version of the principle would be.

I deny this triviality—proof notwithstanding. I hold that the claim 'Objects that bear all of the same properties are identical' is either false, or else substantively true. More precisely, I deny that a trivial version of the PII *can be expressed*. From this, it immediately follows that whatever it is philosophers *have* expressed is no trivial version of the PII. I

[3]The first discussion of the trivial version of the PII occurs in Whitehead and Russell (1952). Other philosophers who endorse triviality include—but are not limited to—O'Connor (1954); Adams (1979); Katz (1983); Hoy (1984); French (1989); Della Rocca (2005); Wörner (2021) and Goodman (Forthcoming). One philosopher who denies that this triviality counts as a version of the PII is Rodriguez-Pereyra (Forthcoming) (on the grounds that objects could satisfy this, while differing only numerically). I do not wish to engage in a debate over whether this counts as a version of the PII here. It is a principle—by whatever name we give it—one that may be trivial or substantive.

[4]I won't go into much depth about the details of this system here. Aside from the following brief derivation, higher-order logic plays a minimal role in my argument. Suffice it to say that this is a typed, higher-order language with λ abstraction and two basic types: a type *e* for entities and a type *t* for sentences. Additionally, for every types τ_1 and τ_2, $(\tau_1 \rightarrow \tau_2)$ is a type, and nothing else is a type. We allow for infinitely many variables of every type—as well as the corresponding λ abstracts needed to bind them. In this language, we can represent 'Any two objects bearing all of the same properties are identical' as '$\forall \lambda x, y. \forall \lambda X.(Xx \leftrightarrow Xy) \rightarrow x = y$.' In this representation—as well as the subsequent derivation—the types of terms are omitted because they are contextually evident.

leave it as an open question whether there is a trivial version of the PII that cannot be expressed; if there is one, then it is a version that we have not (and, indeed, *can*not) state.

The argument for inexpressibility will occupy the entirety of this paper; a bare-bones version of it is this. Consider a world with two indiscernible objects, and suppose we attempt to refer to one of the indiscernibles with the name '*a*' (in order to express the haecceity 'is identical to *a*'). There are two possibilities in this situation: either '*a*' refers ambiguously, or it refers unambiguously. That is, either '*a*' ambiguously refers to both objects, or else it unambiguously refers to only one of them.

If 'a' refers ambiguously, then haecceities do not trivialize the PII.[5] The phrase 'is identical to *a*' either denotes two distinct properties (ones that each of the indiscernible objects bears)—or else a property that both of them bear. Neither alternative is trivializing, for the phrase does not denote a single property that only one object can bear.[6] The previous proof presupposed that '*a*' is unambiguous, and so says nothing about cases where the reference of '*a*' is ambiguous.

If 'a' refers unambiguously, then something disambiguates its reference: something in virtue of which '*a*' refers to the particular object that it does. What this feature is depends on the one, true metasemantics of proper names—but it must be something that distinguishes *a* from all other objects to secure a determinate referent (that is, in order to account for the fact that '*a*' refers to *that* object, rather than some other). Precisely because *a* has this uniquely distinguishing feature, it is not truly indiscernible. And because *a* is not indiscernible, the PII says nothing about this case. To put it another way, if there were a true instance of indiscernibility, it would be impossible to refer to one of the indiscernible objects without thereby referring to all of them. Without the ability to refer to individual indiscernibles, we cannot express haecceities that trivialize the PII.

The structure of this paper is as follows. I begin by surveying the literature on inexpressible ignorance. There are several reasons to start with this discussion. It introduces the phenomenon of inexpressible states of affairs—which is precisely what I claim holds for the haecceities of indiscernibles. But the two are not only analogous, but intimately related; I will also argue that every case of indiscernibility generates inexpressible ignorance. Additionally, this literature serves as an argumentative foil. Some potential objections to my view apply equally well to all cases of inexpressible ignorance. Those who embrace inexpressible ignorance ought not be moved by these objections. In the following section, I argue that the haecceities of indiscernible objects are inexpressible (along the lines outlined above) before arguing that this inexpressibility generalizes to every trivializing property that they bear. I close with some brief remarks on the implications this has for the Identity of Indiscernibles.

[5]Or, if the reader holds that haecceities are *essentially* trivializing, then we do not express a haecceity with 'is identical to *a*'.

[6]Rodriguez-Pereyra (Forthcoming) argues that some properties—such as *is not identical to a*—are trivializing despite being borne by distinct objects. I discuss these sorts of properties in greater depth below.

Inexpressible Ignorance

There are times when our ignorance is *inexpressible*. Although we recognize a sense in which we lack important information, we cannot express a proposition (in the relevant domain) whose truth-value we do not know. What inexpressible ignorance consists of is hotly contested; the easiest way to grasp the phenomenon is via examples.

Newtonian Location

Perhaps the clearest example of inexpressible ignorance concerns our location in Newtonian space. On the Newtonian view, space is a Euclidian substance: one that extends infinitely in all directions, persists throughout time and is independent of the objects located within it. Even if the world were devoid of material things, *something* would still exist: space itself.

It was once widely held that our position in Newtonian space is unknowable. After all, the universe as a whole could have been shifted four feet to the left from its actual location, and our phenomenal experiences would remain unchanged. In fact, for any Newtonian world N, there are infinitely many 'shifted' worlds that differ from N only in that the matter within these worlds is moved uniformly in some-direction-or-other. These worlds are genuinely distinct, but are *indiscernible* from one another. Because our phenomenal experiences are the same in each, we do not know which world we occupy—and so are ignorant of our Newtonian location.

Maudlin (1993) departed from the standard view. He defended two claims—only one of which concerns us at present. He argued that the ostensible ignorance of our Newtonian location is inexpressible, and that all ignorance is expressible. He concluded that we are not ignorant of our Newtonian location. Let us bracket the second claim and focus solely on the first: the claim that we cannot express ignorance of our Newtonian location.

Maudlin maintains that ignorance is expressible just in case there is a sentence that satisfies two conditions: its content describes the state of affairs that we are ignorant of, and we do not know its truth-value. That is, if we are to express our ignorance about whether some state of affairs obtains, there must be some sentence S describing that state of affairs such that 'We do not know whether S' is true. For example, we can express our ignorance of the primality of stars, since 'We do not know whether the number of stars is prime' is true.

In contrast, we cannot express ignorance of our Newtonian location; we know the truth-value of every sentence that describes what our location is. As Maudlin says,

> "Various positional states of the universe as a whole are possible: it could be created so my desk is *here*, or three meters north of here, or 888 meters from here in the direction from Earth to Betelgeuse, and so on. Which is the *actual* state of the world? Now the answer is easy: in its actual state, my desk is here, not three meters north or anywhere else." (Maudlin, 1993, pg. 90)

The basic thought is this: there are number of ways to refer to our Newtonian location. One involves indexicals—like the word 'here.' But we arguably know all of the pertinent indexical truths. I know that 'I am located here in Newtonian space' is true, and that 'I am located four feet to the left of here' is false. After all, there is no doubt that I am located *right here*—on *this very spot*—and nowhere else.

We need not refer to our Newtonian location with overt indexicals. I could, for example, dub the region of space I currently occupy 'region *r*,' and thereafter refer to my location with a proper name. But, once again, I seem to know every pertinent truth. I know that 'I am located at *r*' is true, and that 'I am located four feet to the left of *r*' is false. In fact, there seems to be *no* sentence that describes my Newtonian location whose truth-value I do not know. There is no 'natural origin' of spacetime I can use to identify particular points. Rather, I refer to points via the relation I stand in to them. But I know all of the facts about my relative location—so I cannot express my location in a way that reflects my ignorance.

Of course, this does not *prove* that ignorance of our Newtonian location is inexpressible. But, minimally, in constitutes an invitation to express it: to find some assertoric sentence *S* that describes our location whose truth-value we do not know. To date, no one has identified such a sentence. And so, a new orthodoxy has replaced the old. It is now widely held that we cannot express ignorance of our Newtonian location.[7] Those (like Maudlin) who also hold that all ignorance is expressible thus maintain that we are not ignorant of our location in Newtonian space.

Quidditism

Another example of inexpressible ignorance concerns *quidditism*—the view that properties are not defined by their causal roles.[8] According to the quidditist, the property that plays the mass-role is not defined by the causal effects of mass. Either that property is a primitive (and thus has no definition at all) or else is defined in terms of something that makes no mention of causation. In principle, quidditism allows for properties to 'swap' their causal roles; the property that had played the mass-role could switch places with the property that had played the charge-role (or any other).[9] It is thus a contingent fact that properties play the causal roles that they actually play. Arguments for quidditism are contentious—and engaging with them would take us far afield.[10] For our purposes, the important point

[7]See Pooley (Forthcoming) for a list of adherents to this view. Horwich (1978) draws upon similar reasoning to conclude that there is no fact of the matter regarding our position in Newtonian space.

[8]The term 'quidditism' was first coined by Armstrong (1989).

[9]Lewis (2009) sometimes describes quidditism as the higher-order analog of haecceitism—the view that there is primitive *thisness* for objects. On this conception, there is also a primitive *thisness* for properties—and so we can conceive of a property with one *thisness* exchanging places with one with another.

[10]Lewis (2009), for example, argues that quidditism best aligns with our ordinary modal judgments, while Bird (2007) considers (but ultimately rejects) a regress argument. A bit roughly, if properties are to be identified with their causal roles, then a change in an object's property amounts to a change in its ability to cause changes

is this: *if quidditism is true*, then properties give rise to inexpressible ignorance.

The quidditist holds that we lack important information. There are possible worlds that differ from one another only in that the properties have swapped their causal roles across them. In one world, the property that actually plays the mass-role has swapped places with the property that actually plays the charge-role; in another, it has swapped places with the property that actually plays the weak-nuclear-force-role, etc. These worlds are genuinely distinct, yet are *indiscernible* from one another. Our phenomenal experiences are the same across these worlds—as we only interact with properties via the causal roles that they play. We merely observe that some-property-or-other plays each causal role, and have no way to identify which. Because we do not know which of these worlds we occupy, we do not know which property plays which causal role.

This is an ignorance that we cannot express. While there are sentences that describe which properties play the various causal roles, we arguably know all of their truth-values. We know that 'The property that actually plays the mass-role plays the mass-role' is true and that 'The property that actually plays the mass-role plays the charge-role' is false. Of course, we need not refer to properties by overtly describing their causal roles. We could dub the property that actually plays the mass-role '*m*' and the property that actually plays the charge-role '*c*'—and thereby refer to them with proper names. Once again, however, we know the truth-values of the relevant sentences: that '*m* plays the mass-role' is true and that '*c* plays the mass-role' is false. In fact, there seems to be *no* sentence that describes which property plays each causal role whose truth-value we do not know. As Lewis aptly said, "We cannot answer the question: which property occupies that role? But worse: not only can we not answer that question, we can't even ask it." (Lewis, 2009, pg.15-16).

The problem is that we interact with properties causally. When we refer to a particular property, we do so via the causal role it plays (whether that reference occurs directly—by describing the causal role—or indirectly—by dubbing a property that plays a causal role and thereafter providing a name). But we know perfectly well that the property that plays causal role *r* plays causal role *r*. So, while it seems we are ignorant of which property plays each causal role, it is an ignorance that we cannot express.

Prime Matter

A third example of inexpressible ignorance concerns a branch of scholastic Aristotelianism—in particular, the view that substances are composites of prime matter and essential form, and these substances exhibit various accidental properties. It may be, for example, that Socrates consists of prime matter (essentially) formed as a human, and accidentally bears the property *is the teacher of Plato*.

in other properties. But changes in these properties purely consist in the ability to change other properties, etc. It thus seems that denying quidditism allows for an infinite series of changes in dispositions without actually *doing* anything.

On this view, there is a sense in which we are ignorant; we do not know which bit of prime matter composes each substance. There are possible worlds that differ from one another only in that the various bits of prime matter have been swapped around. In one world, the matter that actually composes Socrates has swapped places with the matter that actually composes Plato; in another, it has swapped places with the matter that actually composes Aristotle, etc. These worlds are genuinely distinct, yet are *indiscernible* from one another. We only observe that some-bit-or-other composes Socrates—and have no independent way of discerning which. Because we do not know which world we occupy, we are ignorant of which bit of matter composes each substance.

Once again, this is an ignorance that we cannot express. As Dasgupta says:

> "One can only refer to the underlying individual (or substance, or bit of prime matter) by demonstration ('*this* one') or by describing its relation to qualities ('the one that underlies this constellation of qualities'). And once we formulate a sentence *S* about which individual underlies the qualities in these terms, there is no problem determining whether it is true." (Dasgupta, 2015, pg. 450)

While there are various ways to refer to bits of prime matter, none reflect our ignorance. I could refer to some as 'that bit that actually composes Socrates,' gesture to Socrates and say '*that* prime matter,' or else dub it 'prime matter *p*.' But under each mode of reference, I know that that prime matter composes Socrates; i.e., I know that 'The bit that actually composes Socrates composes Socrates,' '*That* bit composes Socrates' and 'Prime matter *p* composes Socrates' are all true. Without another way to refer to that bit of prime matter, my ignorance is inexpressible.

I hope that the phenomenon of inexpressible ignorance is sufficiently clear. It arises when our referential resources are impoverished: when we have some way (or ways) of referring to a state of affairs, but none that reflects our ignorance of whether that state obtains.[11]

Some may still deny that ignorance is ever inexpressible—or minimally hold that the jury is still out. After all, each of these examples depend upon a controversial (and, in at least one case, demonstrably false) philosophical or physical theory. Perhaps they only motivate the following conditionals:

[11]There is an ongoing debate over how we ought to analyze inexpressible ignorance: i.e., over what the phenomenon ultimately consists of. Propositionalists (like Langton (2004) and Schaffer (2005)) hold that there are propositions that we cannot express—and we are ignorant of these propositions' truth-values. Nonpropositionalists (like Maudlin (1993) and Dasgupta (2015)) deny propositional ignorance. While Maudlin claims that we are ignorant of nothing at all in these cases, Dasgupta argues that inexpressible ignorance alienates us from the world—where this alienation consists in neither being acquainted with a thing nor knowing its full essence. I do not hope to settle the debate between propositionalists and nonpropositionalists here. For our purposes, the ability to identify the phenomenon itself is what matters—not how to analyze it.

1. *If Newtonian mechanics is true,* then we cannot express ignorance of our location.
2. *If quidditism is true*, then we cannot express ignorance of which property plays which causal role.
3. *If scholastic Aristotelianism is true,* then we cannot express ignorance of which bit of prime matter composes each substance.

If the theories that these examples depend upon are false, then the three conditionals may be true—despite the fact that all ignorance can be expressed.

I think we should not reject inexpressible ignorance prematurely. Minimally, the examples serve to shift the burden of proof onto those who deny the phenomenon. And while there are ample reasons to reject Newtonian mechanics, it is not entirely clear that we ought to do so because it generates inexpressible ignorance. Rather, it ought to be rejected because of its empirical inadequacy. But let us set that point aside. For our purposes, this objection misses its mark, for I only seek to defend a conditional:

4. *If there are indiscernible objects*, then their haecceities are inexpressible.

Those who endorse conditionals 1-3 thus accept all that I require, regardless of their attitude toward their antecedents.[12]

There are important connections between the PII and inexpressible ignorance: connections that philosophers have generally—perhaps even entirely—overlooked. I start with this: *every case of indiscernibility generates inexpressible ignorance*.

Consider a Black-world: one with two (or more) indiscernible objects. In this world, there seems to be something we are ignorant of: which object is which. After all, there is another world that is the same in all respects except that the indiscernible objects have swapped places—one where *this* object is located where *that* object actually is, and where *that* object is located where *this* object actually is. These worlds are phenomenally indistinguishable from one another (after all, the objects that have swapped places appear exactly the same, in virtue of their indiscernibility), so we do not know which world we occupy.

Yet again, this is an ignorance that we cannot express. While we can refer to an object indexically, we arguably know all of the indexical truths.[13] We know that *this* object is located where *this* object is—and that *that* object is not there in its place. And while we

[12]Conditionals 1-4 are all given as material conditionals—but they can easily be swapped for counterfactual conditionals if the reader prefers.

[13]As I argue later, it may actually be that we *cannot* refer to one of these objects indexically. By gesturing to one object—but not the other—we introduce a difference between those objects. One bears the property *was gestured to by me*, while the other does not. This complication does not change the fact that we cannot express our ignorance of which object is which.

can refer to objects with proper names rather than indexicals—i.e., we can dub one of the indiscernible objects '*a*'—these names do not allow us to express our ignorance. I know that 'Object *a* is located where *a* is located' is true and that 'Object *b* is located where *a* is located' is false.

There is no way to refer to the objects *except* via indexicals (whether that indexical is overt, or used to introduce a proper name). Because the objects are indiscernible, there exists no other discriminating feature: one that would allow us to refer to one object but not the other. (We cannot refer to a unique object with 'the object that is *F*' since both objects are *F* if either is). Because we can only refer to these individual objects via indexicals—and because we know all of the indexical truths—we cannot express our ignorance of which object is which. Given that indiscernibility always generates inexpressible ignorance, there is thus a novel argument for the PII: those who maintain that all ignorance is expressible ought to hold that the Identity of Indiscernibles is true. For our purposes, however, the important point is this: indiscernibility impoverishes our referential resources—and this gives rise to an ignorance that we cannot express.

Inexpressible Haecceities

This discussion indicates that indiscernibility generates inexpressible ignorance—but it does not establish that the haecceities of indiscernibles are inexpressible. For all that has been said, it could be that our ignorance of which object is which cannot be expressed—but their haecceities can. A further argument is needed.

At its core, the argument is simple. There is something in virtue of which a name refers to the object it does: something that makes it the case that 'Socrates' refers to Socrates and that 'Aristotle' refers to Aristotle. Philosophers have debated what feature this is since the birth of the analytic tradition—what property an object must have for a name to refer to it, rather than to something else. All that matters, for our purposes, is that there is some sort of 'semantic stickiness' or other: one that accounts for the reference of a name. In cases of indiscernibility, the semantic stickiness of one object is the same as the semantic stickiness of another (after all, if two objects are the same in all respects, then they are just as sticky as one another). When we attempt to name one of these objects, one of two things occurs: either the name ambiguously refers to both objects (in which case both objects can be understood as bearing the haecceity, despite being distinct) or else the name refers to neither (in which case we cannot use the name to express a haeccity that they bear). Neither alternative involves expressing a trivializing property—so we cannot express haecceities that trivialize the PII.

Let us consider this argument in some detail. Consider a world with indiscernible objects. For the sake of concreteness, take a world with two indiscernible spheres. (For our purposes, it would make no difference if there were two spheres or twenty, but let us restrict our attention to the simple case of two). And suppose that we attempt to name one of the spheres '*a*' in order to express the haecceity 'is identical to *a*.' There are two possibilities: either '*a*' ambiguously refers to both spheres or else it unambiguously refers to only one.[14] Let us take these in turn.

Suppose '*a*' Refers Ambiguously

A word is said to be ambiguous if it has distinct meanings. 'Bat' is classic example—as it refers both to a species of mammal and to an artifact used in baseball. A name, in particular, is said to be ambiguous just in case it has multiple referents. 'Mr. Smith' is such a name—as there are many people named Mr. Smith.[15] In supposing that '*a*' refers ambiguously, we thus suppose that each of the qualitatively identical spheres are referents of that name.

The claim that '*a*' is ambiguous is not altogether implausible. If there had there been only one sphere, presumably there would be no obstacle to naming it. We could either gesture to the sphere and dub it '*a*'—or else describe it in terms of a feature it has and everything else lacks. If there is no obstacle in naming a sphere when only one exists, it's hard to see what would prevent us from naming it when there are two. But if we *do* name one of the spheres when both exist, it seems the name will refer to the other as well. After all, whatever property made one sphere the referent of '*a*' is a property had by the other. And if having this property suffices for the name to refer to one, surely it suffices for the name to refer to both. For this reason, we might reasonably think that the name is ambiguous.[16]

If the name '*a*' is ambiguous, then 'is identical to *a*' denotes one of two things. The compound phrase may itself be ambiguous; it may refer to two properties—one that each sphere bears respectively. If this is so, then the phrase does not trivialize the PII. After all, each sphere bears a property denoted by 'is identical to *a*' despite being distinct. Alternatively, 'is identical to *a*' may refer to a single property that both objects bear. With the phrase 'is identical to *a*' we thus fail to describe a property borne by a single object;

[14]As it will emerge, what I think actually occurs in unambiguous cases is that '*a*' refers to neither sphere—but this is made evident by considering what would be required for '*a*' to refer to only one. Perhaps some think that there is yet another alternative: that the reference of '*a*' is *indeterminate*. This is related to—but strictly distinct from—the suggestion that '*a*' is ambiguous. In cases of ambiguity, a name simultaneously refers to two (or more) objects. In cases of indeterminacy, there is no fact of the matter about what the name refers to. This possibility also does not allow us to express trivializing haecceities—as there is no fact of the matter about what property 'is identical to *a*' refers to.

[15]There are numerous types of ambiguity in natural language. The examples of 'bat' and 'Mr. Smith' both involve lexical ambiguity, which occurs when distinct words are co-spelled. Lexical ambiguity is often contrasted with syntactic ambiguity, which occurs when a sentence in natural language is associated with distinct logical forms. 'Everyone loves someone' is such a sentence—as this can either be read as asserting that everyone loves someone-or-other, or else that there is a single person beloved by everyone.

[16]Although this paragraph consists of a defense of ambiguity, it should be clear from this arguments' overall structure that nothing turns on this. Those who deny ambiguous reference are free to take the second horn of this dilemma.

rather, we describe a property borne by two. And because 'is identical to *a*' describes an attribute of distinct objects, it does not trivialize the PII.

But what about logic? At the outset, we *proved* that objects bearing the same haecceity are identical. Proof comes with an air of finality. How could this situation be one where objects are distinct, while the haecceity they bear is the same?

Classical logic presupposes that terms are not ambiguous. Without this assumption, inferences that appear valid lead from truth to falsity. Suppose, for example, that the name '*b*' is ambiguous—and that one referent is *F* while the other is *G*. Classical logic allows us to infer $\exists x(Fx \wedge Gx)$ from *Fb* and *Gb*. Because *b* is *F* and *b* is *G*, there exists something that is both *F* and *G*. But, if '*b*' is ambiguous, this may not be so. Rather than one object that is both *F* and *G*, there may be distinct objects, each of which is a referent of '*b*', and one of which is *F* while the other is *G*. We cannot rely upon classical proof in a language with ambiguous names.

The same problem arises when we attempt to prove that objects bearing 'is identical to *a*' are one and the same. The proof presupposes that '*a*' is unambiguous—and so tells us nothing about a language with ambiguity. If the name is ambiguous, then there is a sense in which each object bears *is identical to a* and a sense in which each bears *is not identical to a*. Each object is identical to itself (and is one of the referents of '*a*')—and each object is distinct from the other (which is also a referent of '*a*'). While it appears that we have arrived at a contradiction (since each object both is, and is not, identical to *a*) this appearance is only due to the ambiguity of the name. For this reason, we cannot establish that '*is identical to a*' trivializes the PII in this sort of language. So, *if 'a' is ambiguous*, we do not express a trivializing property with 'is identical to *a*'.

Suppose '*a*' Refers Unambiguously

Suppose that '*a*' refers unambiguously. This suggestion is also not implausible. As we have already seen, accepting ambiguity forces the rejection of classical logic. This is unwelcome—especially since many classical inferences about the indiscernible spheres seem unobjectionable. If a sphere is grey and it is round, then it is both grey and round. And if a sphere has a mass of 5 kg, then it either has a mass of 5 kg or 10 kg. It is difficult to justify these inferences if classical logic fails when reasoning about the spheres.[17] In

[17]Arguably, the claim 'classical logic fails' is somewhat misleading. Strictly, classical logic is a formal language—one independent of the natural languages we carry out ordinary inferences in. Some might suggest that there are translations between natural languages like English and First-Order Logic—so that 'classical logic fails' amounts to the claim that when we translate ordinary thought into classical logic, the resulting classical inferences fail. But, as Williamson (2003) argues, there may be no perfect translations between natural language and logic. Regardless of whether there are any such perfect translations, there seem to be close correlates: $a \wedge b$ strongly resembles '*a* and *b*' (for example). So perhaps we should understand 'classical logic fails' as claiming that there are close connections between classical inferences and natural thought—and these inferences fail in cases of ambiguity.

order to accommodate these ordinary inferences, perhaps we ought to hold that '*a*' is unambiguous; it refers to one sphere and not the other.

If '*a*' refers unambiguously, then something disambiguates its reference; there is something in virtue of which the name refers to the object that it does. That is, there must exist some feature—or property—of *a* that makes '*a*' refer to *that* object, rather than some other.[18] Regardless of what this feature is, it *must* be a feature that *a* has that all other objects lack. After all, if *two* objects had this feature, then the feature would not secure a determinate reference of the name—it would at best ambiguously refer to those two objects. But we are considering the possibility in which the reference of '*a*' is *not* ambiguous.

So, if '*a*' has a determinate referent, then the object it refers to must have *some feature or other* that no other object has—one that secures the referent of the term. Even prior to the expression of a haecceity, there must exist a difference between objects in order to name them. But because sphere *a* has a feature that no other object has, it is not indiscernible from all other objects.[19] And because the object isn't indiscernible, the PII says nothing about this case. To put it another way, if there were a *true* case of indiscernibility, it would be impossible to name only one of the indiscernible objects—as there would be no attribute of that object which could secure a unique referent of a name. And without the ability to refer to individual indiscernibles, we cannot express haecceities that trivialize the PII.

But what about primitive reference?[20] There was a step in the argument that may seem innocuous, but that can be resisted: the step from "*a*' refers unambiguously' to 'There exists something that disambiguates the reference of '*a*'.' But perhaps there is *nothing at all* in virtue of which the name refers to the object that it does. It simply....refers. If there is nothing in virtue of which a name refers to the object that it does, then perhaps we can name one of the indiscernible objects without thereby presupposing a difference between that object and all others. By providing a sphere a name, we thus do not presuppose that it has some feature that all others lack.

For what it's worth, I think that this is the most promising path to resisting my argument. But there are costs to accepting primitive reference: ones that are worth emphasizing.

First, note that primitive reference not only impacts the PII—*but eliminates inexpressible ignorance entirely*. With primitive reference in our toolkit, it would be possible to refer to particular spacetime points without appeal to the relation we stand in to them—to particular properties without appeal to their causal effects—and to particular bits of prime matter without appeal to the substances that they compose. And if we *could* primitively

[18]Perhaps some suspect that a feature of an object *other* than '*a*' disambiguates its reference—something like the intension of the speaker of '*a*.' But this can easily be restructured to be a property of *a* itself: *a* may bear the property *is the intended referent of the speaker*.

[19]Perhaps some suspect that the reference can be secured relationally; perhaps I can gesture to one of the objects (and not the other) and in so doing dub that object '*a*.' But after this act, only one of the objects bears the property 'is gestured to by me'—which is a difference in properties used to secure a determinate reference.

[20]My thanks to Michael Della Rocca for pressing me on this point.

refer to a particular region of spacetime with 'r' without appeal to the relation we stand in to it, then we could state 'I do not know whether I am located at r,' and thereby express ignorance of our Newtonian location (and similarly so for primitive reference to particular properties or bits of prime matter). But no one has ever suggested primitive reference in these cases—so there is no reason to endorse it here either. To put it another way, any philosopher who finds primitive reference appealing in the discussion of the PII ought to apply it across the board. It is notable that no one has suggested doing so for these inexpressible ignorance.

Second, primitive reference *completely abandons the prospects for a metasemantics of proper names*. Perhaps the earliest puzzle in the analytic tradition concerned the meanings of names. Frege (1892) noted that 'Hesperus is Hesperus' differs in significance from 'Hesperus is Phosphorus.' While the first sentence is knowable *a priori* (say those who countenance *a priori* knowledge) and is a mere instance of the universal truth that all objects are self-identical, the second is only knowable *a posteriori*—and its apparent logical form ('$a = b$') has many false instances. To resolve this puzzle, Frege suggested that names have a *sense*—or way in which they refer. The sense of 'Hesperus' differs from the sense of 'Phosphorus,' which accounts for the difference in significance between the two identity claims. Descriptivists that follow this tradition identify proper names with disguised definite descriptions.[21] Perhaps the meaning of 'Hesperus' is 'the object appearing in *this* position in the evening sky' and perhaps the meaning of 'Phosphorus' is 'the object appearing in *that* position in the morning sky.'

If descriptivism were correct, primitive reference would be a nonstarter. On this view, the name of (an indiscernible) sphere is a disguised definite description of that sphere: something along the lines of 'the grey sphere of such-and-such mass.' But if the sphere truly is indiscernible from another, then both spheres satisfy any description we might provide. Because every description satisfied by one sphere is also satisfied by the other, there is no way to provide a name of one sphere without thereby naming the other. Moreover, the reference of the sphere's name would *not* be primitive—it would hold because a sphere satisfies a particular description.

Descriptivism has largely fallen out of favor in recent decades—but the problem generalizes to every theory of names' meanings. For example, intentionalists (roughly) hold that a name refers to an object in virtue of being the object that the speaker intends to refer to.[22] If two spheres are indiscernible, while it may be possible for a speaker to intend to refer to only one sphere, there is nothing which allows a speaker to intend to refer to a particular one, rather than the other.[23] One sphere thus bears the property *is the intended referent of speaker x* just in case the other does, so there is no prospect, on this metasemantics, for a name to refer to one sphere but not the other. If a speaker were to intend to a

[21]See, canonically, Russell (1905) .

[22]For discussions and defenses of intentionalism, see Stine (1978); Bach (1992). This sort of view is also suggested by Donnellan (1968).

[23]This point is analogous to one made by Putnam (1981).

particular sphere (and not the other), they would thereby introduce a difference between them: akin to painting a red line on one sphere and referring to 'the sphere with the red line.'

Direct reference arguably comes closest to primitive reference.[24] On this view, names directly refer to the objects they denote (rather than doing so via speakers' intensions or descriptions that the objects satisfy). 'Socrates' directly refers to Socrates and 'Aristotle' directly refers to Aristotle. Perhaps the name '*a*' takes one sphere, but not the other, as its referent—and so refers to only one of the spheres. We need not identify the meaning of '*a*' with a property borne by one sphere and not the other. The very fact that the spheres are distinct allows different names to denote different spheres.

This suggestion confuses semantics with metasemantics. Direct reference theorists typically endorse a metasemantic view that explains why a name refers to the object that it does. A common view is causal. When a name is first introduced, it refers to a particular object because the object is dubbed that name. Thereafter, the reference of later uses causally depends on a chain going back to the original dubbing. And it is in virtue of the first utterance that the later uses refer to the object that they do. For example, 'Gottlob' may first have referred to Frege because his parents dubbed him so—and later uses of that name refer to him because they stand in the appropriate causal relation to the initial utterance. On the causal view, it is (once again) impossible to name only one of the indiscernible spheres. Because the spheres are indiscernible, one of the spheres is dubbed a name if and only if the other is—and the spheres stand in precisely the same causal relations as one another. On this metasemantics, it is impossible for a name to denote only one of the spheres.

The upshot is this: if we are to endorse primitive reference (in the sense that a name denotes an object—but there is no property of that object in virtue of which it is the referent of that name) then there no prospect for a metasemantics of proper names. The aim of this metasemantics is to uncover what it is in virtue of that a name refers to the object that it does. Primitive reference amounts to the denial that there is any such feature—so it rejects all such accounts.

In any case, it's not entirely clear that accepting primitive reference will allow us to refer to only one of the indiscernible spheres. Even if it were the case that '*a*' referred to one sphere (but not the other) for no reason at all, there would still be a difference between the spheres. In this situation, one would bear *is the referent of 'a'*, while the other does not. Because the objects differ with respect to their properties, they are not truly indiscernible.

Primitive reference thus offers a path to expressing the haecceities of indiscernible objects—but not a particularly promising one. It has undesirable implications (and may not even succeed).

[24]Direct reference theory traces back to Mill (1867). More recent direct-reference theorists include Marcus (1961); Geach (1969); Donnellan (1970); Kripke (1970).

Generalizing Inexpressibility

So far, I have argued that the haecceities of indiscernible objects are inexpressible. But establishing this does not guarantee that no trivializing property of an indiscernible object can be expressed. While haecceities are the most paradigmatic example of trivialization, it is widely recognize that other properties trivialize the PII as well. Consider the conjunctive property *is red and identical to a*. Any object which bears this property is identical to *a*—and so any two objects that bear the conjunctive property are identical to one another. Or consider the property *is a member of {a}*. All objects bearing this property are identical, as only object *a* is an element of the singleton. We are just scratching the surface; why maintain that none of an indiscernible object's trivializing properties can be expressed?

One strategy parallels Maudlin (1993)'s discussion of inexpressible ignorance. While we have not proven that every trivializing property is inexpressible, this argument constitutes an invitation to express them: to find some property that trivializes the PII not susceptible to the previous argument.[25] Neither *is red and identical to a* nor *is a member of {a}* fit the bill. If it is impossible to express a haecceity, then it is presumably impossible to express a conjunctive property that has that haecceity as a conjunct. After all, expressing the conjunctive property presumably involves expressing its conjuncts—one of which cannot be expressed. And if the reason we cannot express *is identical to a* is that it is impossible to (uniquely) name object *a*, this rules out expressing *is a member of {a}* as well. If there were indiscernible objects, then the singletons these objects belong to would also be indiscernible from one another—and so they cannot be named for reasons already discussed. Until—and unless—we identify an expressible trivializing property, it is reasonable to hold that they are all inexpressible.

There is an independent reasons to suspect indiscernible objects' trivializing properties are all inexpressible. The trivializing properties we have considered so far can only borne by a single object (this, after all, is what guarantees that all objects bearing these properties are identical). But expressing these properties involves referring to their sole bearer—something needed to specify that it is a property of *that* object, rather than some other. In expressing these properties, we thus refer to the object that bears them. But if one object is indiscernible from another, then it is impossible to refer to it without thereby referring to another. For this reason, we cannot express the trivializing properties of indiscernibles.

An Objection—And Reply

Recently, Rodriguez-Pereyra (Forthcoming) has argued that some trivializing properties are borne by multiple objects. If there is any hopeof expressing trivialization, it lies here. If a property *F* is borne by multiple objects, then expressing *F* need not involve referring to its sole bearer (after all, there is no sole bearer to refer to). And if we can trivialize

[25] This resembles Maudlin's argument in that, while he did not prove that ignorance of our Newtonian location is inexpressible, there was an invitation to express it.

the PII without referring to individual indiscernible objects, then perhaps we can express this trivialization. It is valuable to consider these examples in depth. The upshot will be this: I think that there is something fundamentally correct about the examples Rodriguez-Pereyra provides—but there is also room to push back. This pushback is of independent interest, so I will discuss it in some detail. However, regardless of whether we embrace these examples, they do not allow us to express properties that trivialize the PII—as they too cannot be expressed.

One example are properties of difference—those like *is distinct from a* (or, *is not identical to a*). A great many objects bear the same property of difference; Plato and Aristotle both bear *is distinct from Socrates* (as does nearly everything else). But while distinct objects can (and do) bear some of the same properties of difference, none bear all of the same properties of difference. Objects that bear *is distinct from a*, *is distinct from b* etc. for all objects—save one—are identical to one another. Because objects that share all of the same properties of difference are identical, these properties trivialize the PII.

This example serves multiple purposes. Most directly, it undermines the claim that a property F is trivializing just in case it is borne by a unique object. Some other account of trivialization is required. It also suggests a novel path to trivialization. If a property F is trivializing because it can only be borne by a unique object, then the property $\neg F$ must be borne by all objects except one. We can use these negative properties to back into a trivial version of the PII. In effect, they serve to 'rule out' individual objects as their bearers; when all objects but one have been ruled out, the remaining objects must be identical, and the PII is trivially true.

There is room to resist the claim that properties of difference are trivializing—at least by themselves. A natural requirement on trivialization is that it is necessitating; properties that trivialize the PII necessitate the truth of the PII. Properties of difference do not satisfy this requirement. Suppose that a world w contains n objects. Within w, object a bears *is distinct from b, is distinct from c, ... is distinct from n*. It is perfectly true that every object within w that bears all of these properties is identical to a. Consider, however, a world w' that differs from w in that there is an additional object m. In w' there are two distinct objects that bear all of the properties of difference that a bears in world w. That is, both a and m bear *is distinct from b, is distinct from c, ... is distinct from n*. And so, there is a possible world where distinct objects bear all of the properties of difference that a bears within w. For this reason, these properties do not necessitate the truth of the PII. If properties must necessitate the PII in order to trivialize the PII, then these properties of difference are not trivializing.

There are several ways to correct for this. We might, firstly, deny that trivialization requires necessitation: perhaps properties trivialize the PII despite not necessitating its truth. Alternatively, we might countenance properties of difference concerning objects that do not exist. Perhaps even in world w, object a bears *is distinct from m*—despite the fact that m does not exist within w. If this is so, then a and m do not bear all of the same properties of difference in w' that a bears in w; while a bears *is distinct from m*, m does not. Or we might hold that there is a constant domain of objects across possible worlds (*a la* Williamson (2013)). If every object exists in all possible worlds, then there is no pair of worlds such that the latter has more objects than the former. Yet another (and my preferred) resolution is to hold that there are 'totality properties': those like *being such that a, b, c, ... , n* are all of the objects that exist. On this view, properties of distinctness do not trivialize the PII themselves—but they do in conjunction with the appropriate totality property. That is, necessarily, any objects which bear the properties *is distinct from b*, *is distinct from c*, ... , *is distinct from n*, *is such that a, b, c, ... , n are all of the objects that exist* are identical to *a*.

Another example of shared trivializing properties that Rodriguez-Pereyra discusses are disjunctive properties where one disjunct is a haecceity.[26] For example, take the property *is red or identical to a*. A great many objects bear this property—but only one can both bear it and bear *is not red or identical to a*. Collectively, the two are borne by *a* and *a* alone. Because it is impossible for distinct objects to share this pair of properties, they trivialize the PII.

Here too there is room for pushback. Throughout this book, Rodriguez-Pereyra treats trivialization as a monadic, second-order property. It is borne by haecceities—but not *is red* or *is round*. However, these examples suggest that it is more natural to treat trivialization as a *polyadic*, second-order property—rather than a monadic one. That is to say, we may hold that *collections* of properties trivialize the PII—rather than claiming that only individual properties do. Perhaps properties of distinctness (with or without a totality property) *collectively* trivialize the PII, but that no individual property of distinctness does—and perhaps *being red or identical to a* and *being red or not identical to a* jointly trivialize the Principle, but that neither does by itself. On this conception, haecceities are merely the special case of trivializing collections that consist of a single property.

There are (at least) two ways to recover a monadic conception of trivialization from a polyadic one. We could, first, define a conjunctive property for each collection of trivializing properties. That is, if a collection *FF* of properties jointly trivializes the PII, we might claim that the property of bearing all of the *FF*s is trivializing. Using the resources of λ-abstraction, we could represent the trivializing disjunctive property as $\lambda x.((Red(x) \vee x = a) \wedge (\neg Red(x) \vee x = a))$—and we could represent the trivializing property of difference as $\lambda x.(x \neq b \wedge x \neq c ... \wedge x \neq n)$ (perhaps adding a totality conjunct if needed). If this is so, then properties of distinctness and disjunctive properties aren't actually trivializing—but they figure as conjunctive parts of properties that are.

Another way to recover monadic trivialization is to define it in terms of belonging to a trivializing collection (I suspect, though am not sure, that this is Rodriguez-Pereyra's preferred method). Of course, we cannot hold that every property that belongs to a trivializing collection trivializes the PII. For an arbitrary set of properties S that collectively trivialize the PII, $S \bigcup \{F\}$ is also trivializing (for any property F). So, if we held that every

[26] He also argues that disjunctive properties where a disjunct is a property of difference are trivializing.

property that belongs to a trivializing collection is itself trivializing, then all properties would be trivializing.

However, we *can* identify trivializing properties with those that belong to *minimal* collections of trivializing properties. That is, if the members of a set of properties *S* jointly trivialize the PII—and there is no proper subset of *S* whose members jointly trivialize the PII—then each of the members of *S* is trivializing. On this view, properties of distinctness and disjunctions of trivializing properties are indeed trivializing: but what makes them trivialize is that they are members of minimal collections of properties that jointly trivialize.

These objections aside, there are two potential examples of trivializing properties that can be borne by multiple objects: properties of difference, and disjunctive properties where one disjunct is a haecceity. We need not take a stand on whether these examples are convincing, for *these properties of indiscernibles also cannot be expressed*.

If there were two indiscernible spheres, it would be possible to express many properties of difference that they bear. There would be no obstacle to saying '*being distinct from the Eiffel Tower*' or '*being distinct from Mars*.' However, in order to express a trivial version of the PII, it must be possible to state all of their properties of difference—and there is one where problems arise: their distinctness from one another. In order to state '*is distinct from sphere a*,' it must be possible to refer to one sphere but not the other (after all, while sphere *b* is distinct from sphere *a*, it is not distinct from itself). But for the reasons already belabored, this is impossible in cases of indiscernibility. We also cannot express disjunctive properties where one disjunct is a haecceity. Just as expressing a conjunctive property involves expressing its conjuncts, so too expressing a disjunctive property involves expressing its disjuncts. Because it is impossible to express haecceities of indiscernibles, so too it is impossible to express disjunctions of their haecceitistic properties. For this reason, *even if we accept Rodriguez-Pereyra's examples*, we cannot express the trivializing properties that result.

Conclusion

Metaphysicians have long debated whether the PII is true. The discussion provided by Whitehead and Russell (1952) suggested that it is, if only trivially so. I have here argued that this is is not the case: if there is a triviality, it is one that we cannot express. One question has loomed over this discussion: if we do not express a triviality with 'Objects bearing all the same properties are identical,' what is it that we do we express? Nothing in this argument turns on the answer—but, for my part, I suspect that we express something substantive: a claim that may or may not be the case, but one that turns on whether the PII is actually—substantively—true.

References

Adams, Robert. 1979. "Primitive Thisness and Primitive Identity." *Journal of Philosophy* 76(1):5–26.

Armstrong, David. 1989. "The Nature of Possibility." *Canadian Journal of Philosophy* 16(4):575–94.

Bach, Kent. 1992. "Intentions and Demonstrations." *Analysis* 52(3):140–6.

Bird, Alexander. 2007. "The Regress of Pure Powers?" *The Philosophical Quarterly* 57(239):513–34.

Black, Max. 1952. "The Identity of Indiscernibles." *Mind* 61(242):153–64.

Casullo, Albert. 1984. "The Contingent Identity of Particulars and Universals." *Mind* 93(372):527–41.

Dasgupta, Shamik. 2015. "Inexpressible Ignorance." *Philosophical Review* 124(4).

Della Rocca, Michael. 2005. "Two Spheres, Twenty Spheres and the Identity of Indiscernibles." *Pacific Philosophical Quarterly* 86(4):480–92.

Donnellan, Keith. 1968. "Putting Humpty Dumpty Back Together Again." *The Philosophical Review* 77(2):203–15.

Donnellan, Keith. 1970. "Proper Names and Identifying Descriptions." *Synthese* 21(3-4):335–58.

Frege, Gottlob. 1892. "On Sense and Reference." *Zeitschrift für Philosophie und Philosophische Kritik* 100:25–50.

French, Steven. 1989. "Why the Identity of Indiscernibles is not Contingently True Either." *Synthese* 78:141–66.

Geach, Peter. 1969. "The Perils of Pauline." *The Review of Metaphysics* 23(2):287–300.

Goodman, Jeremy. Forthcoming. "Consequences of the Conditional Excluded Middle." *https://jeremy-goodman.com/ConsequencesCEM.pdf* .

Horwich, Paul. 1978. "On the Existence of Time, Space, and Space-Time." *Nous* 12(4):397–419.

Hoy, Ronald. 1984. "Inquiry, Intrinsic Properties and the Identity of Indiscernibles." *Synthese* 61(3).

Katz, Bernard. 1983. "The Identity of Indiscernibles Revisited." *Philosophical Studies* 44(1):37–44.

Kripke, Saul. 1970. *Naming and Necessity*. Harvard University Press.

Langton, Rae. 2004. "Elusive Knowledge of Things in Themselves." *Australasian Journal of Philosophy* 82(1):129–36.

Leibniz, Gottfried Wilhelm. 1991. Discourse on Metaphysics. In *Discourse on Metaphysics and Other Essays*, ed. Translated by Daniel Garbed and Roger Ariew. Hackett Publishing Company pp. 1–40.

Lewis, David. 2009. Ramseyan Humility. In *Conceptual Analysis and Philosophical Naturalism*, ed. David Braddon-Mitchell and Robert Nola. MIT Press pp. 203–22.

Marcus, Ruth Barcan. 1961. "Modalities and Intensional Languages." *Synthese* 13(4):303–22.

Maudlin, Tim. 1993. "Buckets of Water and Waves of Space: Why Space-Time is Probably a Substance." *Philosophy of Science* 60:183–203.

Mill, John Stuart. 1867. *A System of Logic*. Logmans.

O'Connor, D. J. 1954. "The Identity of Indiscernibles." *Analysis* 14(5):103–110.

Pooley, Oliver. Forthcoming. *Substantivalism and Haecceitism*. Unpublish Manuscript.

Putnam, Hilary. 1981. *Reasons, Truth and History*. Cambridge University Press.

Rodriguez-Pereyra, Gonzalo. Forthcoming. *Two Arguments for the Identity of Indiscernibles*. Oxford University Press.

Russell, Bertrand. 1905. "On Denoting." *Mind* 14(6):479–93.

Schaffer, Jonathan. 2005. "Quiddistic Knowledge." *Philosophical Studies* 123:1–32.

Stine, Gail. 1978. "Meaning Other Than What We Say and Referring." *Philosophical Studies* 33(4):319–37.

Whitehead, Alfred and Bertrand Russell. 1952. *Principia Mathematica: Volume 1 (second edition)*. Cambridge University Press.

Williamson, Timothy. 2003. "Everything." *Philosophical Perspectives* 17(1):415–65.

Williamson, Timothy. 2013. *Modal Logic as Metaphysics*. Oxford University Press.

Wörner, David. 2021. "On Making a Difference: Toward a Minimally Non-Trivial Version of the Identity of Indiscernibles." *Philosophical Studies* 178:4261–78.