

HAECCEITIES AND THE TRIVIALITY OF THE IDENTITY OF INDISCERNIBLES

Samuel Z. Elgin¹

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 that bear 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 trivialize the Identity of Indiscernibles 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. Whether the Identity of Indiscernibles is trivially true ultimately turns on whether we can quantify over properties 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 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 containing 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. To the best of my knowledge, no one holds that this is how the world actually is—but it seems perfectly clear that this situation is metaphysically possible. No contradiction arises from spheres existing in such a way. Moreover, 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.

¹I would like to thank Andrew Bacon, Michael Della Rocca, Catherine Elgin, William Ratoff, and Gonzalo Rodriguez-Pereyra for their helpful feedback on earlier drafts of this paper.

²Here, ‘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.

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 acknowledging 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 all 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 ‘ $\lambda x.(x \text{ is red})$ ’ and the property *is identical with 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 if *a* and *b* bear all of the same properties, then they are identical, as follows:

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

The choice of ‘*a*’ and ‘*b*’ is arbitrary; it is possible to construct such a proof for any objects. Haecceities thus ensure that the PII is a simple matter of logic: a triviality 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

³The 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 (2022) (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.

⁴I won’t go into much depth about the details of this system here. Aside from the following derivation—and a brief discussion in the conclusion—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.

would be. A research program has thus emerged that aims identify all of the trivializing properties—simply in order to exclude them from substantive interpretations.⁵

I deny that haecceities trivialize the PII—proof notwithstanding. This denial is not restricted to indiscernible objects' haecceities: I similarly deny that any of their properties are trivializing. More precisely, I deny that their trivializing properties *can be expressed*. From this, it immediately follows that whatever it is philosophers *have* expressed are not properties that trivialize the PII.

The argument for inexpressibility occupies the entirety of this paper, but a bare-bones version of it is this. Consider a world with two indiscernible objects, and suppose that 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.⁶ The phrase 'is identical to *a*' is itself ambiguous; it simultaneously denotes two properties—the property of being identical to one of the indiscernible objects, and the property of being identical to the other. What we express with 'is identical to *a*' is thus compatible with the objects being distinct: it may be that each bears a different haecceity (each of which is a referent of 'is identical to *a*'). The former proof presupposes that '*a*' is unambiguous, and so says nothing about cases of ambiguity.

If '*a*' refers unambiguously, then something disambiguates its reference; there is something in virtue of which '*a*' refers to the object that it does. What this feature is depends upon 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. To put it another way, if there were an instance of indiscernibility, it would be impossible to refer to one of the indiscernible objects without thereby referring to all of them. And without the ability to refer to individual indiscernibles, we cannot express haecceities that trivialize the PII.

This paper lies at the intersection of metasemantics and metaphysics. Most directly, it is an argument concerning the language used to discuss the Identity of Indiscernibles (rather than the PII itself). It emerges that metaphysicians have not expressed properties they have taken themselves to express; that which was held to be trivial is, in fact, substantive. Moreover, those who hold particular views about the relation between language and world will be pressured to accept metaphysical implications. Most notably, if all properties are expressible (and if, as I argue, no trivializing property is expressible) then no properties trivialize the PII.⁷

⁵See, most notably, Rodriguez-Pereyra (2022)—but see, also, Katz (1983); Rodriguez-Pereyra (2006).

⁶If the reader holds that haecceities are essentially trivializing (so that any property which does not trivialize the PII does not count as a haecceity), then we do not express a haecceity with 'is identical to *a*'.

⁷I myself take no stand on whether there are inexpressible properties. Conversationally, some have

The structure is as follows. I begin with a slight digression: by surveying the literature on inexpressible ignorance. There are several reasons to begin with this discussion. It introduces the phenomenon of inexpressible states of affairs—which is analogous to the inexpressibility of indiscernible objects' haecceities. Inexpressibility is neither as novel nor as suspect as it might have seemed. But the two literatures are not only similar, but directly related. I will argue that indiscernibility generates inexpressible ignorance; we cannot express our ignorance of which of the indiscernible objects is which. This literature also serves as an argumentative foil. Potential objections to my view apply equally well to inexpressible ignorance. Those who endorse inexpressible ignorance ought not be moved by these objections. In the following section, I argue that indiscernible objects' haecceities are inexpressible—before arguing that this inexpressibility generalizes to every trivializing property that they bear. I close by discussing the implications this has for the Identity of Indiscernibles. Ultimately, its triviality turns on whether we can quantify over properties that we cannot express.

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.

expressed skepticism about the existence of inexpressible propositions—and if all propositions are expressible, it is natural to hold that all properties are as well.

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. 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, it 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.⁸ Those (like

⁸See 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.

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.⁹ 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).¹⁰ 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.¹¹ For our purposes, the important point 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 independent 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 some sentences state that a given property plays a particular causal role, 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,

⁹The term ‘quidditism’ was first coined by Armstrong (1989).

¹⁰Lewis (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.

¹¹Lewis (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 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.

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 that 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*.

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.¹²

Inexpressibility and Indiscernibility

There are important connections between the PII and inexpressible ignorance: connections that philosophers have generally—perhaps even entirely—overlooked. I start with this: *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.¹³ 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.¹⁴ We know that *this* object is located where *this* object is—and that *that* object is not there in its place. And while we can refer to objects with proper names rather than indexicals—i.e., we can dub one of the

¹²There 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.

¹³This particular claim—that there is a distinct possible world where the objects have swapped places—depends on controversial (meta)physical assumptions. In particular, if spacetime were purely relational and the world consisted of two indiscernible atomic spheres, there may be no distinct world where they switch places (since the relational facts in one situation are the same as they are in the other). But although this view holds that there is no distinct possibility where the two spheres swap places, I hope that we can still recognize a sense in which we are ignorant of which object is which. My thanks to Gonzalo Rodriguez-Pereyra for pressing me on this point.

¹⁴As 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.

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. Indiscernibility impoverishes our referential resources—and this gives rise to an ignorance that we cannot express.

Inexpressible Haecceities

Thus far, I have argued that indiscernibility generates inexpressible ignorance; but this does not establish that the haecceities of indiscernible objects are inexpressible. For all that I have said, it could be that our ignorance of which object is which cannot be expressed—but their haecceities can. Some further argument is needed.

The core 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 long debated what this feature is—what property an object need have for a name to refer to it, rather than to something else. For our purposes, what matters 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 precisely 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 what we express is compatible with the objects bearing distinct haecceities—ones that happen to be denoted by the same phrase) or else the name refers to neither (in which case we cannot use the name to express a haecceity that they bear). Neither alternative involves expressing a trivializing property—so we cannot express haecceities that trivialize the PII.

It is valuable to examine this argument in some detail. Consider a world with indiscernible objects. For the sake of concreteness, take a world with two indiscernible spheres. (For the purpose of this argument, 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.¹⁵ Let us take these in turn.

¹⁵As it will emerge, what I think actually occurs in unambiguous cases is that '*a*' refers to neither sphere—

Suppose 'a' Refers Ambiguously

A word is said to be ambiguous if it has multiple meanings. 'Bat' is classic example—as it refers both to a species of mammal and to a stick 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.¹⁶ 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 been only one sphere, presumably nothing would prevent us from giving it a name. 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 is hard to see what prevents us from naming it when there are two. But if we *do* name one sphere when both exist, it seems the name must refer to the other sphere as well. Because the spheres are indiscernible, whatever property made one sphere a referent of 'a' is a property borne by the other. And if bearing this property ensures that the one sphere is a referent of 'a,' then the fact that both bear it ensures that both are referents of that name.¹⁷

If the name 'a' is ambiguous, then 'is identical to a' is presumably ambiguous as well. It refers both to the property of being identical with one referent of 'a'—and to the property of being identical with the other. What we express with 'is identical to a' is thus compatible with the existence of distinct spheres that bear distinct haecceities (ones that both happen to be referents of the same phrase). And because our expression is compatible with the spheres being distinct, we do not express anything that trivializes the PII.

But what about logic? At the outset, we *proved* that if *a* and *b* bear all of the same properties, then they are identical. Proof comes with an air of finality. How could this situation be one where objects are distinct?

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

but this is made evident by considering what would be required for 'a' to refer to only one. Quite clearly, if 'is identical to a' expresses a haecceity for neither object, then it does not denote one of their trivializing haecceities. 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.

¹⁶There are several 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.

¹⁷Although 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, and suppose that 'a' refers unambiguously.

'*c*' is ambiguous—and that one referent bears property *F* while the other bears property *G*. Classical logic allows us to infer $\exists x(Fx \wedge Gx)$ from *Fc* and *Gc*; because *c* is *F* and *c* is *G*, there exists something that is both *F* and *G*. But, if '*c*' 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 '*c*', 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 occurs when we attempt to prove that objects bearing 'is identical to *a*' are one and the same. The proof presupposes that '*a*' and '*b*' are unambiguous. This assumption is at work in line *iv*—which infers $\lambda x.(x = a)(b)$ from $\lambda x.(x = a)(a) \leftrightarrow \lambda x.(x = a)(b)$ and $\lambda x.(x = a)(a)$. This inference is valid only if the referents of '*a*' and '*b*' are held constant in each of these lines. But, if the terms are ambiguous, their referents need not be constant; '*a*' might refer to one sphere in one line and the other sphere in the next. So, if terms are ambiguous, then we cannot conclude that *b* bears *is identical to a*—in much the way that we cannot conclude that an object is both *F* and *G* from *Fc* and *Gc*. We thus cannot establish that '*is identical to a*' trivializes the PII in this sort of language. 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 prevents classical reasoning. This is unwelcome; many classical inferences 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.¹⁸ In order to accommodate 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 to another.¹⁹ Regardless of what this feature is, it *must* be a feature that *a* has and all other objects lack. After all, if *two* objects had this feature, then the feature would not secure a

¹⁸Arguably, the claim 'classical logic fails' is 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 (if imperfect translations) between classical inferences and natural thought—and these inferences fail in cases of ambiguity.

¹⁹Perhaps 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*.

determinate referent of 'a'; it would, at best, ambiguously refer to those two objects. But we are considering the possibility that 'a' is unambiguous.

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 name. 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.²⁰ 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 only one of the indiscernible objects—as there would be no attribute of that object to secure a unique referent of its name. And without the ability to refer to individual indiscernibles, we cannot express haecceities that trivialize the PII.²¹

*But what about primitive reference?*²² 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'." 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 is 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 refer to a particular region of spacetime with '*r*,' 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). It is notable that no one has suggested this in any debate over inexpressible ignorance. To put it another way, any philosopher who appeals to primitive reference for the PII ought to apply it across the board. Those who accept any cases of inexpressible ignorance should not be moved

²⁰Perhaps 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.

²¹This argument is neutral on what the correct metasemantics is; it merely presupposes that there is some-correct-metasemantics or other. But it also has implications for naming discernible objects once we subscribe to a particular account. That is, if property *F* determines the referent of a name, it will be impossible to provide distinct names for multiple objects that are *F*—even if those vary with respect to property *G* (and so are not indiscernible from one another).

²²My thanks to Michael Della Rocca for pressing me on this point.

by this objection.

Second, primitive reference *completely abandons the prospects for a metasemantics of proper names*. The aim of this metasemantics is to uncover what determines the reference of a name: to accept primitive reference is to reject every account.²³ Even direct-reference theorists (who hold that names refer directly to the objects they denote) often reject primitive reference.²⁴ Such philosophers typically supplement direct reference theory with a causal metasemantics of reference. That is, while they deny that a name means anything other than the object it refers to, they account for *which* object it refers to causally. When a name is first introduced, it refers to a particular object because the object is dubbed that name. The reference of later uses depends on a causal chain of reference; the reference of subsequent uses of that name is determined by how they relate causally to prior uses. This metasemantic account (like all metasemantic accounts) is incompatible with primitive reference; it holds that reference is determined causally, while primitivism holds that nothing determines reference. If primitive reference is true, then every metasemantics of proper names is false. It is thus a heavy price to pay.

Third, it is 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—and the PII once again says nothing about this case.

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).

Generalizing Inexpressibility

I have argued that indiscernible objects' haecceities are inexpressible. In itself, this does not establish that none of their trivializing properties can be expressed. While haecceities are the most well-known trivializing properties, it is widely recognized 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

²³I will not go in to inordinate depth about all of the available metasemantic accounts; the literature is too rich to do justice to here. Descriptivists (like Frege (1892) and Russell (1905)) hold that a name refers to a particular object because it (uniquely) satisfies a particular description. Intentionalists (like Stine (1978) and Bach (1992)) hold that a name refers to an object in virtue of being the object that the speaker intends to refer to. There are many views beyond these two.

²⁴Examples of direct-reference theorists include Marcus (1961); Geach (1969); Donnellan (1970); Kripke (1970).

singleton. We are just scratching the surface; any number of properties might trivialize the PII. Why think that none of them can be expressed?

One strategy parallels Maudlin (1993)'s discussion of inexpressible ignorance. While we have not proven that every trivializing property is inexpressible, we have extended 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 our inability to (uniquely) denote *a* prevents us from expressing *is identical to a*, then we presumably cannot express *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 reason to deny that any trivializing property can be expressed. The trivializing properties we have considered so far can only be 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 of some other. In expressing these properties, we thus refer to the object that bears them. But if two of (or more) objects are indiscernible, then it is impossible to refer to one without thereby referring to another. For this reason, we cannot express any trivializing property of indiscernible objects.

An Objection—And Reply

Recently, Rodriguez-Pereyra (2022) has argued that some trivializing properties are borne by multiple objects. If there is any hope of expressing trivialization, it lies here. After all, if a property *F* is borne by multiple objects, then expressing *F* need not involve referring to its sole bearer (there being no sole bearer to refer to). The previous argument says nothing about these properties; it was an argument to the effect that our inability to refer to indiscernible objects prevents us from expressing properties that only they bear. And if we *can* trivialize the PII without referring to particular indiscernible objects, then perhaps we can express this trivialization. It is valuable to discuss these properties in depth. The upshot will be this: while I think that there is something fundamentally correct about the examples Rodriguez-Pereyra provides, 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.

²⁵This 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.

One such property is an antihacceity—a property like *is distinct from a* (or, *is not identical to a*).²⁶ A great many objects bear the same antihacceity; 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 antihacceities, none bear all of the same antihacceities. All 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 antihacceities are identical, these properties trivialize the PII.

This example serves multiple purposes. Most directly, it undermines the claim that a property is trivializing just in case it is borne by a single 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 single 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. They serve to rule out individual objects as their bearers; when we consider sufficiently many such properties, all objects but one have been ruled out. The remaining objects must be identical, and the PII is trivially true.

I think that there is room to resist the claim that antihacceities are trivializing—at least by themselves. A natural requirement for trivialization is that it is necessitating; properties that trivialize the PII necessitate the truth of the PII. There are a few reasons why we might endorse this requirement. The claim ‘ F trivializes the PII’ seems synonymous with (or, at least, very close to synonymous with) with ‘ F trivializes the truth of the PII’—or perhaps ‘ F trivializes that the PII is true.’ It is challenging to see how F could trivialize the truth of the PII if it is compatible with its falsity. It is difficult to even understand what ‘trivialization’ means if properties trivialize the PII’s truth while being compatible with its falsity. So, if a property F allows for the possibility that the PII is false, then it does not trivialize its truth.

Antihacceities do not satisfy this requirement; they allow for the possibility that the PII is false. Suppose that a world w contains m objects. Within w , object a bears *is distinct from b*, *is distinct from c*, ... *is distinct from m*. And, within w , any object that bears all of these antihacceities is identical to a . But consider a world w' that differs from w in that there is an additional object n . In w' there are two distinct objects that bear all of the antihacceities that a bears in world w . That is, both a and n bear *is distinct from b*, *is distinct from c*, ... *is distinct from m*. And so, there is a possible world where distinct objects bear all of the antihacceities that a bears within w . For this reason, the antihacceities that a actually bears do not necessitate the truth of the PII; they allow for the possibility that the PII is false (there being a possible world where distinct objects share all of these antihacceities). If properties must necessitate the PII in order to trivialize its truth, then antihacceities 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 antihacceities concerning objects that do not

²⁶We might also refer to these properties as ‘properties of difference.’

exist. Perhaps even in world w , object a bears *is distinct from n* —despite the fact that n does not exist within w .²⁷ If this is so, then a and n do not bear all of the same antihacceities in w' that a bears in w ; while a bears *is distinct from n* , n 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 correction is to hold that there are ‘totality properties’: those like *being such that a, b, c, \dots, m are all of the objects that exist*. On this view, antihacceities do not trivialize the PII themselves—but they do in conjunction with the appropriate totality property. Necessarily, any objects which bear the properties *is distinct from b , is distinct from $c, \dots, is distinct from m , is such that a, b, c, \dots, m are all of the objects that exist$* are identical to a .

Another novel example of trivialization concern disjunctive properties—where one disjunct is a hacceity.²⁸ 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 hacceities—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 antihacceities (with or without a totality property) *collectively* trivialize the PII, but no individual antihacceity does—and perhaps *being red or identical to a* and *being red or not identical to a* jointly trivialize the Principle, but neither does by itself. On this conception, hacceities 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 antihacceity as $\lambda x.(x \neq b \wedge x \neq c \dots \wedge x \neq m)$ (perhaps adding a totality conjunct if needed). If this is so, then antihacceities 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

²⁷Note that this suggestion involves a modification—if not a complete rejection—of classical logic. From $a \neq b$, classical logic licenses the inference to $\exists x(a \neq x)$. But if b does not exist, presumably it does not witness this existential claim.

²⁸Rodriguez-Pereyra also argues that disjunctive properties where a disjunct is an antihacceity are trivializing.

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 of properties itself trivializes the PII. For an arbitrary set of properties S trivializes the PII, $S \cup \{F\}$ is also trivializing (for any property F). So, if we held that every 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, antihacceities 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 and modifications aside, there are two potential examples of trivializing properties that can be borne by multiple objects: antihacceities, 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 antihacceities 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 antihacceities—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

It is impossible to refer to individual indiscernible objects—as indiscernibility prevents reference to one without reference to all. This inexpressibility generalizes; we cannot express any trivializing property that indiscernible objects bear. But the fact that we cannot express any trivializing property does not guarantee that the PII itself is nontrivial. While we cannot state what *makes* the Principle a triviality, a triviality it may be.²⁹

I suspect that many who have accepted this much of the argument will be tempted to claim that the PII is substantive (or, at least, claim that we express something substantive

²⁹My thanks to Andrew Bacon for pressing me on this point.

when we state the PII). If it impossible to state that which makes the PII a triviality, why not hold that we make a substantive assertion when stating the Identity of Indiscernibles?

However, it is worth noting that it is possible to prove the triviality of the Identity of Indiscernibles without reference to any trivializing property (or, indeed, directly to indiscernible objects)—at least given certain background logical assumptions. In particular, if we accept a logic that quantifies over properties that we cannot express—and objects that we cannot refer to—then we can prove that the Identity of Indiscernibles is true. Whether the PII is substantive thus turns on whether we can quantify over properties and objects that we cannot express.

The proof given at the outset of this paper relied upon naming indiscernible objects: something we now recognize cannot be done. However, the following principles allow us to prove the PII without referring to individual indiscernibles:³⁰

Classicality	$\vdash p$ if p is a theorem of propositional logic.
β -Equivalence	$\vdash p \leftrightarrow q$ if q is the β -reduction of p
Self-Identity	$\vdash a = a$ for all constants and variables
Gen	If $\vdash p$ then $\vdash \forall x p$
U-Elim	$\vdash \forall x.F(x) \rightarrow Fa$

Classicality asserts that theorems of propositional logic are theorems of this system; it is a theorem that $p \rightarrow (p \vee q)$ and that $(p \wedge q) \rightarrow p$. β -equivalence allows us to perform the operation of β -reduction: $\lambda x.Tall(x)(Sarah)$ entails $Tall(Sarah)$: a bit roughly, the claim that *Sarah* is in the extension of *is tall* entails that *Sarah is tall*.³¹ Self-identity asserts that everything is self identical. Note that its instances include not only constants, but also variables; $x = x$ is a theorem on this system, despite the fact that it is not truth-evaluable due to the variable ‘ x ’ that lacks an assignment. Gen, for its part, allows us to generalize the results of theorems. For example, because $x = x$ is a theorem on this system, so is $\forall x.(x = x)$. U-Elim allows us to infer particular claims from universal ones. It asserts that if all objects bear a property F then a particular a is F . As with self-identity, U-Elim strictly allows us to make inferences not only involving constants, but also involving variables. So, for example, we may infer $F(x)$ from $\forall x.F(x)$. With these assumptions to hand, the

³⁰Many of the λ abstracts and parentheses are suppressed in these formalisms. Additionally, many of these should not be interpreted as principles, but rather as schemata with applications in many types. For example, Self-Identity is not only taken to entail that $a = a$ but also that $F = F$. The types themselves are also omitted, as they are contextually evident.

³¹There has recently been sustained discussion about β -reduction in higher-order logic; it is valuable to state what this principle precisely claims. It is highly controversial whether β -equivalent terms denote the same proposition—whether $\lambda x.Tall(x)(Sarah)$ expresses the same proposition as $Tall(Sarah)$. While Dorr (2016) argues that they do, Rosen (2010) and Fine (2012) endorse principles that entail they do not. We need not take a stand on this debate; as stated, this principle only claims that the two propositions have the same truth-value—not that they are identical. This is entirely uncontroversial (at least in contexts that lack opaque predicates like ‘believes’).

proof proceeds as follows:

<i>i.</i>	$\forall \lambda X.(Xx \leftrightarrow Xy)$	Supposition
<i>ii.</i>	$\forall \lambda X.(Xx \leftrightarrow Xy) \rightarrow (\lambda z.(z = x)(x) \leftrightarrow \lambda z.(z = x)(y))$	U-Elim
<i>iii.</i>	$\lambda z.(z = x)(x) \leftrightarrow x = x$	β -Equivalence
<i>iv.</i>	$\lambda z.(z = x)(y) \leftrightarrow y = x$	β -Equivalence
<i>v.</i>	$x = x \leftrightarrow y = x$	<i>i-iv</i> and Classicality
<i>vi.</i>	$y = x$	<i>v</i> , Self-Identity and Classicality
<i>vii.</i>	$\forall \lambda X.(Xx \leftrightarrow Xy) \rightarrow x = y$	<i>i-vi</i> and Classicality
<i>viii.</i>	$\forall \lambda x, y.(\forall \lambda X.(Xx \leftrightarrow Xy) \rightarrow x = y)$	<i>vii</i> Gen

This concludes that, for all objects, if those objects bear all of the same properties, then they are identical. Notably, this proof never refers to individual indiscernible objects; indeed, there are no constants for *any* object. Nor is there reference to the haecceities of indiscernibles (or any other of their trivializing properties). In effect, this proof uses variables without assignment in order to bypass the need for constants that denote indiscernible objects. Rather than supposing that there are indiscernible objects, denoting them ‘*a*’ and ‘*b*’ and then proving that they are identical, this proof uses variables for objects and their haecceities, then generalizes to conclude that objects bearing all of the same properties are identical. While we cannot express any trivializing property, whether the PII is itself a triviality thus turns on whether the assumptions that underlie this proof are tenable.

For this to establish the triviality of the PII, its quantifiers must quantify over indiscernible objects. Line *viii* makes a claim about all objects; for it to make a claim about indiscernibles, those must fall within that quantifier’s scope. That is to say, while line *viii* asserts that—for all objects, if objects bear all of the same properties, then those objects are identical—in order for this to assert something *about indiscernible objects*, they must fall under the scope of ‘all objects.’ If they do not fall under this quantifier’s scope, this proof says nothing whatsoever about them—in much the way that ‘All men are mortal’ says nothing about John’s mortality if he does not fall under the scope of ‘All men.’ So, this proof only *actually* demonstrates that the PII is trivially true if we can quantify over that which we cannot express.³² There is something odd about suggesting that our quantifiers range over objects we cannot refer to—but the suggestion is not obviously untenable.³³

³²I note, in passing, that there are independent, paradoxical reasons to deny that any quantifier can range over all objects. Unrestricted quantification is controversial—but for a defense of unrestricted quantifiers, see Cartwright (1994); Williamson (2003).

³³Of course, in certain languages this happens all the time. I might quantify over all real numbers and assert that all of them are real numbers—without the ability to denote each real number. But this limitation is due only to our choice of language. We cannot denote certain reals given the finite limitations our language imposes. If we were to operate with a language with a greater cardinality (perhaps by introducing different symbols for every angle between two lines) it would be possible to denote every real. In contrast, the inability to express indiscernibles—or their trivializing haecceities—is not due to our choice of language. They cannot

With such a quantifier to hand, we could assert that there exists an object that is F without the ability to state which object it is. While unintuitive, this is the cost of the triviality of the Identity of Indiscernibles.

What I have directly argued is that none of an indiscernible objects' trivializing properties—be it a haecceity, antihacceity, or any other—can be expressed. In itself, this does not guarantee that the PII is nontrivial: for it may be possible to demonstrate its triviality without referring to any trivializing property. Whether or not this holds turns on whether we can quantify over properties that we cannot express.

be expressed in any language whatsoever.

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.
- Cartwright, Richard. 1994. "Speaking of Everything." *Noûs* 20:1–20.
- 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. 1970. "Proper Names and Identifying Descriptions." *Synthese* 21(3–4):335–58.
- Dorr, Cian. 2016. "To be F is to be G." *Philosophical Perspectives* 30(1):39–134.
- Fine, Kit. 2012. Guide to Ground. In *Metaphysical Grounding*, ed. Fabrice Correia and Benjamin Schneider. Cambridge University Press pp. 37–80.
- 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.
- O'Connor, Daniel. J. 1954. "The Identity of Indiscernibles." *Analysis* 14(5):103–110.
- Pooley, Oliver. Forthcoming. *Substantivalism and Haecceitism*. Unpublish Manuscript.
- Rodriguez-Pereyra. 2006. How not to Trivialize the Identity of Indiscernibles. In *Universals, Concepts and Qualities: New Essays on the Meanings of Predicats*, ed. P. F. Strawson and A. Chakrabati. Ashgate pp. 205–23.
- Rodriguez-Pereyra, Gonzalo. 2022. *Two Arguments for the Identity of Indiscernibles*. Oxford University Press.
- Rosen, Gideon. 2010. Metaphysical Dependence: Grounding and Reduction. In *Modality, Metaphysics, Logic and Epistemology*, ed. Bob Hale and Aviv Hoffman. 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* 17:415–65.

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.