

THE REVIEW OF SYMBOLIC LOGIC
Volume 0, Number 0, Month 2009

Identity and discernibility in philosophy and logic

JAMES LADYMAN

Department of Philosophy, University of Bristol

and

ØYSTEIN LINNEBO

Department of Philosophy, Birkbeck College, University of London

and

RICHARD PETTIGREW

Department of Philosophy, University of Bristol

Abstract. Questions about the relation between identity and discernibility are important both in philosophy and in model theory. We show how a philosophical question about identity and discernibility can be ‘factorized’ into a philosophical question about the adequacy of a formal language to the description of the world, and a mathematical question about discernibility in this language. We provide formal definitions of various notions of discernibility and offer a complete classification of their logical relations. Some new and surprising facts are proved; for instance, that weak discernibility corresponds to discernibility in a language with constants for every object, and that weak discernibility is the most discerning non-trivial discernibility relation.

§1. Introduction There has been much debate in philosophy about the relation between identity and distinctness on the one hand, and various forms of discernibility on the other. For instance, philosophers have debated the truth of the Principle of the Identity of Indiscernibles (PII), which is naturally formulated using a second-order quantifier ranging over some class of properties of particular philosophical significance:

$$\forall P(Px \leftrightarrow Py) \rightarrow x = y \quad (1)$$

However, questions about the relation between identity and different forms of discernibility can also be formulated in the branch of mathematical logic known as *model theory*. What is the relation between these two approaches to such questions? Can the tools of mathematical logic usefully be brought to bear on the philosophical questions?

This paper has two main aims. Firstly, we show how many of the philosophical questions about the relation between identity and discernibility can be ‘factorized’ into two components: a sharper philosophical question about the adequacy of a formal language to the description of the world, and a mathematically precise question about discernibility in this formal language. This shows how a mathematical study of the relation between identity and various forms of discernibility in formal languages can be of philosophical interest. Our second aim is to undertake this sort of mathematical study, building on work by Ketland (Ketland (2006), Ketland (2011)). We provide formal definitions of a variety of notions of discernibility and then break new ground by offering a complete classification of the logical relations

between the resulting forms of discernibility. Some other new and surprising facts are proved as well; for instance, that weak discernibility corresponds to discernibility in a language with constants for every object, and that weak discernibility is the most discerning non-trivial discernibility relation. By contrast, relative discernibility emerges as comparatively unimportant.

§2. Philosophical context The philosophical debate about the relation between identity and discernibility has often centred on PII. So we begin by clarifying this principle. What is the range of the second-order quantifier that is used in the above formulation of PII? If the range includes, for each object in the first-order domain, the property of being identical with that object, then PII comes out trivially true.¹ So attention focuses on whether PII is true when the second-order quantifier has a more restricted range.

The usual response is to restrict the second-order quantifier to qualitative properties. But it is hard to define ‘qualitative’. One might try to define a property as ‘qualitative’ just in case it is not ‘identity-involving’. For instance, *being on Earth* seems ‘identity-involving’ in a way that *being made of iron* does not. However, this notion of ‘identity-involving’ is ambiguous between involving the identity relation and involving some particular object. We shall return to this issue below, where we consider languages where these two notions of identity-involving come apart. For now, it will suffice to operate with the intuitive idea that qualitative properties are observable, either in an everyday sense or in the more idealized sense in which this notion is used in science.

Why be interested in PII restricted to qualitative properties? One reason is an interest in the history of philosophy. Leibniz endorsed a particularly strong version of the principle according to which no two objects can have all the same intrinsic qualitative properties. For instance, every snowflake differs from every other snowflake in its internal structure. Call two such objects *intrinsically discernible*. (A precise definition will be offered below.) However, objects that are not intrinsically discernible may yet satisfy a form of PII where the variable P ranges over extrinsic properties.

Secondly, many philosophers (perhaps including Leibniz) are attracted to PII because it promises to deliver the notion of individuality from metaphysical speculation and make it amenable to empirical verification. If every object has a unique set of qualitative properties, then facts about the identity and diversity of objects become fully accessible to empirical inquiry. If PII is true, we can ask the further question whether this truth is accidental or necessary. Is there any empirical motivation for the claim that PII is a necessary truth? For example, it may be extremely unlikely that two distinct snowflakes, or leaves, or rocks, should be exactly the same with regard to intrinsic qualitative properties. So empiricists may plausibly argue on inductive grounds for the universal generalisation that no two objects share all their intrinsic and qualitative properties. However, the empiricist lacks grounds for the further claim that this universal generalization holds of necessity. Nothing about the way new snowflakes form is causally related to facts about the

¹ In the philosophical literature, the property of being identical with a particular object is known as *the haecceity* of that object.

exact shape of snowflakes long ago and miles away. So although the violation of this version of PII may be highly improbable, it does not appear to be impossible. However, if the principle of impenetrability for classical particles is assumed to be lawlike, then arguably PII will be necessarily true when the range of its second-order quantifier includes spatio-temporal properties. Even so, perhaps the best way for an empiricist to defend the necessity of PII is to construe it as a methodological norm. The motivation for PII may be the empiricist sentiment that we ought to be able to have empirical access to individuation.²

A third reason why philosophers are interested in PII arises in connection with the metaphysical question: ‘What is it that makes an object what it is and not some other object?’. (An answer to this question is often called a ‘principle of individuation’.) Two traditional answers are given by the bundle theory and the idea of transcendent individuality. The former states that an object is individuated by its qualitative properties. Hence, it is often thought that the bundle theory requires PII.³ On the other hand, those who deny PII often use its failure or possible failure to argue that objects must be individuated by something that transcends their qualitative properties.

Readers may be suspicious of this talk of ‘individuation’ and what ‘grounds’ or ‘determines’ facts about identity and distinctness. We share these concerns. However, many of the questions about the relation between identity and diversity on the one hand, and qualitative facts on the other, admit of more deflationary readings. For instance, we can ask whether the identity and diversity relations in a given domain are definable in terms of, or are logically determined by, qualitative properties and relations. These are the sort of ‘sanitized’ questions that will be investigated here.

Recent discussions of PII have been particularly concerned with quantum mechanics and mathematical structuralism.⁴ In both cases, it has been claimed that the putative objects in question violate PII and hence that some kind of transcendent individuality is required. And in both cases, a weak form of discernibility has been invoked to deny that PII is violated after all. Most notably, Simon Saunders Saunders (2006) has defended the importance of a distinction, going back to Quine Quine (1976), between three grades of discernibility: absolute, relative, and weak. The basic ideas are simple (although formal definitions are deferred until Section §5.). We begin by adding a fourth grade: intrinsic.

Two objects are *intrinsically discernible* when there is an intrinsic property that one object has that the other lacks. Examples might include snowflakes if they are

² See, for instance, Saunders (2003).

³ Russell (1911) is an early source for this view. It is explicitly defended in (Armstrong, 1980, 91–97). More recently it is appealed to by E.J. Lowe (Lowe, 2003, 80), and Hawthorne and Sider (Hawthorne & Sider, 2006, 32) write, “The classic bundle theory is generally thought to preclude the possibility of distinct indiscernible particulars”. However, it is not clear that the bundle theory requires any form of PII, let alone a strong one.

⁴ See, for instance, Saunders (2006), Muller & Saunders (2008), Muller & Seevinck (2009), and Bigaj & Ladyman (2010) for discussions of PII in quantum mechanics. And see Burgess (1999), Keränen (2001), Ketland (2006), MacBride (2006), Shapiro (2006), Ladyman (2005), and Leitgeb & Ladyman (2008) for discussions of the principle in mathematical structuralism.

genuinely all distinguished by their internal spatial structure. Or people if each has a distinguishing intrinsic property.

Two objects are *absolutely discernible* when there is a property that one object has that the other lacks. Examples include everyday material objects, natural numbers, and classical particles in a sufficiently asymmetric universe. Note that absolute discernibility does not require intrinsic discernibility: for example, even if there were two snowflakes not discernible by means of their intrinsic properties, they may nonetheless be absolutely discernible by an extrinsic property, such as being more than five miles from each reindeer.

Next, two objects are *relatively discernible* when there is a two-place relation in which the first stands to the second but the second does not stand to the first. Examples include instants of time if time has an intrinsic direction, any two people in a queue, or more generally the elements of a set equipped with a linear order.

Finally, two objects are *weakly discernible* when there is a relation in which the first stands to the second but the first does not stand to itself. Examples famously include Max Black's two spheres Black (1952), which are weakly discerned by the relation ' x is two miles from y '; two fermions in the singlet state of spin, which are weakly discerned by ' x has opposite spin to y '; and the complex numbers i and $-i$, which are weakly discerned by ' $x + y = 0$ '.

We call objects that cannot even be weakly discerned *utterly indiscernible*.⁵ Bosons are often thought to be utterly indiscernible,⁶ and nodes in an edgeless graph provide another instance of utter indiscernibility Leitgeb & Ladyman (2008). If there are bare particulars, they will be utterly indiscernible.

It is controversial whether Quine's generalized forms of discernibility establish that, for example, fermions and i and $-i$ satisfy PII.⁷ Indeed, it is controversial whether they should be called forms of 'discernibility' at all. 'Discerning' has two broad connotations. The first is a monadic notion of picking something out as distinct. For instance, one may discern one's friend in a large crowd. The second is a dyadic notion of distinguishing or observing the difference between objects. For instance, one may discern a minute difference between two (so-called) identical twins. Only the second connotation is compatible with Quine's generalized forms of discernibility. However, whether or not weak discernibility is really 'discernibility' in the sense of PII, the terminology has stuck and we will use it here.

Discernibility relations can be compared as to how discerning they are. Say that one discernibility relation is *more discerning* than another just in case the following holds: there are situations in which two objects are discerned by the former but not by the latter; but in any situation, two objects that are discerned by the latter relation are also discerned by the former. In a different context, we say the same thing of the palates of two gourmands: one palate is more discerning than another if the first can tell apart any two flavours that the second can, while there are flavours that the first can tell apart that the second cannot. It is well

⁵ Ketland (2011) calls this notion 'strong indiscernibility'; we have heard others refer to it as 'weak indiscernibility'. We find both labels somewhat confusing and hence adopt ours, which is proposed by Caulton & Butterfield (ta). Ketland uses 'relative' and 'weak discernibility' in the same sense as we do. But where we (and other parties to the debate) say 'absolute discernibility', he says 'monadic discernibility'.

⁶ However, see Muller & Saunders (2008) for a dissenting view.

⁷ See, for instance, Hawley (2009), MacBride (2006), Bigaj & Ladyman (2010).

known that weak discernibility is more discerning than absolute discernibility: that is, in all situations, two objects that are absolutely discernible are also weakly discernible; but there are situations in which two objects are weakly discernible but not absolutely discernible.

In what follows we develop some logical tools that can be used to analyze the various discernibility relations that come up in the philosophical discussion. Our aim is thus similar to that of Ketland (2006), Caulton & Butterfield (ta), and Ketland (2011), although many of our results are (to the best of our knowledge) new. First, we provide a ‘factorization’ of the traditional philosophical questions about discernibility into a sharper philosophical question about the adequacy of a formal language to the description of the world and a mathematically precise question about discernibility in this formal language. Second, we articulate a precise sense in which weak discernibility is the most discerning nontrivial discernibility relation, thus for the first time providing an explanation of why this notion that has recently received so much attention is actually the most important relation to study. Third, we prove that weak discernibility by relations expressible in a purely qualitative language is capable of discerning surprisingly many pairs of objects because it is equivalent to weak (as well as to absolute) discernibility by relations definable in the much richer language that results from adding a constant for each object of the domain. Weak discernibility by relations definable in a purely qualitative language is thus equivalent to weak discernibility by relations definable in a language that helps itself to singular reference to each of the objects concerned. This highlights what is at stake in the debate about whether the philosophically important notion is weak rather than absolute discernibility. Fourth, we provide a complete classification of the logical relations between a variety of notions of discernibility, not just the four traditional ones mentioned above, but also ones arising from considering languages with constants and languages with an identity predicate. Finally, we establish a curious feature of relative discernibility, namely that it is the only discernibility relation considered here that is not the complement of an equivalence relation.

§3. Discernibility in a structure Questions about discernibility arise both in philosophy and in the branch of mathematical logic called *model theory*. Above, we sketched their origin in philosophy. Let us turn now to model theory, before asking how the two approaches relate to one another.

We will (largely) follow the terminology and notation of Hodges (1997). A *structure* consists of: (i) a *domain*, which is a set of objects; (ii) a collection of distinguished elements from the domain, which we call the *constant elements*; (iii) for each n , a collection of n -ary relations on the domain; (iv) for each n , a collection of n -ary functions on the domain. Each structure comes equipped with a *signature*, which contains a constant for each constant element, a relation symbol for each relation, and a function symbol for each function.

A structure A , together with its associated signature, uniquely determines four different sorts of language that will interest us in this paper.⁸ First, a first-order language without the identity symbol, which we denote \mathcal{L}_A . \mathcal{L}_A consists of the set of first-order formulae without identity whose constants, relation symbols, and

⁸ Of course, the mapping from structures to associated languages will be many-one.

function symbols are amongst those in the signature. Second, a first-order language with identity, which we denote $\mathcal{L}_{\bar{A}}$. The third and fourth languages are determined in the same way by a structure that is closely related to A . It is the structure obtained from A by letting every element of the domain of A be a constant element, and expanding the signature to include a constant for each of the new constant elements. We call this structure A^* and we denote the corresponding first-order languages with and without identity \mathcal{L}_{A^*} and $\mathcal{L}_{\bar{A}^*}$ respectively.

Next, satisfaction. If $\varphi(x_1, \dots, x_n)$ is a formula of \mathcal{L}_A with free variables x_1, \dots, x_n , and a_1, \dots, a_n are elements of the domain of A , we write $A \models \varphi(x_1, \dots, x_n)[a_1, \dots, a_n]$ to mean that $\varphi(x_1, \dots, x_n)$ is satisfied by the structure A when each free variable x_i is assigned a_i . Sometimes, when there is no risk of confusion, we instead write $A \models \varphi(a_1, \dots, a_n)$. And similarly for the languages $\mathcal{L}_{\bar{A}}$, \mathcal{L}_{A^*} , and $\mathcal{L}_{\bar{A}^*}$.

With this in hand, we can give model-theoretic definitions of the grades of discernibility introduced informally above. It is crucial to note that these grades of discernibility are defined relative to a structure A and a language \mathcal{L} , where \mathcal{L} is \mathcal{L}_A , $\mathcal{L}_{\bar{A}}$, \mathcal{L}_{A^*} , or $\mathcal{L}_{\bar{A}^*}$. It is also worth noting here that each of these languages contain only *first-order* formulae. Of course, it is possible to define the grades of discernibility relative to other languages, which perhaps permit higher-order quantification or infinite conjunctions and disjunction. We could then explore the consequences. But we leave this for another time.

We define the grades of discernibility in A relative to \mathcal{L}_A . The definitions of discernibility in A relative to $\mathcal{L}_{\bar{A}}$, and discernibility in A^* relative to \mathcal{L}_{A^*} and $\mathcal{L}_{\bar{A}^*}$, are analogous.

DEFINITION 3.1. (Grades of Discernibility) *Suppose A is a structure and a and b are elements of its domain. Then*

- (a) a and b are intrinsically discernible in A (written $\text{Int}_A(a, b)$) if there is $\varphi(x)$ in \mathcal{L}_A such that φ does not contain any quantifiers or any constants and $A \models \varphi(a)$ but $A \not\models \varphi(b)$.
- (b) a and b are absolutely discernible in A (written $\text{Abs}_A(a, b)$) if there is $\varphi(x)$ in \mathcal{L}_A such that $A \models \varphi(a)$ but $A \not\models \varphi(b)$.
- (c) a and b are relatively discernible in A (written $\text{Rel}_A(a, b)$) if there is $\varphi(x, y)$ in \mathcal{L}_A such that $A \models \varphi(a, b)$ but $A \not\models \varphi(b, a)$.
- (d) a and b are weakly discernible in A (written $\text{Weak}_A(a, b)$) if there is $\varphi(x, y)$ in \mathcal{L}_A such that $A \models \varphi(a, b)$ but $A \not\models \varphi(a, a)$.
- (e) a and b are distinct in A (written $\text{Dis}_A(a, b)$) if $A \models a \neq b$.

We write $\text{Int}_{\bar{A}}$ for intrinsic discernibility in A relative to $\mathcal{L}_{\bar{A}}$, and similarly for $\text{Abs}_{\bar{A}}$, $\text{Rel}_{\bar{A}}$ and $\text{Weak}_{\bar{A}}$. We write Int_{A^*} and $\text{Int}_{\bar{A}^*}$ for intrinsic discernibility in A^* relative to \mathcal{L}_{A^*} and $\mathcal{L}_{\bar{A}^*}$ respectively, and similarly for the other discernibility relations.

Some gloss for Definition 3.1. was provided in Section §2., along with examples. A property of an object is often said to be intrinsic to it if the existence and nature of other objects is counterfactually irrelevant to the object having the property. For this reason, we restrict to formulae without quantifiers in (a).⁹ It is plausible that Leibniz meant to restrict the second-order quantifier in PII to such intrinsic properties. And, so restricted, the principle seems to be true of

⁹ The idea is due to Caulton & Butterfield (ta).

people and everyday objects. However, when we consider classical particles of the same kind, we seem to have examples of distinct objects that are nonetheless not intrinsically discernible. This motivates the relation of absolute discernibility. And similarly, apparent counterexamples to PII when restricted to intrinsic and extrinsic properties motivate the relations of relative and weak discernibility.

In sum: given two objects a and b in a given structure A , we can ask whether they are discernible in A relative to either \mathcal{L}_A or \mathcal{L}_A^- , or in A^* relative to either \mathcal{L}_{A^*} or $\mathcal{L}_{A^*}^-$. In Section §7., we provide an exhaustive account of the relations between the different grades of discernibility relative to these different languages.

§4. Discernibility in philosophy and in model theory When investigating questions about discernibility amongst a given collection of objects, philosophers are typically interested in some restricted class C of properties of those objects and relations amongst them. For instance, empiricists are concerned with so-called qualitative properties, whereas metaphysicians are interested in some class of basic metaphysical properties. What is the relation between the questions about discernibility as they appear in philosophy and as they appear in model theory? Take the class of objects in question. Let it be the domain of a structure A . Then the relation between the notions of discernibility in model theory and in philosophy will depend on how the relations on the domain of A that can be defined by formulae in \mathcal{L} compare with the relations in C , where \mathcal{L} is \mathcal{L}_A , \mathcal{L}_A^- , \mathcal{L}_{A^*} , or $\mathcal{L}_{A^*}^-$. Let us say:

- \mathcal{L} is *expressively sound with respect to C* if every relation that is definable by a formula in \mathcal{L} is in C .
- \mathcal{L} is *expressively complete with respect to C* if every relation in C is definable by a formula in \mathcal{L} .

Examples of languages that are expressively unsound (with respect to appropriate C): If C does not include haecceities for any of the objects in question, then a language in which any can be defined is unsound with respect to C . Similarly, a language in which we can define properties or relations corresponding to purely mathematical predicates in the language of quantum mechanics might be thought unsound if there is a methodological stricture to use only physically meaningful properties, as in Muller & Saunders (2008).¹⁰

Examples of languages that are expressively incomplete (with respect to appropriate C): Any language that fails to define basic physical properties will typically be thought incomplete. Also, if each object does have a haecceity, then languages that cannot define these will be incomplete.

We say that a language \mathcal{L} is *expressively adequate* iff it is expressively sound and complete. This follows usage in model theory. Note that the notion of adequacy here means not satisfactory but rather optimal, perfect, or just right.

Clearly, philosophers disagree about what languages are expressively adequate. Most if not all empiricists will think that a language with predicates for haecceities is unsound, whereas some metaphysicians may think that languages without such

¹⁰ One small observation here: If the class C is not closed under the logical operations of conjunction, disjunction, negation, and quantification, then no language can hope to be expressively sound with respect to C .

predicates are incomplete.¹¹ Empiricists may consider a language that refers to all the qualitative properties of things as adequate, whereas others will insist on predicates for all metaphysically fundamental features of the world.

If the language \mathcal{L} is expressively sound with respect to C , then the logical notion of a particular grade of discernibility with respect to \mathcal{L} entails the philosophical notion of the corresponding grade of discernibility by properties and relations in C . If, on the other hand, \mathcal{L} is expressively complete with respect to C , then the logical notion of some grade of indiscernibility with respect to \mathcal{L} entails the philosophical notion of the corresponding grade of indiscernibility by properties and relations in C . Putting these two observations together, it follows that if the language \mathcal{L} is both expressively sound and complete, each logical notion of discernibility relative to \mathcal{L} is equivalent to the corresponding philosophical notion of discernibility by properties and relations in C .

In this paper we do not take a stand on the class C of properties and relations that may discern two objects. We leave this open and investigate instead various forms of discernibility in a structure relative to a language \mathcal{L} . Once a class C is specified, the question of the expressive soundness and completeness of \mathcal{L} can be addressed, and our results about discernibility relative to \mathcal{L} can be translated into results about discernibility with respect to C , as explained above. Our investigation is thus a way of giving precise mathematical content to the original philosophical questions, relative to a choice of the class C . In this way, the original philosophical question ‘factorizes’ into a philosophical question about the class C , and a mathematical question about discernibility relative to \mathcal{L} .

It is important to realize that model theory itself is neutral on the relation between identity and discernibility. A structure is just a mathematical object composed of a domain along with distinguished elements of the domain, distinguished functions and relations on it. Model theory makes no requirement whatsoever concerning the objects on which the structure is based. For nearly every structure considered in this article, the elements of its domain can be taken from amongst the natural numbers; and all of the natural numbers are absolutely discernible in the natural number structure and are thus compatible with the strictest form of PII. However, using these absolutely discernible ‘building blocks’, we can construct structures in which PII is violated, which will later enable a more informed assessment of whether situations corresponding to these structures are genuinely possible. Here it is essential to distinguish between the discernibility of the objects on which the structure is based (e.g. the natural numbers) and their discernibility in the structure. For instance, take the following structure: its domain is $\{1, 2\}$; it has no constant elements or functions; and it has a single binary relation $\{\langle 1, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 1 \rangle, \langle 2, 2 \rangle\}$. Then the ‘building blocks’ of the model—the numbers 1 and 2—are absolutely discernible in the natural number structure, but they are aren’t even weakly discernible in the structure described.

A comparison with Kripke semantics may be instructive. This too is a mathematical tool that is neutral on the relevant philosophical questions but of great help in articulating and examining these questions. Just as Kripke models provide an innocent way of articulating and comparing what is involved in various metaphys-

¹¹ See (Ketland, 2006, 313-4) for a technical discussion of the relation between haecceities and indiscernibility in a second-order language.

ical views about modality, our use of model theory provides an innocent way of articulating and comparing what is involved in various metaphysical views about the relation between identity and discernibility.

We are now in a position to appreciate a distinction, alluded to above, which is often overlooked. Say that a property is *identity-involving* if its proper analysis involves the identity relation. Say that a property is *object-involving* if its proper analysis makes appeal to particular objects. The identity-involving properties are those definable only in $\mathcal{L}_{\bar{A}}$ or $\mathcal{L}_{\bar{A}^*}$. The object-involving properties are those definable only in \mathcal{L}_{A^*} or $\mathcal{L}_{\bar{A}^*}$. Here are some examples.

	identity-involving	not identity-involving
object-involving	being Mercury	being on Mercury
not object-involving	being the innermost planet	having an iron core

We do not try to adjudicate here the thorny issue of which entry in this table corresponds to the correct explication of ‘qualitative’.

§5. The hierarchy of grades of discernibility We noted above that the philosophical discernibility relations can be compared by how discerning they are. So can their logical counterparts introduced in Section §3.. In this section, we are interested only in discernibility in A relative to \mathcal{L}_A .

We have seen how a logical notion of discernibility D can be applied to any structure A to yield a discernibility relation D_A on the domain of A . For instance, the notion of weak discernibility yields the relation Weak_A of weak discernibility in A relative to \mathcal{L}_A . If D and D' are two notions of discernibility, we write $D \Rightarrow D'$ to mean that, for every structure A and all elements a and b of A , if $D_A(a, b)$, then $D'_A(a, b)$. Thus, $D' \not\Rightarrow D$ means that there is a structure A with elements a and b such that $D'_A(a, b)$ holds, but not $D_A(a, b)$. We say that one notion of discernibility D' is *more discerning* than another D just in case $D \Rightarrow D'$ and $D' \not\Rightarrow D$. Then the following theorem is well-known from the literature:¹²

THEOREM 5.2.

- $\text{Int} \Rightarrow \text{Abs} \Rightarrow \text{Rel} \Rightarrow \text{Weak} \Rightarrow \text{Dis}$
- $\text{Dis} \not\Rightarrow \text{Weak} \not\Rightarrow \text{Rel} \not\Rightarrow \text{Abs} \not\Rightarrow \text{Int}$

Proof. First, we show that each is at most as discerning as the next. Suppose A is a structure and a and b are elements of its domain. Then:

- If $\text{Int}_A(a, b)$, then $\text{Abs}_A(a, b)$. Obvious.
- If $\text{Abs}_A(a, b)$, then $\text{Rel}_A(a, b)$. If $\varphi(x)$ absolutely discerns a and b , then $\psi(x, y) := \varphi(x) \wedge \neg\varphi(y)$ relatively discerns a and b .
- If $\text{Rel}_A(a, b)$, then $\text{Weak}_A(a, b)$. If $\varphi(x, y)$ relatively discerns a and b , then $\psi(x, y) := \varphi(x, y) \wedge \neg\varphi(y, x)$ weakly discerns a and b .
- If $\text{Weak}_A(a, b)$, then $\text{Dis}_A(a, b)$. The contrapositive follows immediately from Leibniz’s Law.

¹² See, for instance, Ketland (2011), Section 3.2.

10 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

Second, we show that each is strictly less discerning than the next.

- Let A be the *edgeless graph with two vertices a and b* :

$$a \quad b$$


Then $\text{Dis}_A(a, b)$, but not $\text{Weak}_A(a, b)$.

- Let A be the *dumbbell graph with vertices a and b* :

$$a \longleftrightarrow b$$


Then $\text{Weak}_A(a, b)$, but not $\text{Rel}_A(a, b)$. (This is the graph-theoretic analogue of Max Black's spheres or entangled fermions.)

- Let A be the *cyclic graph with three vertices a , b , and c* :



Then $\text{Rel}_A(a, b)$, but not $\text{Abs}_A(a, b)$.

- Let A be one of the *smallest non-trivial asymmetric undirected graphs with vertices a, \dots, f* :



Then $\text{Abs}_A(a, b)$, but not $\text{Int}_A(a, b)$. (Cf. Ladyman (2007).)

This completes our proof. □

The definitions of grades of discernibility can be stated equivalently using the model-theoretic notion of a *type*. Suppose A is a structure and suppose a_1, \dots, a_n are elements of the domain of A . As in the definition of grades of discernibility, we define the type of (a_1, \dots, a_n) in A relative to \mathcal{L}_A . It is clear how the corresponding definition will go for the type of (a_1, \dots, a_n) in A relative to \mathcal{L}_A^- , and the type of (a_1, \dots, a_n) in A^* relative to \mathcal{L}_{A^*} or $\mathcal{L}_{A^*}^-$.

The *type of (a_1, \dots, a_n) in A relative to \mathcal{L}_A* is intended to include everything that can be said truly of a_1, \dots, a_n , taken in that order, using only the expressive resources found in \mathcal{L}_A .

DEFINITION 5.3. (Type) *Suppose a_1, \dots, a_n are in the domain of A . Then,*

$$\text{Type}_A((a_1, \dots, a_n)) = \{\varphi(x_1, \dots, x_n) \text{ in } \mathcal{L}_A : A \models \varphi(a_1, \dots, a_n)\}$$

Note that this notion of type is relative to a structure and a language. If we keep the domain of the structure fixed, but vary the constant elements, relations, and functions in it, the type of a sequence of objects will generally vary as well. For

instance, as the structure gains constant elements and relations and functions, its associated language will gain expressive resources, and the type will generally expand. Similarly, if we expand the expressive resources of the language, the type will expand. Thus, the type of (a_1, \dots, a_n) in A relative to \mathcal{L}_A^- will typically contain more formulae than the type of (a_1, \dots, a_n) in A relative to \mathcal{L}_A .

It follows quickly from the definitions above that:

THEOREM 5.4. *Suppose A is a structure and a and b are elements of its domain. Then*

- (1) *a and b are absolutely discernible in A relative to \mathcal{L}_A iff $\text{Type}_A(a) \neq \text{Type}_A(b)$.*
- (2) *a and b are relatively discernible in A relative to \mathcal{L}_A iff $\text{Type}_A(a, b) \neq \text{Type}_A(b, a)$.*
- (3) *a and b are weakly discernible in A relative to \mathcal{L}_A iff $\text{Type}_A(a, b) \neq \text{Type}_A(a, a)$.*

§6. Why weak discernibility is important In this section, we ask how discerning is the relation of weak discernibility. In particular, we try to answer the question: Is weak discernibility the most discerning non-trivial, natural discernibility relation? That is: we know that, in general, weak discernibility is not as discerning as numerical distinctness (witness the edgeless graph with two vertices); but is it the most discerning natural relation that is less discerning than numerical distinctness? We explore this question by following two different strategies that might be expected to lead to more discerning relations. But we will show that neither does, and nor does a combination of the two strategies. We remind readers that we are not taking a stand on the expressive adequacy of these languages, only exploring logical relations.

6.1. Adding a constant for each object Above, we saw that each of the discernibility relations in A relative to \mathcal{L}_A are distinct and form a hierarchy ranging from the most discerning (Weak_A) to the least discerning (Int_A). We might expect that, when we add to our language a constant for every element in the domain of the structure in questions, all the discernibility relations will collapse into the numerical distinctness relation. That is, we might expect that the discernibility relations Int_{A^*} , Abs_{A^*} , Rel_{A^*} , and Weak_{A^*} all discern all pairs of distinct elements in the domain of A . But this isn't so. Consider for instance the edgeless graph with two vertices a and b . Then, as we saw above, a and b are not weakly discernible in the language of graph theory without constants or identity. But neither are they weakly discernible if we add to our language constants \bar{a} and \bar{b} that are interpreted as naming a and b respectively.

If this seems surprising, it may be for the following reason. Given an object a , a *haecceity of a in A* is a formula $\psi_a(x)$ that holds of a , but not of any other object in A . If each element of a structure is equipped with an haecceity in that structure, then every element is intrinsically discernible from every other: if $a \neq b$, then a and b are intrinsically discerned by $\psi_a(x)$ in A . Thus, in the presence of haecceities for every object, all discernibility relations collapse into numerical distinctness. It is often thought that, by introducing a constant to name a given object, we thereby introduce a haecceity for that object, in which case all discernibility relations relative to \mathcal{L}_{A^*} would collapse into numerical distinctness. But this thought is incorrect. In order to introduce haecceities for all objects, we would require constants for every object *as well as the identity relation*: we could

12 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

then define $\psi_a(x) := x = a$. Thus, all discernibility relations relative to $\mathcal{L}_{A^*}^-$ collapse into numerical distinctness. But without the identity relation, there need not be all (or indeed any) haecceities, so there is no collapse. Here we see the importance of distinguishing between discernibility by means of object-involving properties and discernibility by means of identity-involving properties. We will encounter this distinction again below.

So numerical distinctness is more discerning than weak discernibility relative to \mathcal{L}_{A^*} . But it might seem that, even if the new expressive resources provided by adding constants for every object are not always sufficient to discern *all* pairs of distinct objects, they should nonetheless be sufficient at least sometimes to discern *more* pairs of objects than are discerned by weak discernibility relative to \mathcal{L}_A . The following perhaps surprising theorem shows that this is not true: exactly the same pairs of objects are discerned by weak discernibility relative to \mathcal{L}_A , and by weak discernibility relative to \mathcal{L}_{A^*} . Also surprising is that in A^* relative to \mathcal{L}_{A^*} , absolute, relative, and weak discernibility all discern exactly the same elements.

THEOREM 6.5. *The following are equivalent:*

- (1) *a and b are weakly discernible in A relative to \mathcal{L}_A .*
- (2) *a and b are absolutely discernible in A^* relative to \mathcal{L}_{A^*} .*
- (3) *a and b are relatively discernible in A^* relative to \mathcal{L}_{A^*} .*
- (4) *a and b are weakly discernible in A^* relative to \mathcal{L}_{A^*} .*

Proof.

- (1) \Rightarrow (2). Suppose a and b are weakly discernible in A relative to \mathcal{L}_A . That is, there is a formula $\varphi(x, y)$ of \mathcal{L}_A with two free variables such that:

$$A \models \varphi(x, y)[a, b] \quad \text{and} \quad A \not\models \varphi(x, y)[a, a].$$

Then let \bar{a} be the constant in the signature of A^* that names a . Then $\varphi(\bar{a}, y)$ is a formula of \mathcal{L}_{A^*} with one free variable such that:

$$A^* \models \varphi(\bar{a}, y)[b] \quad \text{but} \quad A^* \not\models \varphi(\bar{a}, y)[a].$$

Thus, a and b are absolutely discernible in A^* relative to \mathcal{L}_{A^*} .

- (2) \Rightarrow (3); (3) \Rightarrow (4). Similar to analogous implications in Theorem 5.2..
- (4) \Rightarrow (1). Suppose a and b are weakly discernible in A^* relative to \mathcal{L}_{A^*} . That is, there is a formula $\varphi(x, y)$ in \mathcal{L}_{A^*} with two free variables such that

$$A^* \models \varphi(x, y)[a, b] \quad \text{and} \quad A^* \not\models \varphi(x, y)[a, a]$$

Let d_1, \dots, d_m be all and only those elements of the domain of A such that the constants $\bar{d}_1, \dots, \bar{d}_m$ that name them occur in the formula $\varphi(x, y)$. Then

$$A^* \models \varphi(x, y, \bar{d}_1, \dots, \bar{d}_m)[a, b] \quad \text{and} \quad A^* \not\models \varphi(x, y, \bar{d}_1, \dots, \bar{d}_m)[a, a]$$

Now, the formula

$$\psi(x, y) := \exists z_1 \dots \exists z_m (\varphi(x, y, z_1, \dots, z_m) \wedge \neg \varphi(x, x, z_1, \dots, z_m))$$

is in \mathcal{L}_A . Moreover, it is easy to verify that we have $A^* \models \psi(x, y)[a, b]$ but $A^* \not\models \psi(x, y)[a, a]$ and thus also $A \models \psi(x, y)[a, b]$ and $A \not\models \psi(x, y)[a, a]$. This shows that a and b are weakly discernible in A relative to \mathcal{L}_A .

This completes our proof. \square

The theorem shows that weak discernibility is surprisingly discerning. Indeed, it discerns exactly as much as can be discerned in a language equipped with a constant for every object. Does this undermine the claim that, by appealing to weak discernibility, we can save PII from the apparent mathematical and physical counterexamples? Does it provide ammunition for those philosophers—for instance, those mentioned in footnote 7—who have been suspicious of such philosophical uses of weak discernibility? One might think that, since it is illegitimate to discern two objects by appealing to a relation defined by a formula that involves singular reference to either object, and since it turns out that this is possible exactly when those objects can be weakly discerned without such resources, then it is illegitimate to say that two objects are discernible in a structure A when they are only weakly discernible in A relative to \mathcal{L}_A . But of course the argument could be run the other way: the legitimacy of discernibility relative to \mathcal{L}_{A^*} might be inferred from the alleged legitimacy of weak discernibility relative to \mathcal{L}_A by means of the theorem just proved.

We have seen that we cannot define a more discerning relation than weak discernibility by adding constants to our language. Anything we can discern in any way by adding constants was already weakly discernible before we added them. In the next section, we explore a different strategy for finding a weaker discernibility relation than weak discernibility. However, first, we consider a question that arises naturally from the results in this section.

6.2. Substitution salva veritate We wish to know whether there is a natural discernibility relation that is less discerning than numerical distinctness and more discerning than weak discernibility relative to \mathcal{L}_A . Recall the definition of that discernibility relation: a and b are weakly discernible in A relative to \mathcal{L}_A if there is a formula $\varphi(x, y)$ in \mathcal{L}_A such that $A \models \varphi(a, b)$ and $A \not\models \varphi(a, a)$. But why allow only formulae with two free variables to do the discerning? Perhaps we can find a more discerning relation if we allow also formulae $\varphi(x, y, z)$ with three free variables. For instance, we might say that a and b are very weakly discernible in A relative to \mathcal{L}_A if there is $\varphi(x, y)$ in \mathcal{L}_A such that $A \models \varphi(a, b)$ and $A \not\models \varphi(a, a)$ or if there is $\varphi(x, y, z)$ in \mathcal{L}_A such that $A \models \varphi(a, b, a)$ and $A \not\models \varphi(b, a, a)$. In fact, why choose that ordering of the a 's and b 's in the three-place formula? And why not allow four-place formulae as well as two- and three-place? Indeed, why not allow formulae of any arity? Perhaps there is an infinite hierarchy of discernibility relations between numerical distinctness and weak discernibility, each of which is determined by the arity of formula that is allowed to discern the objects, and the order in which the objects must appear in those formulae. If this is so, the most discerning relation in this hierarchy will be the relation that permits *any* formula in \mathcal{L}_A to discern the objects, and allows the objects to appear in *any* order. We call this *the very weak discernibility* relation.

DEFINITION 6.6. (Very weak discernibility) *Suppose A is a structure and a and b are elements of the domain of A . Then*

- a and b are very weakly discernible in A relative to \mathcal{L}_A (written $VWeak_A(a, b)$) if there is $\varphi(x_1, \dots, x_n)$ in \mathcal{L}_A and $c_1, \dots, c_n, d_1, \dots, d_n$ in the domain of A such that

14 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

each c_i and d_i is either a or b and

$$A \models \varphi(c_1, \dots, c_n) \quad \text{and} \quad A \not\models \varphi(d_1, \dots, d_n)$$

Does this discernibility relation lie strictly between numerical distinctness and weak discernibility? The following theorem shows that it does not. Moreover, it shows that the infinite hierarchy of discernibility relations described in the previous paragraph collapses with it into weak discernibility.

THEOREM 6.7. *The following are equivalent:*

- (1) a and b are weakly discernible in A relative to \mathcal{L}_A .
- (2) a and b are very weakly discernible in A relative to \mathcal{L}_A .

Proof.

- (1) \Rightarrow (2). Trivial.
- (2) \Rightarrow (1). Suppose a and b are very weakly discernible in A relative to \mathcal{L}_A . That is, there is $\varphi(x_1, \dots, x_n)$ and $c_1, \dots, c_n, d_1, \dots, d_n \in \{a, b\}$ such that

$$A \models \varphi(c_1, \dots, c_n) \quad \text{and} \quad A \not\models \varphi(d_1, \dots, d_n)$$

Then construct $\psi(x, y)$ as follows. First, let φ^c be obtained from φ by replacing each free variable x_i in φ by the free variable x if $c_i = a$ and by the free variable y if $c_i = b$. Thus, $A \models \varphi^c(x, y)[a, b]$. Next, let φ^d be obtained from φ by replacing each x_i by x if $d_i = a$ and by y if $d_i = b$. Thus, $A \not\models \varphi^d(x, y)[a, b]$. Then let $\psi(x, y) := \varphi^c(x, y) \wedge \neg\varphi^d(x, y)$. Then:

$$A \models \psi(x, y)[a, b] \quad \text{and} \quad A \not\models \psi(x, y)[a, a].$$

So a and b are weakly discernible in A relative to \mathcal{L}_A .

This completes our proof. □

Thus, again, we have failed to identify a discernibility relation that lies strictly between numerical distinctness and weak discernibility.

Let's finally consider whether one might do better by combining the strategy of this section with that of the previous section. That is, is very weak discernibility relative to \mathcal{L}_{A^*} more discerning than weak discernibility relative to \mathcal{L}_A ? The answer turns out to be negative, as can be seen by a fairly straightforward adaptation of the proof of Theorem 6.5..

6.3. Discerning using atomic formulae The results of the previous two sections suggest that weak discernibility relative to \mathcal{L}_A is indeed the most discerning natural discernibility relation after numerical distinctness. We will now try to give this idea more precise technical content.

THEOREM 6.8. *The following are equivalent:*

- (1) a and b are weakly discernible in A^* by an atomic formula of \mathcal{L}_{A^*} or a negation thereof.
- (2) a and b are weakly discernible in A^* by a quantifier-free formula of \mathcal{L}_{A^*} .
- (3) a and b are weakly discernible in A^* relative to \mathcal{L}_{A^*} .

Proof.

- (1) \Rightarrow (2); (2) \Rightarrow (3). Trivial.

- (3) \Rightarrow (1). The contrapositive is established by induction on the complexity of the formulae in \mathcal{L}_{A^*} . The atomic case is given, and the inductive step is straightforward.

This completes our proof. □

If one wants the discerning properties to be ‘natural’ and is concerned that ‘naturalness’ may be lost in the transition from atomic formulae to logically complex ones, then this theorem addresses one’s concern.

By Theorems 6.7. and 6.8., we have that weak discernibility relative to \mathcal{L}_A discerns exactly the elements of the domain of A that are discerned by the atomic part of \mathcal{L}_{A^*} . This fact gives us a robust sense in which weak discernibility is the most discerning natural discernibility relation after numerical distinctness.

6.4. *Discernibility in the object language* Our discernibility relations have so far been defined using the meta-language: their definitions have involved quantification over formulae as well as the satisfaction relation between formulae and models. This raises the question whether we can define any of our discernibility relations in the object language. We now show that this is possible for weak discernibility in a restricted class of languages.

We begin with a definition due to Hilbert and Bernays Hilbert & Bernays (1934).

DEFINITION 6.9. (Hilbert-Bernays discernibility) *Suppose A is a structure containing only finitely many relations. Then, for each n -ary relation symbol R in the signature of A , define*

$$\text{Ind}_R(x, y) := \forall z_1 \dots \forall z_{n-1} ((R(x, z_1, \dots, z_{n-1}) \leftrightarrow R(y, z_1, \dots, z_{n-1}) \wedge \dots \wedge R(z_1, \dots, z_{n-1}, x) \leftrightarrow R(z_1, \dots, z_{n-1}, y))$$

Next, define

$$HB_A(x, y) := \neg \left(\bigwedge_R \text{Ind}_R(x, y) \right)$$

where R ranges over the relation symbols in the signature of A .

We say that a and b are Hilbert-Bernays discernible in A if $A \models HB_A(a, b)$.

We can now prove that, for the languages for which it is defined, Hilbert-Bernays discernibility is equivalent to weak discernibility.¹³

THEOREM 6.10. *The following are equivalent:*

- (1) a and b are Hilbert-Bernays discernible in A .
- (2) a and b are weakly discernible in A relative to \mathcal{L}_A .

Proof.

- (1) \Rightarrow (2). Suppose $A \models HB_A(a, b)$. It is clear that $A \not\models HB_A(a, a)$. Thus, since $HB_A(x, y)$ is a formula in \mathcal{L}_A , a and b are weakly discernible in A relative to \mathcal{L}_A .

¹³ Our notion of Hilbert-Bernays discernibility corresponds to Ketland (2011)’s notion of ‘first-order indiscernibility’, and the following theorem, to his Theorem 3.17.

16 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

- (2) \Rightarrow (1). Suppose a and b are weakly discernible in A relative to \mathcal{L}_A . Then Theorem 6.8. ensures that a and b are discernible by an atomic formula in \mathcal{L}_{A^*} . It is easy to see that it follows from this that $A \models HB_A(a, b)$.

This completes our proof. □

§7. Discernibility in languages with identity In section 6.1., we considered what more can be discerned if we allow the discerning formulae to use constants for each element of the domain. We noted that the discernibility relations do not collapse into numerical distinctness because, in the absence of the identity relation, we cannot use the constants to define haecceities for each object. In this section, we ask what happens if we add the identity relation, but no constants. *Prima facie*, this may seem philosophically inappropriate and technically in danger of collapsing all notions of discernibility to numerical distinctness. However, we will now see that both worries are unfounded. To allay the latter fear, we note the following analogue of Theorem 5.2.:

THEOREM 7.11.


- $\text{Int}^{\bar{=}} \Rightarrow \text{Abs}^{\bar{=}} \Rightarrow \text{Rel}^{\bar{=}} \Rightarrow \text{Weak}^{\bar{=}} \Rightarrow \text{Dis}$
- $\text{Dis} \Rightarrow \text{Weak}^{\bar{=}} \not\Rightarrow \text{Rel}^{\bar{=}} \not\Rightarrow \text{Abs}^{\bar{=}} \not\Rightarrow \text{Int}^{\bar{=}}$

Proof. The proof is exactly the same as the proof of Theorem 5.2.. $\text{Dis} \Rightarrow \text{Weak}^{\bar{=}}$ since $x \neq y$ is a formula in $\mathcal{L}_A^{\bar{=}}$. □

Next we investigate the logical relations between weak, relative, and absolute discernibility in A relative to \mathcal{L}_A and in $\mathcal{L}_A^{\bar{=}}$.

THEOREM 7.12. *Each discernibility relation relative to \mathcal{L}_A is strictly less discerning than its counterpart relative to $\mathcal{L}_A^{\bar{=}}$.*

Proof. Trivially, $\text{Abs} \Rightarrow \text{Abs}^{\bar{=}}$, $\text{Rel} \Rightarrow \text{Rel}^{\bar{=}}$, and $\text{Weak} \Rightarrow \text{Weak}^{\bar{=}}$. The following graph-theoretic example shows that $\text{Abs}^{\bar{=}} \not\Rightarrow \text{Abs}$ and $\text{Rel}^{\bar{=}} \not\Rightarrow \text{Rel}$:




Here we have $\text{Abs}_A^{\bar{=}}(a, b)$ and $\text{Rel}_A^{\bar{=}}(a, b)$, but not $\text{Abs}_A(a, b)$ or $\text{Rel}_A(a, b)$. The problem is that, without the identity relation, there is no way of saying that a is related to one object while b is related to two. However, if we add identity to this language, it is possible to say this.

In the example of the edgeless graph with two vertices a and b , we have $\text{Weak}_A^{\bar{=}}(a, b)$ but not $\text{Weak}_A(a, b)$. □

Does the notion of discernibility in A relative to $\mathcal{L}_A^{\bar{=}}$ have any philosophical interest? If the philosophical project is to study how identity facts ‘supervene on’ or are ‘grounded in’ non-identity-involving facts, then discernibility relative to $\mathcal{L}_A^{\bar{=}}$ is clearly not the relevant notion, since this language simply presupposes the relation we are trying to ground. This sort of project requires that the discernibility take place in a language that only expresses non-identity-involving facts.

However, studying discernibility relative to $\mathcal{L}_{\bar{A}}$ may still be of some interest, both philosophically and for the sake of technical completeness. There is at least one notion of identity for which discernibility relative to $\mathcal{L}_{\bar{A}}$ is relevant, namely ‘identity’ in the sense of ‘defining characteristics’. For instance, sociologists may say that it is part of someone’s identity that she is the mother of three children or the murderer of twelve innocent victims. Such claims are naturally expressed in $\mathcal{L}_{\bar{A}}$.

The following diagram summarizes our results so far (together with the results of Corollary 9.21. below). This represents all the entailment relations between the various discernibility relations relative to \mathcal{L}_A , $\mathcal{L}_{\bar{A}}$, and \mathcal{L}_{A^*} :



§8. Quotients The results of the previous sections suggest that there is a sense in which weak discernibility is the most discerning relation that the structure can ‘see’. In this section, we make this precise in a particular way. In what follows, when a and b are not weakly discernible in A relative to \mathcal{L}_A , we say that they are *utterly indiscernible in A* and write $a \approx_A b$. Throughout this section and the next, we assume that \mathcal{L}_A contains no individual constants and no function symbols. The associated structures are called *relational*.

We begin by generalizing an observation of Ketland’s (Theorem 3.12, Ketland (2011)). Ketland shows that, if a structure A contains pairs of objects that are utterly indiscernible in A , then there is another structure \bar{A} in which any two utterly indiscernible object are ‘identified’, and such that exactly the same sentences of \mathcal{L}_A are true in A as are true in \bar{A} : when two structures have the property of making exactly the same sentences of \mathcal{L}_A true, we say that they are *elementarily equivalent*.

We note that Ketland’s result remains true if we seek a structure \bar{A} in which not *all* pairs of utterly indiscernible elements are identified, but only those pairs that both belong to the same set C_i from a family $\{C_i \subseteq \text{dom}(A) : i \in I\}$ of disjoint subsets of $\text{dom}(A)$, where any two elements of C_i are utterly indiscernible.

THEOREM 8.13. *Suppose that A is a relational structure. Suppose that $\mathcal{F} = \{C_i \subseteq \text{dom}(A) : i \in I\}$ is a family of disjoint subsets of the domain of A such that $a, b \in C_i$ implies $a \approx_A b$. Then there is a structure \bar{A} and a surjective function $\tau : \text{dom}(A) \rightarrow \text{dom}(\bar{A})$ such that:*

- (1) *For all $a, b \in \text{dom}(A)$, then $\tau(a) = \tau(b)$ iff there is $i \in I$ such that $a, b \in C_i$.*
- (2) *A and \bar{A} are elementarily equivalent with respect to \mathcal{L}_A .*

Proof. Our proof employs the notion of a *quotient structure*. We develop this notion quite generally at first; then we apply our findings to prove our generalization of Ketland’s theorem.

Let A be a structure. An equivalence relation \sim on its domain is said to be an \mathcal{L}_A -congruence on A if it respects the interpretation of the predicates of \mathcal{L}_A , in the sense that for any n -place predicate P in \mathcal{L}_A and objects $a_1, \dots, a_n, a'_1, \dots, a'_n \in A$ with $a_i \sim a'_i$, we have: $A \models P(a_1, \dots, a_n)$ iff $A \models P(a'_1, \dots, a'_n)$. When \sim is an \mathcal{L}_A -congruence, we define the *quotient of A under \sim* (written A/\sim) as follows:

- (a) The domain of A/\sim is the set of equivalence classes of \sim in $\text{dom}(A)$: that is, $\{[a] : a \in \text{dom}(A)\}$ where $[a] = \{b \in \text{dom}(A) : a \sim b\}$.
- (b) If R is a relation in A , then we define R_\sim in A/\sim as follows: $R_\sim = \{([a_1], \dots, [a_n]) : (a_1, \dots, a_n) \in R\}$. This is well defined because \sim is a congruence.

Then we have

- (1') If $a, b \in M$, then $[a] = [b]$ iff $a \sim b$. Thus, the function $\tau : a \mapsto [a]$ is surjective and identifies all and only objects that are related by \sim .
- (2') By induction on the complexity of formulae, it is straightforward to see that A and A/\sim are elementarily equivalent.


Thus, to establish claims (1) and (2) of our generalization of Ketland's theorem, it suffices to show that the relation

$$x \sim_{\mathcal{F}} y := (\exists i \in I)(x, y \in C_i) \vee x = y$$

is an \mathcal{L}_A -congruence. By the disjointness of the C_i , it follows that $\sim_{\mathcal{F}}$ is an equivalence relation. And by Theorem 6.8., we know that $\sim_{\mathcal{F}}$ respects the interpretation of the predicates of \mathcal{L}_A . Thus, we let $\bar{A} = A/\sim_I$ and $\tau : a \mapsto [a]$. \square

This result is particular to weak discernibility. If we required only that any $a, b \in C_i$ not be *relatively* discernible, then the result would not hold. The reason is that relative indiscernibility does not make $\sim_{\mathcal{F}}$ into a congruence. This is witnessed by the dumbbell graph, letting the only C_i be the set $\{a, b\}$ that contains both vertices. Then we have $a \sim_{\mathcal{F}} b$, as well as Eab and Eba , but not Eaa or Ebb . Similarly for absolute discernibility.

What about the converse of the theorem? Suppose that A and A' are elementarily equivalent structures and $\tau : \text{dom}(A) \rightarrow \text{dom}(A')$ is a surjective function that identifies a and b and nothing else. Does it follow that $a \approx_A b$? The answer turns out to be negative, as is witnessed by the following two structures A (on the left) and A' (on the right) and the surjective function τ such that $\tau(a_1) = a' = \tau(a_2)$, $\tau(b_i) = \tau(b'_i)$, and $\tau(c) = c'$:



For A and A' are elementarily equivalent, yet a_1 and a_2 are weakly discernible in A .

However, a reverse of the model construction is possible and has interesting philosophical consequences. In our generalization of Ketland's result, we showed

that, for any structure A , there is always an elementarily equivalent structure \bar{A} in which sets of utterly indiscernible objects are identified. In the following theorem, we show that there is always an elementarily equivalent structure \tilde{A} that replaces any chosen elements from the domain of A with any desired number of utterly indiscernible objects.

THEOREM 8.14. *Suppose that A is a relational structure. Suppose $\{\kappa_i : i \in I\}$ is a family of cardinal numbers indexed by some $I \subseteq \text{dom}(A)$. Then there is a structure \tilde{A} with the same signature as A , a family $\mathcal{F} = \{C_i \subseteq \text{dom}(\tilde{A}) : i \in I\}$ of disjoint subsets of $\text{dom}(\tilde{A})$ indexed by I , and a surjective function $\tau : \text{dom}(\tilde{A}) \rightarrow \text{dom}(A)$ such that:*

- (1) for all $i \in I$, $|C_i| = \kappa_i$;
- (2) for all $a, b \in \text{dom}(\tilde{A})$, if there is $i \in I$ such that $a, b \in C_i$, then $a \approx_{\tilde{A}} b$;
- (3) for all $a, b \in \text{dom}(\tilde{A})$, then $\tau(a) = \tau(b)$ iff there is $i \in I$ such that $a, b \in C_i$;
- (4) A and \tilde{A} are elementarily equivalent.

In fact, if the operation $B \mapsto \bar{B}$ is defined as in Theorem 8.13., then $\bar{\tilde{A}}$ is isomorphic to A .

Proof. We construct $\text{dom}(\tilde{A})$ as follows. First, let $\mathcal{F} = \{C_i : i \in I\}$ be a family of disjoint sets that are also disjoint from $\text{dom}(A)$ and such that $|C_i| = \kappa_i$ for each $i \in I$. This gives (1). Let $\text{dom}(\tilde{A})$ be $(A - I) \cup \bigcup_{i \in I} C_i$. Suppose P is an n -place predicate symbol in \mathcal{L}_A . Define \tilde{A} so that $\tilde{A} \models P(a_1, \dots, a_n)$ iff $A \models P(a'_1, \dots, a'_n)$, where for each j , either $a_j \in A - I$ and $a'_j = a_j$, or $a_j \in C_i$ for some $i \in I$ and $a'_j = i$. Note that if $a, b \in C_i$, then no atomic formula or negation thereof weakly discerns a and b in \tilde{A}^* (by definition). By Theorem 6.8., a and b are weakly indiscernible in \tilde{A}^* . By Theorem 6.5., they are weakly indiscernible in \tilde{A} . This proves (2). Now consider $\tau : \tilde{A} \rightarrow \bar{\tilde{A}}$ as defined in Theorem 8.13.. It is straightforward to show that $\bar{\tilde{A}}$ is isomorphic to A . Hence (3) and (4) follow easily. \square

Theorem 8.14. has some philosophical significance: it reveals a limitation on the sort of facts that we might come to know by empirical investigation. If all of our empirical evidence is expressible in \mathcal{L}_A , then Theorem 8.14. tells us that this evidence will never help us distinguish between two models where one is obtained from the other by adding many utterly indiscernible objects in place of a single object at various points. In order to distinguish two such models, we must appeal to theoretical considerations such as ontological parsimony.¹⁴

Thus, the upshot of this Theorem 8.14. is similar to the upshot of the Löwenheim-Skolem theorems. The latter tell us that, if our empirical evidence is expressible in $\mathcal{L}_{\bar{A}}$, then it will never help us distinguish between two models of different infinite cardinality (if there are any such models). In order to distinguish two such models, we must appeal to explanatory power or simplicity or some other theoretical virtue. The argument from Theorem 8.14. is both more and less general than this argument. On the one hand, it requires the stronger assumption that our empirical evidence is expressible in \mathcal{L}_A . On the other hand, it applies equally to finite models and to

¹⁴ See Baker (2003) for one such attempt.

infinite ones, and it provides greater control over the construction of alternative models.¹⁵

So what of the assumption that our empirical evidence is expressible in \mathcal{L}_A ? Perhaps it will be objected that some of our evidence can only be expressed in \mathcal{L}_A^- . If this is conceded, our argument fails. However, it is not clear that the empiricist should be happy with the idea that some of the evidence we obtain can only be expressed using a primitive identity predicate. Identity facts, the empiricist might hold, must be inferred from our theory; they cannot be given to us in our evidence.

§9. Discernibility and symmetries A symmetry of a model is a permutation of the elements of that model that ‘respects’ the properties the objects have and the relations in which they stand.

DEFINITION 9.15. (Symmetry) *Suppose A is a relational structure.*

- (a) *A permutation of A is a bijection $\pi : \text{dom}(A) \rightarrow \text{dom}(A)$. We denote by π_{ab} the permutation that swaps a and b and leaves everything else unchanged.*
- (b) *A permutation π of A is a symmetry of A if, for all relation symbols R in A and elements a_1, \dots, a_n of the domain of A , $Ra_1\dots a_n$ iff $R\pi(a_1)\dots\pi(a_n)$.*

Symmetries are usually called *automorphisms* in model theory. Here are two examples. In Max Black’s example considered above, a permutation that swaps the two spheres is a symmetry because it ‘respects’ the structure of the universe of that example. Similarly, in Leibniz’s example, if the universe is infinite in extent, then to move every point one metre to the left would be to effect a permutation of the points that is a symmetry.

A straightforward induction on the complexity of formulae gives the following basic result:

LEMMA 9.16. *Suppose π is a symmetry of A . Then, for all $\varphi(x_1, \dots, x_n)$ in \mathcal{L}_A or \mathcal{L}_A^- , and all $a_1, \dots, a_n \in M$ we have:*

$$A \models \varphi(a_1, \dots, a_n) \text{ iff } A \models \varphi(\pi(a_1), \dots, \pi(a_n))$$

It is often thought that there is a close connection between symmetries and discernibility relations.¹⁶ In this final section, we explore this connection. We show

¹⁵ Indeed, Theorem 8.14. can be read as an Upwards Löwenheim-Skolem theorem that applies equally to finite models, provided that the elementary equivalence that is required is defined relative to \mathcal{L}_A rather than to \mathcal{L}_A^- . That is, Theorem 8.14. shows that, given a theory in \mathcal{L}_A with a model of cardinality κ , where κ is finite or infinite, this theory has an elementarily equivalent model of cardinality λ for all $\lambda > \kappa$.

¹⁶ This is suggested by much discussion of the Burgess-Keränen objection to *ante rem* structuralism, according to which this form of structuralism is incompatible with structures containing distinct objects that are not absolutely discernible, such as the structure of the complex field, where the two square roots of -1 are not absolutely discernible. Although the problem is thus officially concerned with the existence of indiscernible objects, it is frequently described as one of non-trivial automorphisms (e.g. in several of the discussions of mathematical structuralism cited in footnote 4). Some authors even suggest that the existence of an automorphism of a structure that swaps a and b is equivalent to the claim that a and b are not absolutely discernible (Keränen, 2001, p. 318) (although see also p. 323).

that facts about utter indiscernibility entails the existence of symmetries, which in turn entails facts about objects not being relatively and absolutely discernible, but that the converse implications generally fail. The connection between symmetries and discernibility relations is thus not as close as one might imagine.

We begin by providing a sufficient condition for a permutation π to be a symmetry.¹⁷

THEOREM 9.17. *Suppose A is a relational structure. If π is a permutation of A and $a \approx_A \pi(a)$ for all $a \in \text{dom}(A)$, then π is a symmetry of A . But the converse does not hold in general.*

Proof. We establish the first claim by observing that for any relation R in A and any objects a_1, \dots, a_n from the domain of A we have:

$$Ra_1 \dots a_n \Leftrightarrow R\pi(a_1)a_2 \dots a_n \Leftrightarrow \dots \Leftrightarrow R\pi(a_1) \dots \pi(a_n)$$

The dumbbell graph provides a counterexample to the converse. For the permutation π_{ab} is a symmetry of this graph, but we have $a \not\approx_A b = \pi_{ab}(a)$. \square

An important corollary of this:¹⁸

COROLLARY 9.18. *Suppose $a, b \in \text{dom}(A)$ and $a \approx_A b$. Then π_{ab} is a symmetry of A .*

We now turn to necessary conditions for π to be a symmetry. But first a useful definition.

DEFINITION 9.19. (Symmetric elements) *Let a and b be elements of the domain of some relational structure A .*

- (a) *If π is a symmetry of A such that $\pi(a) = b$, we say that a and b are symmetric relative to π .*
- (b) *If π is a symmetry of A such that $\pi(a) = b$ and $\pi(b) = a$, we say that a and b are fully symmetric relative to π .*
- (c) *We say that a and b are (fully) symmetric iff there is a symmetry π of A relative to which they are (fully) symmetric.*

The first two results of the following theorem are well-known from the philosophical literature on identity and indiscernibility.¹⁹ The latter two are familiar from the model theory literature, though only recently noted by philosophers. They were brought to our attention independently by Butterfield and Caulton and Kate Hodesdon—see Comment (2), §2.2.2, Caulton & Butterfield (ta).

THEOREM 9.20. *Let a and b be elements of the domain of some relational structure A .*

- (1) *If a and b are symmetric, then a and b are not absolutely discernible in A relative to \mathcal{L}_A^- .*

¹⁷ See all Comment (1), §2.2.2, Caulton & Butterfield (ta).

¹⁸ See also Ketland (2011), Theorem 3.23.

¹⁹ Our Theorem 9.20.(1) is Theorem 1 of Caulton & Butterfield (ta); the first half of our Theorem 9.20.(4) is Theorem 2 of Caulton & Butterfield (ta); the first half of our Theorem 9.20.(3) is noted in the discussion following Theorem 2 of Caulton & Butterfield (ta).

22 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

- (2) *If a and b are fully symmetric, then a and b are not relatively discernible in A relative to \mathcal{L}_A^- .*
- (3) *The converses of (1) and (2) do not hold in general. That is, it is not the case that, for every structure A , if a and b are not absolutely (relatively) discernible in A , then a and b are (fully) symmetric.*
- (4) *The converses of (1) and (2) do hold if $\text{dom}(A)$ is finite.*

Proof. (1) and (2) are immediate corollaries of Lemma 9.16.. To prove (3), let \mathcal{L}_A^- be the first-order language with identity with a single two-place non-logical predicate $<$. Let A be the structure obtained by taking the disjoint union of the structures $\langle \mathbb{Q}, <_{\mathbb{Q}} \rangle$ and $\langle \mathbb{R}, <_{\mathbb{R}} \rangle$ (where ‘ $<$ ’ is interpreted as the union of the relations $<_{\mathbb{Q}}$ and $<_{\mathbb{R}}$). Let $a_1, \dots, a_n, a'_1, \dots, a'_n$ be rational numbers from either \mathbb{Q} or \mathbb{R} such that, for every i , a_i and a'_i are the same rational number, just one from \mathbb{Q} and the other from \mathbb{R} . Then an easy induction on complexity of formulae show that for every ϕ from \mathcal{L}_A^- we have:

$$A \models \phi(a_1, \dots, a_n) \text{ iff } A \models \phi(a'_1, \dots, a'_n)$$

This entails that $0_{\mathbb{Q}}$ and $0_{\mathbb{R}}$ are not relatively (and thus also not absolutely) discernible in A relative to \mathcal{L}_A^- . However, cardinality considerations show that there is no symmetry π of A such that $\pi(0_{\mathbb{Q}}) = 0_{\mathbb{R}}$. This establishes the two claims that make up (3). For the proof of (4), we refer the reader to Caulton & Butterfield (ta). \square

An interesting consequence of Corollary 9.18. and Theorem 9.20.(1) and (2) is the following:

COROLLARY 9.21. *If a and b are absolutely or relatively discernible in A relative to \mathcal{L}_A^- , then a and b are weakly discernible in A relative to \mathcal{L}_A .*

Proof. Suppose a and b are not weakly discernible in A relative to \mathcal{L}_A . Then, by Corollary 9.18., π_{ab} is a symmetry of A . Thus, by Theorem 9.20.(1) and (2), a and b are not absolutely discernible nor relatively discernible in A relative to \mathcal{L}_A^- . \square

Let’s return to our question about the relation between the notion of a symmetry and our three notions of discernibility. An answer is provided by our Theorems 9.17. and 9.20., which show that the former notion cannot in general be characterized in terms of any of the latter notions, except in the case of finite structures, where two objects cannot be absolutely (relatively) discerned just in case they are (fully) symmetric. Our results thus provide a salutary reminder that the connection between these two families of notions isn’t as straightforward as one might initially think.

On reflection, this loose connection should not be too surprising. The two families of notions track completely different philosophical ideas. The discernibility notions are concerned with the extent to which the dyadic relations of identity and distinctness are fixed or determined by relations expressible in a language \mathcal{L}_A or \mathcal{L}_A^- . These questions arise within one particular structure, which represents a way the world could be. The notions of discernibility are in this sense concerned with ‘intra-world’ matters. By contrast, permutations are concerned with what happens when objects *change* their positions in a structure, and symmetries are just the special case where this change of positions preserves all structural relations. Permutations and symmetries raise hard philosophical questions about the extent to which objects are

independent of their positions in a relational structure. If objects are independent in this way, then the result of carrying out a non-trivial permutation will result in a new and distinct way the world could be; and if not, not. The philosophical questions raised by permutations and symmetries are thus *transworld* questions, concerned with the identification of objects across different possible worlds or situations.


§10. A curious fact about relative discernibility So far, we have defined and investigated a small number of different discernibility relation, many of which have been discussed in the literature. But we have not offered any general definition of a discernibility relation; nor has anyone else, as far as we are aware. Can we do better?

A natural minimal requirement is that any notion of *indiscernibility* should give rise to an equivalence relation on the domain of any structure to which the notion is applied. This would entail a minimal requirement on any notion D of discernibility, namely that the discernibility relation D_A to which D give rise when applied to a structure A be the complement of an equivalence relation on A . Our final theorem shows that, while absolute and weak discernibility satisfy this proposed minimal requirement, relative discernibility rather surprisingly does not.

THEOREM 10.22.

- (1) *The relations Abs_A and Weak_A of absolute and weak discernibility in A relative to \mathcal{L}_A are complements of equivalence relations on the domain of any structure A ; and likewise for absolute and weak discernibility relative other other languages.*
- (2) *There are structures A such that the relations Rel_A and Rel_A^- of relative discernibility in A relative to respectively \mathcal{L}_A and \mathcal{L}_A^- are not the complements of an equivalence relation on the domain of A .*

Proof. By Theorem 6.5., we have that $\text{Weak}_A = \text{Abs}_{A^*}$ and that Weak_A^- is just the relation Dis of distinctness. And by Theorem 5.4., the absolute discernibility of a and b in some structure relative to some language is a matter of a and b having distinct types in this structure relative to this language. Since sameness of types is an equivalence relation, so are the relevant relations of not being absolutely discernible, which establishes (1). For (2), consider the following graph:



A routine verification establishes that the following two permutations:

$$(ac)(bd)(a'c')(b'd') \quad \text{and} \quad (bc)(a'd')(b'd)(ac')$$

are symmetries of the graph. Theorem 9.20.(2) then ensures that neither a and c , nor b and c , are relatively discernible in \mathcal{L}_A or \mathcal{L}_A^- . However, a and b are relatively

24 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

discerned by an edge of the graph. So it follows that the relation of not being relatively discernible fails to be transitive. (However, the relation is easily seen to be reflexive and symmetric.) \square

Theorem 10.22. shows that the proposed minimal requirement on the notion of a discernibility relation disqualifies relative discernibility from being a genuine form of discernibility. One response to the theorem is thus to abandon the proposed requirement. A more hard-nosed response would be to uphold the requirement and the ensuing verdict on relative discernibility, despite the fact that this would be a revision of existing practice, which regards relative discernibility as a genuine form of discernibility. We shall not here defend one of the two responses; the issue is anyway in part semantic.

However, we observe that the revisions required by the hard-nosed response would be comparatively modest. For relative discernibility is not nearly as important as absolute or weak discernibility. Absolute discernibility is philosophically important if one wants a monadic notion of discernibility, as discussed in Section §2.. And if one is willing to accept a dyadic notion, one may as well go for weak discernibility rather than relative. From a technical point of view as well, absolute and weak discernibility are more important than relative. The two former notions are connected with the important notion of the type of an object in a structure relative to languages respectively with and without constants. Moreover, our results in Section §6. highlighted the importance of weak discernibility as the most discerning non-trivial discernibility relation.

We conclude that little would be lost if philosophers were to pay less attention to relative discernibility.

§11. Conclusion A philosophical question about identity and discernibility ‘factorizes’ into two parts: first, a philosophical question about the adequacy of a formal language to a particular class of properties and relations in the world; and second, a mathematical question about the discernibility of objects in this language. We hope to have shed some light on the latter sort of question by proving various new theorems and extending some old ones.

Acknowledgements We would like to thank Leon Horsten, John Mayberry, and the other members of the Structuralism Seminar at University of Bristol, as well as an editor and two anonymous referees for this journal for valuable comments and suggestions. Thanks also to Jeremy Butterfield, Adam Caulton, and Jeffrey Ketland for sharing with us work that had not yet been published, and to the audience at the conference Structure and Identity in Bristol, where an earlier version of the paper was presented. All three authors were supported by the AHRC-funded project *Foundations of Structuralism*. Additionally, Øystein Linnebo received support from an ERC Starting Grant, and Richard Pettigrew, from a British Academy Postdoctoral Fellowship.

BIBLIOGRAPHY

Armstrong, D. (1980). *Universals and Scientific Realism*, Volume 1. Cambridge: Cambridge University Press.

- Baker, A. (2003). Quantitative Parsimony and Explanatory Power. *British Journal for the Philosophy of Science* **54**(2), 245–259.
- Bigaj, T., & Ladyman, J. (2010). The Principle of the Identity of Indiscernibles and Quantum Mechanics. *Philosophy of Science* **77**, 117–136.
- Black, M. (1952). The Identity of Indiscernibles. *Mind* **61**, 153–64.
- Burgess, J. P. (1999). Review of Stewart Shapiro, *Philosophy of Mathematics: Structure and Ontology*. *Notre Dame Journal of Formal Logic* **40**(2), 283–91.
- Caulton, A., & Butterfield, J. (t.a.). On Kinds of Indiscernibility in Logic and Metaphysics. *British Journal for the Philosophy of Science*.
- Hawley, K. (2009). Identity and indiscernibility. *Mind* **118**, 101–119.
- Hawthorne, J., & Sider, T. (2006). Locations. In John Hawthorne, *Metaphysical Essays* (Oxford University Press, 2006).
- Hilbert, D., & Bernays, P. (1934). *Grundlagen der Mathematik*. Berlin: Springer.
- Hodges, W. (1997). *A Shorter Model Theory*. Cambridge: Cambridge University Press.
- Keränen, J. (2001). The Identity Problem for Realist Structuralism. *Philosophia Mathematica* **9**(3), 308–330.
- Ketland, J. (2006). Structuralism and the Identity of Indiscernibles. *Analysis* **66**(4), 303–315.
- Ketland, J. (2011). Identity and indiscernibility. *Review of Symbolic Logic* **4**(2), 171–185.
- Ladyman, J. (2005). Mathematical structuralism and the identity of indiscernibles. *Analysis* **65**(287), 218–221.
- Ladyman, J. (2007). Scientific Structuralism: On the Identity and Diversity of Objects in a Structure. *Proceedings of the Aristotelian Society* **81**(1), 23–43.
- Leitgeb, H., & Ladyman, J. (2008). Criteria of Identity and Structuralist Ontology. *Philosophia Mathematica* **16**(3), 388–396.
- Lowe, E. (2003). Individuation. In Loux, M. & Zimmerman, D., editors, *Oxford Handbook of Metaphysics*, pp. 75–95. Oxford: Oxford University Press.
- MacBride, F. (2006). What Constitutes the Numerical Diversity of Mathematical Objects? *Analysis* **66**, 63–69.
- Muller, F. A., & Saunders, S. (2008). Discerning Fermions. *British Journal for the Philosophy of Science* **59**(3), 499–548.
- Muller, F. A., & Seevinck, M. P. (2009). Discerning Elementary Particles. *Philosophy of Science* **76**, 179–200.
- Quine, W. (1976). Grades of Discriminability. *Journal of Philosophy* **73**(5), 113–116.
- Russell, B. (1911). On the relation of particulars and universals. *Proceedings of the Aristotelian Society* **12**, 1–24. Repr. in his *Logic and Knowledge*, ed. R. C. Marsh, London: George Allen & Unwin, 1956.
- Saunders, S. (2003). Physics and Leibniz’s Principles. In Brading, K. & Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press.
- Saunders, S. (2006). Are Quantum Particles Objects? *Analysis* **66**, 52–63.
- Shapiro, S. (2006). Structure and Identity. In MacBride, F., editor, *Identity and Modality*, pp. 109–145. Oxford: Clarendon.

26 JAMES LADYMAN, ØYSTEIN LINNEBO, AND RICHARD PETTIGREW

Department of Philosophy,
University of Bristol
9 Woodland Road, Bristol. BS8 1TB, UK
E-mail: James.Ladyman@bristol.ac.uk
E-mail: Richard.Pettigrew@bristol.ac.uk

Department of Philosophy
Birkbeck College,
Malet Street, London. WC1E 7HX, UK
E-mail: O.Linnebo@bbk.ac.uk