Philosophers' Imprint

VOLUME 23, NO. 1
APRIL 2023

THE METAPHYSICS OF OPACITY

Catharine Diehl, Beau Madison Mount

Leiden University
King's College London

1. Introduction

Leibniz's Law is the principle that a true identity claim involving genuine singular terms licenses substitution of the term on one side of the identity sign for the term on the other side *salva veritate*. It is *prima facie* a fundamental logical principle, present in all classical first- and second-order systems with an identity symbol.[1] Recently, however, two sophisticated challenges to Leibniz's Law have emerged — one due to Michael Caie, Jeremy Goodman, and Harvey Lederman (2020) (henceforth CGL), the other due to Andrew Bacon and Jeffrey Russell (2019) (henceforth BR). Motivated by putative ordinary-language counterexamples, BR and CGL engage in logical revision: BR restrict quantifier rules, whereas CGL develop a system in which identity is coarser-grained than other definable equivalence relations. In this paper, we argue that their proposals are incompatible with the widely held assumption that the world contains individuals; moreover, if identity and a finest-grained equivalence relation are both present, there is strong theoretical pressure for them to coincide. The first part of our paper sets out these arguments; the second proposes a metaphysical picture — a version of stuff ontology — on which the rejection of individualism makes sense; finally, we provide a model theory that demonstrates the coherence of this picture.

That there is a connection among identity, individualism, and Leibniz's Law should be unsurprising: Leibniz's Law appears to capture the core theoretical role of identity precisely because it accounts for a thing's being *one* thing, regardless of how many names refer to it. Moreover, even in languages without an identity symbol, an implicit mastery of the notion of identity is required in order to make sense of quantification: as Hawthorne (2003) remarks, to grasp the difference between the truth conditions of $\exists x \exists y(Fx \wedge Gy)$ and those of $\exists x(Fx \wedge Gx)$, one must grasp the difference between one thing and two things.

For our purposes, we shall take the following to be the canonical

1. See, for example, Church (1956).

form of Leibniz's Law:

$$\mathbf{t}_1 = \mathbf{t}_2 \rightarrow (\phi(\mathbf{t}_1) \leftrightarrow \phi(\mathbf{t}_2)), \tag{LL}$$

where $\mathbf{t}_1$ and $\mathbf{t}_2$ are terms of type e (names, variables, expressions for the value of functions, or the like) and $\phi(\mathbf{t}_1)$ and $\phi(\mathbf{t}_2)$, as usual, represent the result of uniform capture-free substitution by the displayed term.

Nonetheless, violations of (LL) appear to abound. The most prominent cases of such putative failures include belief and attitude reports: it seems, for instance, that we cannot infer from

(1) George Sand = Amantine Dupin

and

(2) George Sand is believed by Madeline to be the author of the novel she is reading.

to

(3) Amantine Dupin is believed by Madeline to be the author of the novel she is reading.

Representing (1) as $s = d$, (2) as Bs, and (3) as Bd, in a case where the inference fails, we have

(4) $s = d \wedge Bs \wedge \neg Bd$,

a violation of (LL). We term putative counterexamples to (LL) such as (4) cases of *opacity*.

Philosophers have also argued that (LL) fails in the case of simple sentence pairs in a wide range of instances. For one example, consider the pair

(5) Clark Kent walked into the phone booth and Superman walked out

and

(6) Superman walked into the phone booth and Clark Kent walked out.

Prima facie, when applied to a single situation, these appear to differ in truth value (Saul, 2007), even though Superman is identical to Clark Kent — thus yielding a candidate counterexample to (LL).

Leibniz's Law also comes under pressure from the ease with which it can be used to proliferate objects in metaphysics. It is frequently employed to argue for the non-identity of things that naive intuition might take to be identical (Magidor, 2011). Consider, for example, the following case from Fine (2003):

(7) That statue is Romanesque

and

(8) That clay is not Romanesque,

where the statue and clay occupy the same spacetime region, which is ostended to introduce the demonstratives. Fine urges us to conclude from this and similar arguments that the statue is not the clay. Those who eschew multiply occupied regions of spacetime must object to this use of (LL). But even statue/clay dualists may baulk at proliferating entities in cases of role-relative predication, which arise because a single entity can occupy two different roles. For instance, the Chief Justice of the United States is empowered to issue writs of habeas corpus, but the Chancellor of the Smithsonian Institution is not empowered to issue writs of habeas corpus. Now, it happens that the same man — John Roberts — plays the role both of Chief Justice and of Chancellor of the Smithsonian. Thus, it seems that:

(9) Justice Roberts is empowered to issue writs of habeas corpus

and

(10) Chancellor Roberts is not empowered to issue writs of habeas corpus.

Nevertheless, it would certainly be ontologically profligate to conclude that Roberts-the-Justice is not identical to Roberts-the-Chancellor.[2]

One might worry that the expressions in (7)–(8) and (9)–(10) are not genuine singular terms. Although demonstratives, such as 'this statue' or 'this clay' are often taken to be directly referential, with their reference fixed by ostension (Kaplan, 1989), this is not universally accepted. Moreover, titles, such as 'Justice' and 'Chancellor' seem to occupy an intermediate place between ordinary bare proper names and descriptions. But, at the cost of introducing slightly more backstory, it would be straightforward to modify these examples to use ordinary language proper names. We could introduce proper names for the statue and clay by ostension,[3] and we can imagine a swearing-in ceremony's including baptism with a new first name for each role.[4] Once these names are established, sentence pairs with names corresponding to (7)–(8) and (9)–(10) would display the same behaviour as our examples.[5]

The tension between the theoretical importance of (LL) and the pervasiveness of intuitively compelling counterexamples to it has divided responses into two camps.[6]

Members of the first camp employ a variety of strategies to explain away the apparent counterexamples. For instance, Fregeans about proper names will argue that, in certain contexts, names do not denote their referents but rather their senses. Fregeans then use this to explain away apparent failures of substitutivity in such contexts: when the underlying logical form, rather than the surface structure, of the sentence is spelled out, there is no true failure (Frege, 1892; Church, 1951a, 1951b). Similarly, on Bertrand Russell's view, ordinary-language proper names are not logically proper names at all but rather disguised descriptions, which are in turn abbreviations of quantified expressions, so natural-language apparent counterexamples do not threaten Leibniz's Law (Russell, 1918, 1919a, 1919b, 1919c). Others, such as Nathan Salmon (1986) and Scott Soames (1989), argue that our initial judgments about truth values are mistaken and result from confusing the claim with a nearby but quasi-quotational statement.

Those in the second camp, by contrast, suggest modifying Leibniz's Law, either by restricting the range of contexts to which it applies or jettisoning it entirely. For instance, Carnap (1956) restricts Leibniz's Law to extensional contexts (those not containing modal operators) and introduces a separate principle for intensional formulas: $\Box \mathbf{t}_1 = \mathbf{t}_2 \to (\phi(\mathbf{t}_1) \leftrightarrow \phi(\mathbf{t}_2))$.[7] But, until recently, there have been few precedents for systems allowing failures of Leibniz's Law for singular terms in formulas of arbitrary structure (as appears to occur in natural-language cases such as (1)–(10)).

Recently, however, CGL and BR have sought to address this lacuna by providing sophisticated formal discussions of systems in which (LL) fails. Both approaches are situated within a higher-order language in which classical identity — that is, identity that obeys (LL) — can be formally defined as *Leibniz equivalence*, which we will denote by $\approx$:

$$\mathbf{t}_1 \approx \mathbf{t}_2 =_{\mathrm{Df.}} \forall X(X\mathbf{t}_1 \leftrightarrow X\mathbf{t}_2). \qquad \text{(LEquiv)}$$

The central difference between BR's and CGL's approaches is that BR accept that (LEquiv) characterises genuine identity, while CGL dispute this. In order to maintain the equation of identity with Leibniz equivalence, BR modify both first- and higher-order principles of Universal

2. On the related phenomenon of role-relative predication, see Fine (1982), Landman (1989a, 1989b), Szabó (2003), and Loets (2021).
3. See Gibbard (1975) for such a strategy.
4. Compare, for example, Prince Albert's becoming George VI on his accession in 1936.
5. We thank an anonymous referee for raising this point.
6. See Magidor (2011) for a slightly different taxonomy of responses, as well as more detailed discussion.

7. Thanks to an anonymous reviewer for reminding us of Russell's and Carnap's approaches. Carnap, in our view, codifies the Fregean thought that in certain contexts singular terms denote concepts or senses. His approach is a case of the second strategy, however, because occurrences in intensional contexts count as genuine uses of singular terms.

Instantiation,

$$\forall x\, \phi(x) \to \phi(\mathbf{t}), \tag{UI1}$$
$$\forall X\, \phi(X) \to \phi(\mathbf{T}) \tag{UI2}$$

(where $\mathbf{t}$ is a first-order and $\mathbf{T}$ a higher-order term). Doing so permits violations of (LL) without allowing any instance of $\mathbf{t}_1 = \mathbf{t}_2 \wedge \exists X(X\mathbf{t}_1 \wedge \neg X\mathbf{t}_2)$.[8] In effect, their strategy is an implementation of positive free logic, with the quantifier restricted from the perspective of the metalanguage. Because they do not allow unrestricted second-order generalization, they can accept a failure of the inference in (1)–(3) as a true violation of (LL) and regiment it using (4) without accepting that *there is* a property by which Sand and Dupin differ.

CGL adopt a different strategy: they permit cases in which $\mathbf{t}_1$ and $\mathbf{t}_2$ flank the identity sign but nonetheless fail to be Leibniz-equivalent. The distinctive commitment of this view is what we term SUBMAXIMAL FINENESS:

(SUBMAXIMAL FINENESS) There exists a unique finest-grained equivalence relation, but identity is not it: genuine identity is a coarser-grained relation.

This allows them to retain full (UI1–2). They call their view *Classical Opacity*, because it preserves the core of classical higher-order logic, while allowing violations of (LL) (Caie et al., 2020, 532). For clarity, we term the coarser-grained relation denoted by the identity sign on their view *weak associative equivalence* (WA-equivalence), denoted by $\equiv$: we can thus express the distinctive claim of the Classical Opacicist as the thesis that $\equiv$, not $\approx$, expresses genuine identity.

BR and CGL offer contrasting justifications for their views concerning whether genuine identity must track Leibniz equivalence. BR argue that, if we allow genuine identity to float free of Leibniz equivalence, we will lose our grip on the subject matter (Bacon and Russell, 2019, p. 84).[9] CGL (Caie et al., 2020, p. 540) respond by claiming that BR err in according too much evidential force to a theoretical desideratum—the requirement that identity be the finest-grained equivalence relation—rather than the common-sense data (which they take to include, for example, the claim that Hesperus is identical to Phosphorus and the claim that they vary in properties, in that one was known by the ancients to rise in the evening and the other was not). Because of these different desiderata, CGL and BR offer contrasting modifications to the classical picture: CGL argue that the identity relation is coarser-grained than Leibniz equivalence, whereas BR restrict the quantifier (relative to the 'external' realm of objects for which there are singular terms).

Neither BR nor CGL, however, offer any account of the underlying metaphysics of their systems for opacity—the way the world would have to be for their accounts to be right. They take the apparent linguistic counterexamples as data and develop candidate logics to accommodate them, building models to demonstrate consistency, but they do not explain the intended interpretations of their systems in metaphysical terms.

In this paper, we show that there is a tension between these systems incorporating opacity and widely held, intuitively compelling metaphysical principles. In particular, we show that the core commitments of a metaphysics of individuals conflict with accepting a case of (LL) violation, such as (4). First, in §2, we address the objection that (1)–(3) and similar cases are merely superficial and can be explained away metalinguistically. We then raise the objection that acceptance of (4) conflicts with adherence to a metaphysics of individuals. In order to avoid this result, one must adopt a highly revisionary ontology: either an aspect-based metaphysics, such as that proposed by Donald Baxter (2018) or an ontology on which individuals do not play a primary role.

8. In addition, they also reject any instances of $\exists x \exists y(x = y \wedge \neg(\phi(x) \leftrightarrow \phi(y)))$ and $\exists x \exists y(x = y \wedge \exists X(Xx \wedge \neg Xy))$. We can take the existential quantifier here as a defined abbreviation; were it primitive, modified axioms of Existential Generalization would also be required.

9. Compare Williamson (2002).

We then show that Baxter's view cannot be used to solve the problem.

We next raise a metalogical objection to CGL's attempt to pry apart Leibniz equivalence and identity: we argue that Submaximal Fineness is incompatible with standard accounts of counting and the logicality of identity. We concentrate on CGL here, because their denial that identity is Leibniz equivalence is a particularly novel and radical view and because their acceptance of classical quantifier rules makes their position particularly attractive to those who take logical conservatism as a definitive advantage. After concluding our argument against existing ways of denying (LL), we develop in §4 an alternative, positive proposal that accommodates violations of Leibniz's Law within a metaphysically well-motivated system. On our view, in order to deny (LL) intelligibly, one must both downgrade the status of individuals and reject Submaximal Fineness by positing an infinite hierarchy of ever-finer-grained equivalence relations.

We propose an ontology of stuff that would motivate this combination of views. We do not mount a thoroughgoing case for this position, but we show that it provides a coherent metaphysical story. Although we do not develop the idea here, it is worth noting that this metaphysics has a claim to capture the intuition behind at least two historically prominent views — Anaxagoras's "theory of extreme mixture" (Marmodoro, 2017) and Leibniz's account of the phenomenal world as an infinitely variegated, infinitely divisible multitude.[10]

Finally, in the Appendix, we give a model theory that corresponds heuristically to a stuff metaphysics of this type.

2. Metalinguistic Leibniz's Law

Our motivating cases, such as (1)–(3), might seem merely superficial and thus unworthy of extensive logical and metaphysical attention. After all, in the absence of any restrictions on allowable values of ϕ, $\mathbf{t}_1$, and $\mathbf{t}_2$, natural-language (LL) can fail in explicitly or implicitly quotational contexts, but these appear to be tangential, cheap cases. In this section, we show that a metalinguistic version of (LL), developed by Timothy Williamson (2002) to exclude such cases, does not provide an easy way to avoid CGL's and BR's challenges. CGL, in particular, are committed to genuinely identical objects differing in properties in the only sense of 'property' that they countenance (that is, as the value of the term generated by λ-abstraction from any open formula).

Let's consider Williamson's (2002, pp. 288–89) formulation, which uses sameness of semantic value on minimal pairs of variable assignments in place of intersubstitutability *salva veritate*. A (first-order) assignment for a language is a function mapping the language's variables to objects; we use α, β, and variants for assignments. We say that β is an $^{a_2}/_{a_1}$-variant of α just in case α and β differ only in that there exists one variable $\mathbf{x}$ such that $\alpha(\mathbf{x}) = a_1$ and $\beta(\mathbf{x}) = a_2$.

Williamson's principle (slightly paraphrased) is:

(CLL) For all a_1, a_2, if there exist assignments α and β such that β is an $^{a_1}/_{a_2}$-variant of α and a formula ϕ is true on α but not true on β, then $a_1 \neq a_2$.

Quotational cases do not provide even arguable counterexamples to (CLL). For instance, ' "George Sand" was chosen to trade on the gender expectations of nineteenth-century readers' is true but the result of substituting 'Amantine Dupin' for 'George Sand' in this sentence is false. But there is no way to turn the sentence ' "George Sand" was chosen to trade on the gender expectations of nineteenth-century readers' into an open formula 'x was chosen to trade on the gender expectations of nineteenth-century readers' whose truth-value differs on a minimal assignment pair where one assignment takes x to George Sand and the other takes x to Amantine Dupin. (Of course, the truth-value would differ on a minimal assignment pair involving the expressions 'George Sand' and 'Amantine Dupin', but no one has ever claimed that the ten-letter expression 'George Sand' is identical to the fourteen-letter

10. For passages where Anaxagoras discusses infinite divisibility, see Diels and Kranz (1951, fr. B3, B6); on universal mixture, see Diels and Kranz (1951, fr. B17). For Leibniz on the phenomenal realm, see 'Primary Truths' (1686?), trans. in Ariew and Garber (1989, p. 30) and *Monadology* §§69–70 (1714), trans. in Ariew and Garber (1989, p. 220).

expression 'Amantine Dupin'.)[11]

As we will now show, CGL are committed to the existence of violations of (CLL) as well as (LL). To see this, note that they accept violations of the quantified principle

$$\forall x \forall y (x \equiv y \rightarrow \forall X (Xx \leftrightarrow Xy)). \tag{ULL}$$

For instance, they take (4) to be a witness to the negation of (ULL). But any such case will provide a counterexample to (CLL) as well, for the truth of the universally quantified formula is defined in terms of the satisfaction of its instances on all assignments. Let α and β differ (if at all) only in that α assigns Dupin to x and Sand to y, whereas β assigns Sand to both. Consider the open formula 'x is believed by Madeline to be the author of the novel she is reading'; on CGL's view, this is false on α and true on β. It is unavailing to object that intensional predicates cannot be treated by the usual apparatus of assignments: CGL also are forced to accept that $x \approx y$, i.e. $\forall X(Xx \leftrightarrow Xy)$, which is free from all intensional vocabulary, is false on α and true on β.

This argument shows that Classical Opacicism violates even the metalinguistic (CLL). This creates a problem for the Classical Opacicist, for (CLL) appears to be what Williamson terms "a special case of a trivial theorem of classical mathematics":

(CMT) Let f and f^* be functions on a domain D such that for some $d \in D$, $f(e) = f^*(e)$ whenever $d \neq e$. Suppose that some object x has R to f but not to f^*. Then, $f(d) \neq f^*(d)$. (Williamson, 2002, p. 289).

This theorem is indeed *classically* trivial, but it can only be derived by reasoning that CGL would reject. For it is proved from an extensionality assumption for functions:

(11) If, for all $e \in D$, $f(e) = f^*(e)$, then $f = f^*$,

and the principle

(12) If there exists an object x that has R to f but not to f^*, then $f \neq f^*$.

But (12) is itself a contraposed instance of Leibniz's Law that CGL are obliged to reject, when $=$ is understood as $\equiv$, for $f \equiv f' \rightarrow (Raf \leftrightarrow Raf')$, and more generally $f \equiv f' \rightarrow (Ff \leftrightarrow Ff')$, has false instances in their system. Thus, CGL must reject certain statements of classical mathematics in full generality — and they have the resources, on their own terms, to motivate that rejection. But they can offer restricted versions of these statements that will hold when R is *transparent* — that is, when it does not induce a case of opacity. They suggest some principles that would allow canonical logical vocabulary to be transparent and they might likewise provide a restricted form of (CMT) that would hold for mathematical — i.e., purely extensional — functions (Caie et al., 2020, pp. 534, 536–40).

As for BR, the situation is more complicated because they accept (ULL) — it is a trivial variation of what they term LL(xy) — even though they countenance what seem to be counterexamples to it (Bacon and Russell, 2019, p. 101). (Rejecting unrestricted universal generalization allows them to hold these two positions consistently.) But they will still view (1)–(3) as providing a counterexample to a schematic version of (CLL),

11. Parallel reasoning also holds in 'quasi'-quotational cases that include covertly semantic vocabulary. Consider, for instance, Quine's (1961) example of the sentence pairs 'Giorgione is so-called because of his size' and 'Barbarelli is so-called because of his size', which intuitively differ in truth value. Nevertheless, they do not provide a counterexample to (CLL). It is clear that, for a singular term $\mathbf{t}$ of any type, the formula $\ulcorner \mathbf{t}$ is so-called because of $\phi \urcorner$ is true relative to an assignment α just in case $\ulcorner \mathbf{t}$ is called "$\mathbf{t}$" because of $\phi \urcorner$ is true relative to α (cf. Quine, 1961, p. 140). In this case, we form the open formula 'x is so-called because of his size' and apply the equivalence to obtain 'x is called "x" because of his size'. Let α be an assignment taking 'x' to Barbarelli and every other variable to zero, β an assignment taking 'x' to Giorgioni and every other variable to zero; 'x is called "x" because of his size' is false on both α and β, for the painter is not called the twenty-fourth letter of the alphabet at all; thus no non-identity claim follows. Compare Williamson's discussion of an artificial 'definitely' operator (2002, pp. 286–87).

(CLL′) If β is a $\mathfrak{t}_1/\mathfrak{t}_2$-variant of α and a formula ϕ is true on α but not true on β, then $\mathbf{t}_1 \neq \mathbf{t}_2$,

since the open formula Bx will be true on an assignment taking x to George Sand but not on one taking x to Dupin. They will thus have to reject the corresponding schematic instance of (CMT).

Since ordinary mathematics assumes full classical logic, and thus unrestricted universal instantiation, no distinction is standardly made between (CMT) and its schematic analogue. If BR wish to avoid restricting (CMT) in the way that CGL do, they will need to introduce a distinction exogenous to mathematical practice.

The preceding discussion shows that neither CGL's nor BR's systems can be accommodated within a conciliatory strategy on which their motivating cases merely involve violations of the letter but not the spirit of Leibniz's Law. The metaphysical questions cannot be evaded.

3. Two Challenges for Metaphysical Opacity

We will now develop our two main negative arguments: the first is directed against the combination of (LL) violations with an intuitively appealing metaphysical picture, the second against views that claim that there is a finest-grained equivalence relation but identity is not it. In the first argument, we set out four principles about the relationship between language (on the assumption of reasonably perspicuous logical form) and the world that, we claim, are unavoidable consequences of the individualist picture of objects and properties. We challenge both CGL and BR to specify an alternative metaphysical picture on which the rejection of one of these principles makes sense.[12] The second argument applies to CGL in particular: we show that any system on which there is a finest-grained equivalence relation that is not identity must reject the logicality of identity and of the cardinality quantifiers on the standard criteria. We thus conclude that the only reasonable choice for the individualist is to accept a finest-grained equivalence relation that is identity and with it (LL); the alternative, as we discuss later, requires rejecting the individualist picture and the existence of any finest-grained equivalence relation.

12. The basic idea behind the argument is not novel: see (Wiggins, 1967, pp. 4–5) for a precursor.

3.1 The S-Argument

We begin our argument by assuming, for reductio, that a violation of (LL) such as that set out in (4) holds:

(A1) $s = d$;
(A2) Bs;
(A3) $\neg Bd$.

We now introduce the following theses about the relations between true sentences and the objects and properties in virtue of which they obtain.[13]

We will term these complexes of objects and properties *obtaining states of affairs* (SOAs), but we intend the terminology to involve only a minimal commitment to the existence of a world with 'ontological structure'. We adopt Jason Turner's characterization of such structure as akin to a pegboard:

> Such descriptions of the world implicitly suppose that it has a certain sort of structure—an *ontological structure*. Ontological structure is the sort of structure we could adequately represent with a pegboard and rubber bands. The pegs represent things, and the rubber bands represent ways these things are and are interrelated.
>
> To say 'Bertrand thought about language', for instance, is to hang the *thought about language* rubber band on the peg labeled 'Bertrand'. And to say 'Some logicist admired every philosopher

13. We do not intend to endorse any strong version of truthmaker theory; we merely require that there is a subclass of true sentences, among them (A1)–(A3), whose truth obtains in virtue of the way objects and properties stand. We take this to be a minimal commitment of any realist version of individualism.

who didn't notice the inconsistency in Basic Law V' is to say that, somewhere on the pegboard, there is a peg which (a) has a *logicist* rubber band hanging on it, and (b) has an *admires* rubber band stretching from it to each of the pegs with the *didn't notice the inconsistency in Basic Law V* band on. (Turner, 2011, p. 5)[14]

We use 'state of affairs' for the structure, whatever it is, that is picked out in Turner's metaphor by the pegs and rubber bands: however one wishes to interpret the pegboard model, one can interpret talk of SOAs accordingly. In particular, the minimal commitment involved in SOAs is compatible with a variety of detailed accounts about the nature of ontological structure. Someone who views SOAs as structured entities akin to Russellian propositions can accept these premisses, but so can someone who construes talk of SOAs as a mere *façon de parler* for non-nominalised facts about objects and properties.[15]

We take the following four platitudes to be fundamental principles governing the interaction between language and world on the individualist's picture: (1) true monadic atomic predications obtain in virtue of an individual's being a certain way; (2) true negations of monadic atomic predications obtain in virtue of an individual's being a certain way (contrary to that involved in the unnegated sentence); (3) if a simple identity statement holds, then the same object is picked out by each of the terms flanking the identity connective; (4) a single individual cannot both be a certain way and be a contrary way. The individualist picture is committed to the 'pegboard' of individuals and ways for them to be. The platitudes merely reflect our ability to use a language in which those individuals are picked out in ways compatible with ontological structure, where we use a mapping that conforms with the usual syntactic categories and meanings of the logical connectives. We will show that properly formalised versions of these platitudes conflict with (A1)–(A3).[16]

The principles are:

(S1) If $\Phi\mathbf{t}$, then there obtains an SOA σ whose only constituents are an object x denoted by $\mathbf{t}$ and a property P designated by $\ulcorner(\lambda x \,.\, \Phi x)\urcorner$;

(S2) If $\neg\Phi\mathbf{t}$, then there obtains an SOA σ whose only constituents are an object x denoted by $\mathbf{t}$ and a property $\bar{P}$ (the complement of P) designated by $\ulcorner(\lambda x \,.\, \neg\Phi x)\urcorner$;

(S3) If $\mathbf{t}_1 = \mathbf{t}_2$ and there are SOAs σ_1 and σ_2 with constituents denoted by $\mathbf{t}_1$ and $\mathbf{t}_2$, respectively, then there exists a unique object x such that x is denoted by both $\mathbf{t}_1$ and $\mathbf{t}_2$ and is a constituent of σ_1 and σ_2; and

(COMP) For all x and P, it is not the case that there obtain SOAs σ_1 and σ_2 such that the only constituents of σ_1 are x and P and the only constituents of σ_2 are x and $\bar{P}$.

The first principle, (S1), operates on the assumption that there are

14. For an earlier use of the pegboard model, see Armstrong (1989, pp. 64–65).

15. One reason to construe talk of SOAs periphrastically is provided by a cardinality problem: if (1) properties are as abundant as classes of objects, (2) whenever an object has a property, a corresponding SOA obtains, and (3) SOAs are individuated as finely as object-property tuples, then (4) there will be more SOAs than objects. So SOAs cannot be objects. In this case, we could, for instance, associate objects with their singleton properties and paraphrase SOAs as third-order entities, coding an object-property tuple by a property of properties. Higher-orderisation of this type can also assuage any nominalistic qualms about reifying properties. Of course, someone who believes only in sparse properties, rejecting (1), can consider SOAs to be objects and need not resort to higher-orderisation. We take no stance on this question.

16. Our regimentations of the platitudes in (S1)–(S3) and (COMP) use the idioms of 'denoting' and 'designating' to express relationships between terms and components of SOAs. On standard semantic assumptions, these relationships will be fairly straightforward: the term 'Jones' in 'Jones plays the piano' will denote Jones, and the term '... plays the piano' will designate the property of playing the piano, and these entities will be components of the SOA that makes it the case that Jones plays the piano. But we wish to leave room for more complex semantic accounts, on which the connection between the truth-conditions of the sentence and the truthmaker is indirect: all that we require is that there be some relation through which terms correspond systematically to constituents of SOAs in such a way that an SOA's obtaining underpins a true sentence's being true. The reader who thinks this relation does not merit such names is welcome to read 'corresponds to' or some more neutral locution instead.

objects and properties that make atomic sentences true and claims that for any such true sentence there is some way the world is involving an appropriate object and property. We do not assume that there is a *unique* such object and property pair. In particular, we do not assume that SOAs are individuated extensionally: it is compatible with our assumptions that there be two distinct obtaining states of affairs with the same object and property as constituents (for instance, if SOAs can manifest several different modes of combination). Nor do we require the converse assumption that if σ_1 and σ_2 are identical SOAs, then they will have the same objects and properties as constituents (itself an instance of (LL)).

The motivating idea behind (S2) is that for F not to hold of a is for a to have a negative property — the property of not being F. Though one might have metaphysical scruples about negative properties, they do seem to be countenanced by natural language: I have the property of not being in Hong Kong if and only if it is not the case that I am in Hong Kong, and so on. (For those who continue to have misgivings about negative properties, we will soon offer an alternative argument that does without such properties.) (S3) just requires that if $\mathbf{t}_1 = \mathbf{t}_2$ and x is the object denoted by $\mathbf{t}_1$, then x is also the object denoted by $\mathbf{t}_2$. This is, to our mind, a truism about the relation between objects and names. Finally, (COMP) is a metaphysical consistency constraint: the world does not contain obtaining SOAs that conjoin a single individual to incompatible properties.

As noted, we take (S1)–(S3) and (COMP) to be formalisations of platitudes about the language-world relation on an individualistic picture, rather than deliverances of any particular logical theory. It might be objected, however, that these principles covertly presuppose Leibniz's Law or similar principles of classical logic that CGL and BR reject, and thus that their use in this context is question-begging.[17] On our view, this objection gets the order of explanation wrong. We do not first adopt a collection of axioms and rules of inference and then work our way towards the metaphysical picture of (S1)–(S3) and (COMP); rather, we choose our principles to reflect a framework to which the individualist is *independently* committed. We can compare Kripke's (1971) intuition-based arguments for the necessity of identity: although the key premises *could* be regimented as instances of the necessity principle itself, which would render the arguments question-begging in form, there is no reason to force them into this mould. The intuitive picture is intended to stand on its own and should be allowed to do so. Likewise, we claim merely that if one thinks that the world contains individualistic ontological structure and holds that this structure can be captured in a suitable language, then one will be thereby committed to these principles, for they merely delimit how the objects and ways of being to which one is committed make corresponding atomic sentences true.

But (A1)–(A3), (S1)–(S3), and (COMP) are jointly inconsistent.

Proof. By (A2) and (S1), there obtains an SOA σ_1 whose only constituents are an object x_1 denoted by 's' and a property P designated by '$(\lambda x \,.\, Bx)$'. By (A3) and (S2), there obtains an SOA σ_2 whose only constituents are an object x_2 denoted by 'd' and a property $\bar{P}$ designated by '$(\lambda x \,.\, \neg Bx)$'. By (A1) and (S3), there is a single object x denoted by both 's' and 'd' and figuring in both σ_1 and σ_2; thus there is an x such that σ_1 has as its only constituents x and P and σ_2 has as its only constituents x and $\bar{P}$. But, by (COMP), this is impossible. □

This result generalises to all putative counterexamples to (LL): (S1)–(S3) and (COMP) entail (LL). We will now consider several ways in which opacicists might resist the argument based on (S1)–(S3).

First of all, it might be objected that, in demonstrating this incompatibility, we are interpreting the reuse of the variable x in (S3) *classically*; a hard-core follower of CGL might insist that (S3) can only be understood as:

(S3*) If $\mathbf{t}_1 = \mathbf{t}_2$ and there are SOAs σ_1 and σ_2 with constituents denoted by $\mathbf{t}_1$ and $\mathbf{t}_2$, then there exists a unique object x such

17. We thank an anonymous referee for raising this worry, as well as the issues about extensionality raised earlier.

that some $y \equiv x$ is denoted by $\mathbf{t}_1$ and some $z \equiv x$ is denoted by $\mathbf{t}_2$ and y is a constituent of σ_1 and z is a constituent of σ_2,

on the grounds that every repetition of a variable is to be cashed out in terms of what CGL claim to be identity — i.e., WA-equivalence.

But this refusal to countenance any vocabulary not cashed out in terms of $\equiv$ — ostrichism about opacity — is unavailing; if (S3) is thus modified, then the variables in (COMP) should be modified in the parallel way, yielding

(COMP*) For all x and P, it is not the case that there obtain SOAs σ_1 and σ_2 such that the only constituents of σ_1 are an object y such that $y \equiv x$ and P and the only constituents of σ_2 are an object z such that $z \equiv x$ and $\bar{P}$,

which suffices to recover the contradiction with (S1), (S2), and (S3*).

If the classical opacicist is determined to adopt ostrichism, he is better off sticking to his convictions and denying (COMP) outright, on the grounds that the right notion of identity, WA-equivalence, simply does not preclude worldly incompatibility. But this position strikes us as metaphysically unintelligible: (COMP) does not invoke identity explicitly, but merely concerns the most basic aspects of objects and properties. To deny it is, in effect, to endorse a form of worldly dialetheism. Moreover, even without (COMP), unacceptable worldly dialetheism results from (A1)–(A3) and (S1)–(S3) in the presence of a plausible aggregation principle:

(AGG) For all x, P_1, P_2, if there obtain SOAs σ_1 containing only x and P_1 and σ_2 containing only x and P_2, then there obtains an SOA σ_3 containing x and $(P_1 \& P_2)$.

The principle

(NC) For all P, there obtains no SOA containing $(P \& \bar{P})$

is, if anything, even harder to deny than (COMP), and (A1)–(A3), (S1)–(S3), (AGG), and (NC) are inconsistent.

The opacicist might also try flat-footedly denying (S3) without reinterpreting repeated variables. This option is the one most in keeping with CGL's model theory, on which the 'internal' identity relation for a type does not correspond to the 'external' identity of the model. It is also a result suggested by CGL's view that identity does not have to cut as finely as Leibniz equivalence. We will argue against this directly in the next section, but we are already in a position to note that allowing for a finer-grained identity relation in the metalanguage raises the suspicion that it is this metalinguistic identity that plays the true identity role. Of course, it is open to CGL to say that the metalinguistic 'identity' to which they appeal is merely heuristic and thus not in competition with object-language identity. But this would require us to have an independent purchase on why their object-language identity is the genuine article, and they have provided no such elucidation.

It is worth noting that BR would also be tempted to deny (S3), but for different reasons: if $s = d$, as in (A1), then since they assume a disquotational account of denotation (Bacon and Russell, 2019, pp. 88–90), they should be prepared to accept that there is a state of affairs σ_1 containing an object denoted by '*s*' and a state of affairs σ_2 containing an object denoted by '*d*'. But they would deny the quantified claim that there is any one thing that figures in both states of affairs.

In the framework of SOAs, however, the plausibility of this strategy of logical revision comes under pressure.[18] The motivation for rejecting (UI1–2) derives from features of natural languages, but our principles for states of affairs derive from the metaphysical intuitions summed up in the 'pegboard' picture. To make this vivid, we can imagine supplementing English with 'Lagadonian' terms for constituents of SOAs (with the relevant objects and properties standing for themselves). Sentences in such a language would be merely sequences consisting of

18. The disquotationalist might reject SOAs — as an anonymous referee has suggested — as extraneous ontological scaffolding, much as they reject facts. Our SOAs, however, merely capture the commitments of the 'pegboard' model of ontological structure. If the disquotationalist adheres to this model, then she is stuck with SOAs (in the modest sense explained in fn. 16).

the disaggregated components of the states of affairs: in such a case, it is difficult to deny that *one and the same thing* features in both σ_1 and σ_2. This strategy does not require a *universal* Lagadonian language, with everything serving as a name for itself (which might reasonably be resisted on cardinality grounds): all that is required is that, for any given case, some appropriate expansion of English can be generated.[19]

BR remain free to reject the pegboard model and the Lagadonian expansion that accompanies it: for instance, they might hold that objects are insufficiently individuated to serve as their own names. If they do so, they may have principled reasons to deny (S3). But the adherent of the pegboard picture should accept this modest restricted Lagadonian expansion: after all, the objects — the 'pegs' according to the metaphor — should be sufficiently distinct to serve as their own names. Our central point is that individualist platitudes must give way if the S-argument is to be resisted. In §4, we sketch an alternative picture that would forgo these platitudes.

This leaves denial of (S1) and (S2) as live options. The only plausible motivation for denying (S2) but not (S1) is suspicion of negative properties. This would also motivate rejection of (COMP). But one who rejects (S2) and (COMP) on the grounds that $\bar{P}$, if it existed, would be a negative property, and no such things exist, should still accept the following modification of (COMP):

(COMP**) For all x, P_1 and P_2, if P_1 and P_2 are incompatible positive properties, then it is not the case that there obtain SOAs σ_1 and σ_2 such that σ_1 contains x and P_1 and σ_2 contains x and P_2.

If there exist violations of (LL) of the form

(A1*) $a = b$,
(A2*) Fa,
(A3*) Gb,

19. We thank an anonymous referee for raising this worry and related issues.

where '$(\lambda x\,.\,Fx)$' designates P_1 and '$(\lambda x\,.\,Gx)$' designates P_2, with P_1 and P_2 incompatible positive properties, then a contradiction can be derived using (S1) and (COMP**) alone.

Such an example can plausibly be found in the following case. Assume that the opacicist wishes to accept an (LL) violation in the case of coincident objects in order not to have to posit the distinctness of the statue Goliath and the clay Lumpl composing it. Suppose that the sort of clay from which Goliath is made is exceptionally sensitive to slight pressure: a misplaced touch will cause it to lose its shape entirely and become an amorphous ball. We have:

(13) Lumpl = Goliath,
(14) Goliath is fragile,
(15) Lumpl is robust.

If 'fragile' and 'robust' denote genuine incompatible positive properties — as is plausible on many theories of dispositions — then (13)–(15), (S1), and (COMP**) yield a contradiction.

Moreover, we can give a formulation of the S-argument that would be acceptable to a nominalist — a version that directly appeals to the incompatibility between a predicate's obtaining and its not obtaining, without requiring a metaphysical notion of either property complementation or the existence of incompatible positive properties. Objections to negative properties are entirely unavailing against this version:

(S1′) If $\Phi\mathbf{t}$, then there obtains an SOA σ such that σ contains an object x that is the denotation of $\mathbf{t}$ and all of σ other than x corresponds precisely to $\ulcorner(\lambda x.\Phi x)\urcorner$;
(S2′) If $\neg\Phi\mathbf{t}$, then there obtains an SOA σ such that σ contains an object x that is the denotation of $\mathbf{t}$ and all of σ other than x corresponds precisely to $\ulcorner(\lambda x.\neg\Phi x)\urcorner$;
(S3) If $\mathbf{t}_1 = \mathbf{t}_2$ and there are SOAs σ_1 and σ_2 with constituents denoted by $\mathbf{t}_1$ and $\mathbf{t}_2$, respectively, then there exists a unique object x such that x is denoted by both $\mathbf{t}_1$ and $\mathbf{t}_2$ and is a constituent of σ_1 and σ_2; and

(COMP′) It is not the case that there obtain SOAs σ_1 and σ_2 both containing some object x such that all of σ_1 other than x corresponds to $\ulcorner(\lambda x.\Phi x)\urcorner$ and all of σ_2 other than x corresponds to $\ulcorner(\lambda x.\neg\Phi x)\urcorner$.

(A1–A3), (S1′), (S2′), (S3), and (COMP′) yield a contradiction.

Thus, rejections of (S2) on the basis of qualms about negative properties are unavailing.

We now consider three ways of denying (S1): No Corresponding Property, Missing Components, and Extra Components. (Adherents of some of these strategies may also deny further principles, but we focus on (S1) for convenience.)

No Corresponding Property denies the claim that all predicates figuring in true atomic statements designate properties figuring in SOAs. (As an example, the defender of No Corresponding Property might say that broadly epistemic predicates fail to correspond to properties figuring in SOAs.) Missing Components claims that, for at least some instances of (S1), the denotatum of the term **t** does not figure in an SOA. Some who adopt this strategy might claim that only certain objects — perhaps the physically fundamental ones — figure in SOAs. Others might take a more radical approach and argue that there are no denotata figuring in SOAs, because SOAs do not contain individuals at all. Extra Components claims that SOAs contain additional material beyond simply objects and properties and that this somehow accommodates opacity.

We shall argue that only a radical version of Missing Components has any plausibility.

No Corresponding Property — denying (S1) on the ground that predicates in true atomic sentences that display opacity lack corresponding properties figuring in obtaining SOAs — has a number of unpalatable consequences. In order to block the argument based on (S1′), the objector must hold not only that there is no property corresponding to a term of the form $\ulcorner(\lambda x\,.\,\Phi x)\urcorner$, but that the semantic value of this expression makes no contribution to any obtaining SOA in virtue of which the sentence is true. One method to do this would be to deny that we could extract the component $\ulcorner(\lambda x\,.\,\Phi x)\urcorner$ from $\Phi\mathbf{t}$ — but to do this is, in essence, to reject β- and η-conversion, which CGL correctly note are standard components of type-theoretical frameworks that support classical reasoning (Caie et al., 2020, p. 529).[20]

An alternative version of the No Corresponding Property response would deny (S1) by rejecting outright the ideology of 'exact correspondence' used to formulate the contribution of $\ulcorner(\lambda x\,.\,\Phi x)\urcorner$ in minimal, non-property-theoretic terms. But one who does so owes us some explanation of what it is that makes true sentences involving opacity-inducing terms true: the mere claim that their truth is not anchored in obtaining SOAs does nothing to bear this explanatory burden. In the absence of such an explanation, No Corresponding Property remains unmotivated.

A moderate version of Missing Components — that some **t** in true, atomic predications denote objects figuring in SOAs, while others do not — might be motivated by the claim that only the fundamental objects of our perfected physical science figure in SOAs. This response, however, is unavailing as a general strategy to preserve opacity, because it seems we can have incompatible epistemic stances towards physical fundamentalia. To explain violations of (LL) involving conflicting epistemic attitudes towards fundementalia, the proponent of this strategy would have to turn to a limited version of No Corresponding Property and deny that epistemic predicates correspond to components of SOAs. But, as we have argued, the No Corresponding Property strategy is unpromising; moreover, combining Missing Components and No Corresponding Property in this fashion would introduce disunity into the account.

20. The advocate of sparse properties might be thought to possess principled reasons for rejecting β- and η-conversion for non-fundamental properties. There are, however, plausible candidate Leibniz's Law violations involving unimpeachably fundamental properties: the account of 'statue' as an aesthetic sortal term that motivates (7) and (8) also motivates the pair 'that statue is not matter' and 'that clay is matter'.

This leaves only a radical form of Missing Components — on which either singular terms only denote objects not figuring in SOAs or they never denote objects at all. The first of these options is extremely *ad hoc*: it would render the denotation of singular terms a pure epiphenomenon, unconnected to what makes sentences containing them true. The second, most radical stance — on which no singular term denotes a genuine object—lies behind the stuff metaphysics we shall defend in the next sections.

Extra Components — the option of denying (S1) on the grounds that the SOA corresponding to $\ulcorner \Phi \mathbf{t} \urcorner$ consists not just of an object and a property, but two objects and a two-place relation—seems initially appealing but is difficult to develop in a plausible fashion. If the claim is merely that there exists an unpronounced term for the missing object, then it does not seem that (A1)–(A3) are faithful formulations of their natural-language counterparts: the contradiction vanishes, but so does the case for genuine opacity. So the objector must have in mind an additional sort of object, beyond the one 'ostensibly' denoted by **t**, that figures in the SOA making a genuinely singulary predication of the form $\Phi\mathbf{t}$ true. What kind of object would this be? Presumably not the sense of the term **t**, or the syntactic object resulting from quoting the name **t**: to say this is to attempt to import a Fregean framework, on which the surface form of $\ulcorner \Phi \mathbf{t} \urcorner$ is deceptive, into a setting with a completely different motivation. Instead, presumably the object must be something of the same kind as the object denoted by **t**, intimately related to that object, which explains how it is that the predication is true under that guise. The obvious candidate for this is something like an object-under-an-aspect.

The most developed version of an aspect-based approach has been formulated by Baxter (2018). As we shall show, however, even a Baxterian strategy cannnot render Extra Components plausible. On Baxter's account, individual objects are identical to their aspects. For instance, Socrates is identical to Socrates-insofar-as-he-is-wise, as well as to Socrates-insofar-as-he-is-snubnosed. These aspects, however, can differ in their properties: Socrates-insofar-as-he-is-wise is not snubnosed, while Socrates-insofar-as-he-is-snubnosed is. According to Baxter, however, Leibniz's Law ranges only over individuals, so aspects do not provide a counterexample to it. Were it extended to include aspects, however, (LL) would be violated. But an adherent of (A1)–(A3) might be tempted to adopt Baxterian metaphysics to motivate a denial of (S1) and (S2), because if the constants denote aspects, the underlying SOA that makes them true would include something other than an individual and a property. This provides a seemingly compelling story for the opponent of (LL).

Such a move, however, faces an unpalatable dilemma in developing Baxter's view. This concerns whether aspects are object-like (in that they are denoted by terms of type e) or not. If they are, then the combination of two of Baxter's explicit background principles with a plausible constraint on the individuation of aspects leads to a contradiction. The background principles are the following (Baxter, 2018, pp. 910–11): (1) Baxter accepts that whenever x is F, there is an aspect $x_y[Fy]$ (read 'x insofar as it is F'); (2) Baxter formulates a notion of 'aspect identity', $\approx$, defined as Leibniz-equivalence. Baxter also appears to assume the following constraint on the individuation of aspects: (3) if F is not coextensive with G, then $x_y[Fy] \not\approx x_y[Gy]$; indeed, it is difficult to see how aspects could be fine-grained enough to do the needed work unless something akin to this principle held.

These three claims generate a contradiction when aspects are treated as the referents of type-e terms, given standard second-order comprehension (and the existence of more than two objects).[21] If aspects are not object-like, however, then there is no natural place for them in the type hierarchy; furthermore, it is unclear how they could be used to evade the S-argument. (The latter issue is not a problem for Baxter

21. This is simplest to demonstrate model-theoretically by a cardinality argument in a standard model: given κ objects, there are $2^{\kappa-1}\kappa$ aspects if κ is finite and 2^{κ} aspects if κ is infinite. But the same result can be shown in the object language, using only second-order comprehension on the basis of the three claims; the proof is parallel to the proof of the inconsistency of Frege's Basic Law V.

himself, of course, since he is not engaged in the project of justifying (A1)–(A3) or a similar (LL) violation.)

These considerations show that all ways of rejecting (S1) (No Corresponding Property, Missing Components, and Extra Components) that retain fundamental individuals are unsuccessful. Moreover, we have earlier argued that rejecting (S3), as we suppose CGL and BR will do, appears unmotivated. The S-argument thus suggests that we should adopt a more radical approach which would deny that individuals are required to underlie true atomic sentences. In the next subsection, we will give an independent, non-metaphysical argument that no account on which there is a finest-grained equivalence relation can plausibly deny that identity is this relation. These two results will serve as constraints in motivating the ontology of stuff we will propose in §4.

3.2 *The Finest-Grained Equivalence Relation?*

Our non-metaphysical argument against classical opacity focuses on the problems that arise with Submaximal Fineness, the assumption that there exists a finest-grained equivalence relation but it does not play the identity role. CGL accept the counterintuitive consequences of Submaximal Fineness but argue that the force of these is diminished when one appreciates the motivating intuitions of opacity (Caie et al., 2020, pp. 546–48). We think, however, that the problems with the position are far more pervasive than they acknowledge and provide compelling additional reasons to reject classical opacicism.

First, the Fregean link between quantification and identity will be broken. One consequence of this is to cleave an assignment-based count of the number of objects from an identity-based count within the language: if we use $\equiv$ and quantifiers to count the objects, there will be strictly fewer than if we count by assignments, since assignments track Leibniz-equivalence. This seems unappealing since it suggests that the (LL) violations are only created by a limitation in the object language.

Second, as CGL note, given Submaximal Fineness, certain extremely appealing inferences will fail to be valid: we shall no longer be able to infer from

(16) Something is F and G

and

(17) Something is F and not G

to

(18) There are at least two things that are F.

They argue, however, that the problem is not as bad as it seems once we appreciate the motivating force of the counterexamples to (LL) (Caie et al., 2020, p. 547).

Breaking the quantification–identity link, however, not only means that certain inferences that were previously assumed to be valid will come out invalid; in addition, connectives which have standardly been taken to be logical constants, on the ordinary criterion of logicality, will come out as nonlogical.

The problem arises already with identity itself. The standard account of logicality understands the logicality of a term through *permutation invariance* (Tarski, 1986; Sher, 1991; McGee, 1996). A permutation on a set D is a bijection $\pi : D \to D$. If D is the domain of individuals in a model of type theory, a permutation π on D induces, at each higher type, a permutation on entities of that type. An operator meets the permutation invariance criterion of logicality if, for every model M and permutation π on the domain of M, the permutation of appropriate type induced by π leaves the semantic value of the operator unchanged.[22] In particular, an operator $\odot$ of type $(e, e \to t)$ (the type of the identity connective) is logical just in case, for every model M and permutation π on M, if $\langle x, y\rangle$ satisfies $\ulcorner \mathbf{t}_1 \odot \mathbf{t}_2 \urcorner$, then so will $\langle \pi(x), \pi(y)\rangle$.

22. McGee (1996, pp. 576–78) gives persuasive reasons for adopting a more stringent criterion of bijection invariance, but for our purposes nothing is lost by choosing the simpler notion of permutation invariance.

As standardly understood, in type $(e, e \to t)$, only identity and nonidentity (and the total and empty relations) validate this criterion. But for CGL, neither WA-identity nor WA-nonidentity meets the permutation invariance criterion of logicality: let π be a permutation on the universe that maps Sand to the Eiffel Tower, the Eiffel Tower to Sand, Elena Kagan to Dupin, Dupin to Elena Kagan, and leaves everything else unchanged. The pair $\langle \text{Sand}, \text{Dupin} \rangle$ satisfies $\ulcorner \mathbf{t}_1 \equiv \mathbf{t}_2 \urcorner$, but $\langle \pi(\text{Sand}), \pi(\text{Dupin}) \rangle$ (i.e., $\langle$the Eiffel Tower, Kagan$\rangle$) does not. (In contrast, Leibniz equivalence remains permutation-invariant.)

This is highly revisionary: it seems that first-order logic with identity is the paradigm case of a genuine *logic*, and identity is a logical constant if anything is. Moreover, the divergence differs in kind from the usual proposals of non-classical logicians. Ordinarily, the non-classical logician differs from the classicist about the *behaviour* of the standard logical constants, not about whether they are logical in the first place. The proponent of strong Kleene logic, for instance, will disagree with the classicist about the existence of tautologies; she can nonetheless use a generalised version of the permutation-invariance criterion to verify that the classical connectives still count as logical on her models. In contrast, the sort of deviant account of identity that CGL endorse involves a more fundamental disagreement.

Similarly, the cardinality quantifiers $\exists_{\geq \kappa}$ will fail to be permutation-invariant (assuming that they track identity in the obvious way). But it seems highly counterintuitive that, although 'there exists a ...' is logical, 'there exist at least two ...' is not.

CGL could try to avoid this result by reformulating the invariance criterion using a narrowed class of acceptable permutations that would exclude, for example, the permutation π in the example above: a constant would then count as logical just in case its semantic value is unchanged under *permissible* permutations—permutations that preserve identity.

This would be wholly circular as a method of defending the logicality of identity, for it amounts to saying that identity is logical because it is one of the things preserved in all permutations that preserve identity. In any case, the permutations that would have to be rejected include ones built up trivially from functions that CGL will need for other purposes, if they view their system as providing a practical framework for reasoning about identity. If Gomez and Strachan both believe Hesperus to be distinct from Phosphorus, and Gomez's favourite luminous body is the Morning Star but Strachan's is the Evening Star, then the map from luminous objects to their lovers will not preserve identity. But there is a tension between using this function in talking about attitudes and denying that a permutation constructed from it counts as legitimate for purposes of determining logicality: if it is legitimate in the one case, why not the other?

Finally, accepting the divergence would not accord with CGL's aim to accommodate (LL) violations within a classical setting. The Fregean-inspired problems are not specific to CGL's framework but attend any system that accepts Submaximal Fineness: it will always be the case that assignment counts differ from identity counts, that common numerical inferences fail, and that identity and the numerical quantifiers cease to be logical. It is our view that these costs are sufficiently heavy that they should be avoided if at all possible.

These results reinforce the Williamsonian point that being the finest-grained equivalence relation (if such a thing exists) characterises the role of identity. Similarly, it is a maxim that identity is the relation that holds only between something and itself (if the world is a world of things). If there is a finer-grained equivalence relation than identity, however, then identity holds between an individual and something that can be distinguished from it by the standard resources of semantic theory such as assignments. If only individual objects are related by the identity relation, then there is nothing for the finer-grained equivalence relation to separate: the only way for the two relations to diverge is if what is called 'identity' within the system in fact holds among multiple individuals.

These departures from orthodoxy are not decisive: indeed, the system we will suggest in the next section has some of these consequences. The distinctive drawback of CGL's approach is that it incurs these costs without any motivating metaphysical story. As the S-argument shows, their picture is metaphysically mysterious. In contrast, we think that the revisionary proposal we shall canvass corresponds to one natural picture of the world. If this is correct, then the divergences from orthodoxy should be seen as surprising consequences of an underlying ontological perspective.

4. Stuff Ontology

In order to avoid these two objections, the opponent of (LL) must reject both individualism and Submaximal Fineness. Is there any metaphysical picture on which these commitments make sense? We argue that a *stuff ontology* provides an attractive approach for the opponent of (LL), because the principles governing the division of stuff motivate both rejecting (S1) and denying the existence of a finest-grained equivalence relation. Our purpose is not to defend such a stuff ontology: we merely wish to elucidate it and show that it motivates a coherent anti-(LL) posture.

There have been a number of recent discussions of stuff ontologies (Zimmerman, 1997; Markosian, 2004; Kleinschmidt, 2007; McKay, 2015), which focus primarily on applying the notion of stuffs to solve coincidence problems.[23] Here, in contrast, we try to motivate a particular large-scale metaphysical picture, concentrating on the hierarchical property structure of the stuff-universe; it is this, on our view, that explains why stuffs can help the foe of (LL).

The thing/stuff distinction is motivated by the observation that natural languages contain not only count nouns, such as 'table', 'cat', and 'electron', but mass nouns, such as 'water', 'wine', and 'cheese'. Count nouns admit indefinite and numerical determiners ('a'/'an', 'one', 'six', 'many'); mass nouns admit only mass determiners ('some', 'much'). Clearly, in not every case does the count/mass distinction track something of metaphysical significance, and it is open to the orthodox individualist to provide paraphrases in a counting idiom at the level of logical form. For instance, she could explain reference to 'some cheese' in terms of a set of cheesy things, or a mereological sum of cheesy things, or a maximal portion of cheese, or some such account.[24] The stuff ontologist, however, thinks that these strategies do not succeed in all cases; at least sometimes, she maintains, that which mass nouns pick out has unique structural properties that no object-based approach can capture. It is not our purpose here to convince the dedicated defender of orthodoxy, but we think there is value in exploring the metaphysics that results from taking our pervasive non-individualistic natural language commitments as a guide to underlying ontological structure.

The key distinctive property of stuff is that it cannot be counted. We can say that there are six chairs in the room, but we cannot literally say that there are four waters or five woods. At most, this is elliptical for 'four units of water' or 'five kinds of wood': in such cases, we count not stuff itself but portions or kinds of stuff.

We can leave open whether, in these ordinary examples, a reduction to an enumerable object (in the sense that the stuff is nothing over and above that object) is possible. With some mass nouns it clearly is: although we cannot ask how many 'furnitures' there are in the room, 'furniture' seems to refer to a derivative entity that can be exhaustively ontologically accounted for through such things as chairs and tables. With water and wine the case is less clear: perhaps the water is nothing over and above the water molecules, or perhaps it is composed of them without being ontologically reducible to them.

But, crucially, the serious stuff ontologist maintains that at least for some stuffs—*fundamental stuffs*—there is no reduction to things. And the very serious stuff ontologist maintains that these fundamental stuffs are what ultimately make up the world.

23. A notable exception is provided by Laycock (2006), who offers an in-depth discussion of stuff that represents an important precursor to the present account.

24. See, e.g., Burge (1972), Moravcsik (1973), Montague (1973), Link (1983), Bunt (1985), and Moltmann (1997) for classic examples of these and similar strategies. For a recent survey, see the essays in Kiss et al. (2021).

It is this sort of very serious stuff ontologist who we think is in a position to reject (LL). She starts from the observation that fundamental stuff requires a new quantifier. The existential quantifier, $\exists$, is informally taken to stand for 'there is a ...', which results in a conflict with the grammar of mass nouns: while we can say 'there is a table' and symbolise this as $\exists x\ Tx$, we cannot symbolise 'there is cheese' as $\exists x\ Cx$, for this would be to say that there is a discrete individual x that is a certain way, but 'there is cheese' is different from 'there is a portion of cheese' or 'there is a kind of cheese', or, indeed, any formulation that would be committed to *one* thing that is cheese.

This argument has not been unopposed: Ned Markosian 2015, p. 685 has claimed that the grammatical feature is superficial: we can use the existential quantifier indifferently for both stuff and things, allowing the natural language paraphrase to vary. On our view, however, Markosian fails to take into account the link between quantification and counting. If we accept that 'there is both hard cheese and non-hard cheese' can be perspicuously regimented as $\exists x \exists y (Cx \wedge Cy \wedge Hx \wedge \neg Hy)$, then, on the standard account of quantifiers and counting, 'how many cheeses?' questions will clearly be coherent, and it is precisely this which the stuff ontologist must reject.

We can think of the serious stuff ontologist as speaking a language that includes, in lieu of ordinary existential and universal quantifiers, primitive stuff quantifiers, which form sentences from terms for stuff-kinds without using bindable variables: read $\sum W$ as 'there is some wine' and $\prod C$ as 'everything is wine'. Optional location parameters can be added to express claims like 'there is some wine at Karla's' ($\sum^{\alpha} W$); we can also introduce *measure* quantifiers, which correspond to numerical quantifiers for objects. Thus, the existence of at least two litres of wine at Karla's entails the existence of at least one litre of wine at Karla's: $\sum^{\alpha}_{\geq 2\mathrm{L}} W \vdash \sum^{\alpha}_{\geq 1\mathrm{L}} W$. (The fact that stuff ontology allows for a natural treatment of these inferences is one of its theoretical advantages and provides a motivation for distinguishing between portions of stuff and genuine objects.)

We have used mundane examples — wine, cheese, and the like — to explain the framework, but the very serious stuff ontologist's fundamental stuffs may look quite different: they may be the fields postulated by physics or perhaps a neutral monist stuff with both a physical and mental aspect or a stuff of some other sort. (A positive defence of the view would have to say something about this question; for our purposes, however, we can remain neutral.)

In principle, a variety of views about kinds of stuff and their interdependence are compatible with very serious stuff ontology. But there is a general theoretical impetus, on grounds of elegance and simplicity, towards reducing fundamental ontological commitments and hence towards priority monism (Schaffer, 2010) that applies particularly strongly to stuff ontology. The thing ontologist faces countervailing pressure from the existence of clear boundaries between discrete objects; it is characteristic of fundamental stuff that it has no such boundaries. Once we have moved from thinking of discrete things to some stuff, it is a short step to the view that all unbounded stuff is derived from a single, variegated basic world-stuff.

If the basic world-stuff is variegated, then there will be subkinds of stuff, corresponding to each of the ways in which it varies. (As we use the term here, a 'subkind' of stuff is all the stuff of that particular type — some but not all of the variegated stuff — not a universal or other abstract object used to categorise the stuff.[25])

The subkinds might be related to the world-stuff in a number of ways. There could be two levels: a fundamental world-stuff and then a plethora of subkinds, all on a par. But this would make the world-stuff very unlike our familiar stuffs, which come in hierarchies of determination. For instance, wine is a subkind of liquid, claret a subkind of wine, Graves a subkind of claret, and so on; important structure would be lost if all of these divisions were posited at once. It thus seems better motivated to assume that at the first stage the world-stuff would

25. Kinds as universals, in addition to the kinds we posit here, may be required for other theoretical purposes — for instance, to express the claim that there could have been more world-stuff than there actually is. We leave this question open.

be partitioned into a few, relatively fundamental kinds; these kinds would then be further divided in turn, and so on. The process might culminate in atomic kinds — kinds that admit no subkinds — or it might go on forever; the serious stuff ontologist endorses the latter view.

The serious stuff ontologist thus claims that stuff is never perfectly homogeneous; however finely we separate out kinds of stuff, what remains is still variegated. In effect, she adopts Leibniz's picture of an infinitely divisible universe of limitless qualitative diversity (though without the underlying atomic realm of monads he takes as its metaphysical basis). This picture has two consequences: first, stuffs are infinitely numerous; second, the relations of stuff to substuff always go infinitely deep: there is no final level of substuff. Were there a terminal subkind, it would parcel out homogeneous stuff. Because she views stuff as ineradicably mixed, the serious stuff ontologist has a natural affinity for a conception of the hierarchy as 'indefinitely extensible' (Dummett, 1963; 1991) rather than given all at once. The serious stuff ontologist views the stuffs in the hierarchy as very different from individual substances as conceived on the pegboard model of ontological structure: whereas substances would play the role of pegs, stuffs are constituted by their qualitative properties.[26]

In addition to the primary hierarchy of kinds, stuff is secondarily divided into portions as well: for instance, stuffs derivatively form portions such as a glass of water or a decanter of wine.[27] For the serious stuff ontologist, mere portions of stuff are not metaphysically privileged in the way that kinds of stuff are: kinds are explanatorily fundamental and account for the constitutive structure of the world, but portions are merely accidental divisions.

The kind hierarchy and the portion hierarchy cut across one another: a portion of stuff may contain many kinds of stuff, and a kind of stuff may comprise numerous scattered portions. For any carving of the world-stuff into kinds, each kind can be portioned out, and, correspondingly, for each portion of world-stuff, we can separate out different kinds within it.

These portions have the status of quasi-things: we can refer to them using singular terms, but unlike the substance ontologist's notion of a genuine thing, which is always no more and no less than one thing, portions come in amounts. For the very serious stuff ontologist, all singular reference must ultimately be explained through these portions: she does not eschew the ordinary $\exists$ and $\forall$, but she views facts expressible using them as derivative on those facts we capture with $\sum$ and $\prod$.

The infinitely descending hierarchy of qualitative properties is the *only* fundamental posit for the very serious stuff ontologist. As a consequence, all relations must be constructed out of monadic, qualitative properties. In particular, every equivalence relation is exactly as fine as the available configurations of monadic properties on some level. So there can be no finest-grained equivalence relation, since there is no ultimate level in the hierarchy of stuffs and substuffs. But for the very serious stuff ontologist, who admits no individuals prior in nature to the stuffs, identity just is the equivalence relation on portions induced by a partitioning of stuff at a certain level.[28] So, for every relation that can serve as a candidate for identity, there will be a stronger one, corresponding to a level further down in the hierarchy.

This picture thus meets the two desiderata set out at the beginning: the very serious stuff ontologist can reject (S1), because she is not committed to individuals figuring in states of affairs. Nevertheless, she is free to use singular terms to designate quasi-individuals, portions of stuff. Using the qualitative divisions available at a level to form equivalence classes, she can then introduce an identity symbol to express that portion *a* is within the same level-relative equivalence class as portion *b*. Thus, Leibniz's Law is not a mere vacuity. These identity statements,

26. See Turner (2011) and our discussion of this model in §3.

27. We need not decide whether any collection of stuff, no matter how scattered, makes up a portion: our general metaphysics is compatible with a variety of positions about conditions on portions.

28. One might try to use the regions occupied by portions of stuff in order to define identity for stuffs. The very serious stuff ontologist, however, will reject the substantivalist presupposition that regions are ontologically distinct from the portions that occupy them.

however, will always be evaluated relative to some bounded depth in the hierarchy. For this reason, violations of (LL) may arise: at any level ℓ, a portion a of stuff may be ℓ-identical to a portion b of stuff (i.e., they stand in the finest-grained equivalence relation existing at ℓ), even though a and b can be differentiated by properties that arise only at a level beyond ℓ.

The very serious stuff ontologist will give something like the following story to explain how singular reference comes about: I ostend a particular portion of stuff, and pick it out by one of its kinds, for instance as 'this watery stuff.' I then baptise the portion so picked out by saying: 'This watery stuff is named "Lake Ontario."' The content of subsequent uses of 'Lake Ontario' will then be *dthat (this watery stuff)*. In something like this manner, singular reference can come about through apportionments of stuff.

This account of singular reference explains how the singular terms a and b come to denote portions of stuff that are level ℓ-identical but perhaps distinguishable at a level beyond ℓ. For instance, consider a particular portion of liquid that is both winey and watery. (We use these in our toy example as placeholders for ultimate stuffs; real wine and water, of course, are better thought of as a mixture.) Applying the account above, we can imagine that we have come to name this portion of liquid in two ways: first, we might use the definite description 'this watery stuff' and form a singular term *dthat (this watery stuff)*; then, we might use 'this winey stuff' and form the singular term *dthat (this winey stuff)*. We now have two singular terms referring to this portion of liquid, derived from the conceptual distinction between watery and winey stuff. If, however, this conceptual distinction does not correlate to a distinction among qualitative properties present at level ℓ, then *dthat (this watery stuff)* will be ℓ-identical to *dthat (this winey stuff)*, but they will be distinct at the first level at which watery stuff is distinguished from winey stuff.

Although it is not our goal to give a positive defence of stuff ontology, it is worth discussing a key potential objection. It might be claimed that successful physical theories, such as the standard model of particle physics, involve *prima facie* commitment to individual particles and hence an individualistic ontology; thus, the very serious stuff ontologist runs afoul of the naturalistic constraint that metaphysics should reflect the deliverances of the natural sciences.[29]

We first note a methodological problem with the objection: the quantum field theory in which the standard model is couched is an *effective* field theory and does not purport to track fundamental structure; in fact, since a high-energy cutoff is built into QFT, it would violate the theory's own presumptions to try to apply it at sub-Planckian scales. And, as David Wallace (2011, p. 211) has noted:

> Whatever our sub-Planckian physics looks like (string theory? twistor theory? loop quantum gravity? non-commutative geometry? causal set theory? something as-yet-undreamed-of?) there are pretty powerful reasons *not* to expect it to look like quantum field theory on a classical background spacetime. As such, what QFT (of any variety) says about the nature of the world on length scales below $\sim 10^{-43}$ m [...] doesn't actually tell us anything about reality.

It should be emphasised that the problem is not merely that the standard model is not final physics. Rather, the standard model itself says that we cannot coherently extend it to arbitrarily small scales: unlike the theory of Newtonian gravity (for example), it is not even a complete picture of a world with different laws from our own.

Nonetheless, we do not find the objection compelling even as applied to the standard model. Following David Baker (2016), we can distinguish two broad families of interpretations of QFT — 'particle-theoretic' and 'field-theoretic'. Particle-theoretic interpretations vary in the extent to which the core commitments of the classical notion of 'particle' remain intact.

Even the most stringently particle-theoretic interpretations, however, preserve little of the intuitive conception of particles as localised indi-

29. We thank an anonymous reviewer for raising this point.

viduals that can be counted. Localisation is ruled out by no-go theorems (Malament, 1996; Halvorson and Clifton, 2002). As for counting, given a specific Fock space representation of a system in QFT — the most hospitable framework for particle interpretations — there will, in general, be a well-defined 'particle number observable', but the system may be in a superposition of this observable. This already rules out any direct equation between the particle 'number' operator and a literal count of individuals as specified using existential quantification.

Moreover, this hospitable framework is often unavailable. Even in the toy case of a free (noninteracting) boson field, Minkowski and Rindler representations — corresponding to inertial and accelerated observers, respectively — have different particle number observables, and each representation interprets the other's observable as yielding an expectation value distinct from its own (Clifton and Halvorson, 2001; Ruetsche, 2011, pp. 191–219). As Baker (2016, p. 7) remarks,

> Since there is nothing physically privileged about either observer's definition, it must be that the number of particles is a perspective-dependent fact. If particles belonged to QFT's fundamental ontology, the number of particles would not be dependent on one's perspective — here we assume, plausibly, that the number of fundamental entities in the universe is an objective fact. So if this argument succeeds, it rules out fundamental particle interpretations.

Still worse, it is not always the case that there are Fock space representations at all: Doreen Fraser (2008) has shown that none exist for a large class of interacting theories, a result which, as Baker (2016, pp. 8–9) notes, undermines even very weak interpretations that attempt to find a place for particles in non-fundamental, much less fundamental, ontology.

In contrast, field-theoretic interpretations abandon any commitment to particles, even in a minimal sense, and instead take the physical realm to consist in properties of spacetime points or regions (Baker, 2016, p. 9). There are several proposals to implement this idea: the simplest ('wavefunctional') field-theoretic interpretations treat the quantum field as (roughly) composed of configurations of classical fields, represented in a complex-valued Hilbert space; as such, their ontological commitments are not qualitatively different from those of classical field theories. While standardly formulated using spacetime points, there are alternative, empirically equivalent formulations of classical field theory using gunky spacetime (Arntzenius, 2003, 2008). Purged of commitment to spacetime points, classical field theories would provide a hospitable setting for very serious stuff ontology. Baker (2009) does argue that wavefunctional interpretations confront a difficulty parallel to that faced by particle interpretations, but his preferred alternative approach, which uses algebras of operators, remains fundamentally field-theoretic (in that it assigns properties to spacetime regions) and thus does not present any new challenges to very serious stuff ontology.

Of course, this is not a demonstration that stuff ontology is compatible with physics; such a result must await the development of a theory integrating gravity and quantum field theory. But at present we have no reason to suspect that our basic scientific commitments rule out a stuff ontology.[30]

This concludes our sketch of a principled metaphysical stance on which violations of (LL) would be countenanced without falling afoul of the arguments developed in §3. At this point, however, it is worth asking whether a consistent theory can be developed that would enable us to reason about stuffs without reintroducing a finest-grained equivalence relation. In the next section, we show that this is possible using a toy theory: we describe a class of hierarchical models and provide a semantics, complete with level-relative notions of evaluation, for a second-order language with identity over those models. These models are not fully faithful to the very serious stuff ontologist's metaphysics, for they are (as is standard) constructed in set theory, and thus built out of sets of individuals. But they nonetheless share the stuff ontologist's

30. We thank Tushar Menon for exceptionally helpful discussions of the relationship between physics and stuff ontology.

decisive commitment: there is no way to single out individuals from within the object language. The toy semantics thus provides a mathematically tractable way to demonstrate the internal coherence of the stuff ontologist's account.

5. Model Theory for Stuff

The central idea behind our toy model for stuffs is to employ infinite sets to simulate stuff. Nothing in the very serious stuff ontologist's fundamental ontology corresponds to the elements of this set: they are merely an artefact of the model theory, but they give us a structure that captures the infinite division of stuff.

For convenience, we build the model over the natural numbers, but any infinite collection could be taken as a first-order domain.[31] The fundamental requirements are that the second-order domain (standing in for the hierarchy of kinds) contain not *every* subset of the first-order domain, but only carefully chosen sets: every set in the second-order domain is infinite (or empty), sets subdivide as we progress down the levels, and any two distinct elements are ultimately separated on some level. In this way, we obtain a progressive winnowing of sets corresponding to the infinite divisibility of stuff: at a given level, only those sets are available which correspond to the kinds produced at that level.

We start from the set of all natural numbers, standing in for the world-stuff, and we partition at each level by separating out the even-numbered from the odd-numbered elements (under the obvious ordering). The zeroth level contains just $\{0, 1, 2, 3, \ldots\}$; the first, $\{1, 3, 5, 7, \ldots\}$ and $\{0, 2, 4, 6, \ldots\}$. The second level contains four sets, $\{1, 5, 9, 13, \ldots\}$, $\{3, 7, 11, 15, \ldots\}$, $\{0, 4, 8, 12, \ldots\}$, and $\{2, 6, 10, 14, \ldots\}$. The third contains eight infinite sets, the fourth sixteen, and so on *ad infinitum*.

In order to turn the hierarchy of levels into a domain containing the values of monadic predicate expressions (modelling kinds of stuff, whether these are cheese, wood, and wine, and so on, or something more esoteric), we take the union of all the levels and close under complements and finite unions and intersections. We do *not* close under countable unions and intersections, for this would lead to singletons: we do not wish the model to be able to discriminate out any nonempty sets with finitely many members. (The very serious stuff ontologist admits no divisions that cannot be further divided.)

Thus we have a first-order domain, D_0, which is just ω, and a second-order singulary domain, D_1^1, formed of finite unions and intersections and complements of the subsets of ω constructed at various levels by the splitting procedure. We now need relational domains D_1^n for $n > 1$; crucially, D_1^2 contains all the two-place relations we accept, and thus all the equivalence relations: these provide the only acceptable candidates for the identity relation. Here the idea is that relations should cut no finer than properties: everything in D_1^n can be constructed from a collection of sets in D_1^1, one for each argument place, and no two individuals which the sets do not distinguish can be distinguished by the relation. Crucially, each equivalence relation will be constructed from a finite number of sets in D_1^1. As a result, there will always be a level ℓ such that the relation cuts no finer than the partition on the natural numbers induced at that level. Since we do not close D_1^1 under infinite intersections, there is no way to construct a 'level-ω' equivalence relation that would capture all of the divisions: we are limited to an infinite chain of equivalence relations, each of finite level.

The collection of domains forms a fixed *frame*. We use the frame to provide a semantics for a second-order language $\mathscr{L}$: $\mathscr{L}$ has first- and second-order variables, possibly constants, the usual connectives and quantifiers, and equality. A model for $\mathscr{L}$ is constructed from the frame by adding an interpretation: the interpretation assigns to each first-order constant an element of D_0 and to each second-order constant an element of the appropriate D_1^n. Similarly, a variable assignment on

31. For instance, in illustrating how a stuff ontologist could make sense of a field theory on gunky space, we might produce a heuristic model whose first-order domain is an open interval I of real numbers and whose second-order domains contain, as we move down the levels, collections of finer and finer subintervals of I. Here some care is required in choosing an appropriate linear ordering for the levels, but the basic principle stays the same.

the model maps first-order variables to elements of D_0 and second-order variables to elements of the corresponding D_1^n. The values of second-order variables on an assignment may occur at any level in the hierarchy.

We interpret formulas in L relative to two level parameters: a formula holds or does not hold in a model, on an assignment, *at an equality-level* ℓ_1 and *at a quantifier-level* ℓ_2. Specifically, the formula $\ulcorner \mathbf{t_1} = \mathbf{t_2} \urcorner$ holds on a given assignment at equality-level ℓ_1 just in case the values of $\mathbf{t}_1$ and $\mathbf{t}_2$ on that assignment are not separated by the finest-grained equivalence relation constructible at level ℓ_1. Similarly, $\ulcorner \forall \mathbf{X}\ \phi \urcorner$, with $\mathbf{X}$ a second-order variable, holds at quantifier-level ℓ_2 just in case ϕ holds on every relevant assignment mapping $\mathbf{X}$ to a value located at or above level ℓ_2. (First-order quantifiers are treated standardly. In contrast to bound second-order variables, there are no level restrictions on *free* second-order variables.)

Evaluating identity and second-order quantified formulas relative to a level captures the serious stuff ontologist's commitment to the indefinite extensibility of the stuff hierarchy. For her, it makes no sense to evaluate a formula 'from the outside', independently of levels, for the hierarchy is not given all at once: it develops as the variegated stuff is separated out into kinds and subkinds.

The two level parameters give us a variety of well-behaved notions of truth-in-a-model and validity. (And logical consequence, although for simplicity we consider only the single-formula case.) The notions collapse for identity-free first-order sentences, but the differences are crucial for sentences with identity or second-order variables. We introduce six such notions in the appendix; here we explain the four most important (holding *simpliciter*, holding on pairs, holding on pairs pointwise in the limit, and holding on pairs uniformly in the limit).

A formula *holds simpliciter* in a model if it holds on all assignments at every equality-level and every quantifier-level; it is valid *simpliciter* if it holds *simpliciter* on all models.

This is the most demanding notion, and Quantified Leibniz's Law,

$$\mathbf{t_1} = \mathbf{t_2} \to \forall X\ (X\mathbf{t_1} \leftrightarrow X\mathbf{t_2}), \tag{QLL}$$

fails to be valid *simpliciter*. Consider a model and assignment on which $\mathbf{t_1}$ denotes 4, $\mathbf{t_2}$ denotes 6, and the formula is evaluated at quantifier-level 2 and equality-level 1. The finest-grained equivalence relation extant at level 1 treats $\{0, 2, 4, 6, \ldots\}$ as a partition cell; but 4 and 6 are separated by the set $\{0, 4, 8, \ldots\}$, which is available at level 2 to be the value of the quantified variable $\mathbf{X}$.

Since (QLL) is not valid *simpliciter*, the schematic version of Leibniz's Law,

$$\mathbf{t_1} = \mathbf{t_2} \to (\phi(\mathbf{t_1}) \leftrightarrow \phi(\mathbf{t_2})), \tag{LL}$$

is also not valid *simpliciter*.

In addition, second-order existential generalisation (and its dual, universal instantiation),

$$\Phi(\mathbf{T}) \to \exists X\ \Phi(X) \tag{EG2}$$

also fails: consider a model and assignment where $\mathbf{T}$ has as its semantic value a set at level 10 of D_1^1, but the formula is evaluated at a quantifier-level of 5, and no set at level 5 or above of D_1^1 satisfies the open formula Φ.

Another natural, but less demanding, semantic notion is *holding on pairs*: a formula holds on pairs in a model if it holds on every assignment in every circumstance where the equality-level is the same as the quantifier-level; a formula is valid on pairs if it holds on pairs in every model.

Since the identity-level was irrelevant to the counterexample, (EG2) still fails to be valid on pairs. (LL) is also not valid on pairs: if we continue to let $\mathbf{t_1}$ denote 4 and $\mathbf{t_2}$ denote 6, and we evaluate at equality-depth and quantifier-depth 1, the instance $\mathbf{t_1} = \mathbf{t_2} \to (X\mathbf{t_1} \leftrightarrow X\mathbf{t_2})$ will fail when the free variable X is assigned a set below level 1 that splits 4 and 6.

But (QLL) is valid on pairs. A proof is given in the appendix, but the reasoning is clear: when the quantifier and equality levels move in tandem, no set that is available to be the value of a bound second-order variable can separate a pair unless an equivalence relation available at that level already separates them.

A formula holds on pairs uniformly in the limit in a model if there exists some ℓ such that for every $\ell' > \ell$, for every assignment, the formula holds in the model on that assignment at identity- and quantifier-level ℓ'. A formula holds pointwise on pairs in the limit in a model if, for every assignment α, there exists some ℓ such that for every $\ell' > \ell$, the formula holds on α in the model at identity- and quantifier-level ℓ'. A formula is pointwise (resp. uniformly) valid on pairs in the limit if it holds pointwise (resp. uniformly) on pairs in the limit in all models.

Since validity on pairs requires that a formula holds on all pairs, whereas both limit notions require only that the formula holds on certain pairs, any formula that is valid on pairs will be both pointwise and uniformly valid on pairs. Thus (QLL) is both pointwise and uniformly valid on pairs in the limit. Neither (LL) nor (EG2) is uniformly valid on pairs in the limit, since for any candidate ℓ there will be some assignment α falsifying the conditional using a value below ℓ.

But (LL) and (EG2) do hold pointwise on pairs in the limit: for every assignment α we can find an equivalence relation that distinguishes everything that can be distinguished by finitely many values of variables on that assignment: when this equivalence relation is taken to be identity, (LL) holds. Similarly, when the quantifier is evaluated at this level, (EG2) holds.

It is important to be clear about the limitations of these notions: they all represent 'external' perspectives on the totality of levels, and thus — despite their heuristic utility — none of them represents what an inhabitant of the serious stuff ontologist's world can directly access. Nonetheless, even as representational aids, there is a sense in which the notions of holding simpliciter and holding on pairs are more faithful to the programme of stuff ontology than the limit notions. All notions of truth-in-a-model require quantification over all levels (or pairs of levels), but holding simpliciter and holding on pairs permit truth-at-a-level (or pair-of-levels) to be evaluated purely locally at the first step. In contrast, the limit notions require a potentially unbounded search through finer and finer levels before even this step can be completed.

Hence, (LL) (and, on one variant, (QLL)) fail on our toy models when we use the semantic notions that make the most sense for the system. Unlike CGL's framework, however, our models do not deliver the existence of a finest-grained equivalence relation. These models demonstrate the consistency of an approach that denies (LL), while endorsing the conditional claim that if there is a finest-grained equivalence relation, it must be identity.

The model provides a guide to the proper interpretation of identity claims about portions of stuff from the perspective of the very serious stuff ontologist. In general, from the claim that this portion of stuff is identical to that portion of stuff, it will not follow that if this portion of stuff is some way, then that portion of stuff is that way as well. Consider, for instance, the case of a portion of liquid, composed of a mix of wine and water, which we pick out twice by direct reference (through rigidified descriptions, the first based on its winey nature and the second on its watery nature): 'this stuff is that stuff' is true at an identity-level above that on which water and wine are separated and false below that level. Even at a level on which the identity claim does hold, this stuff will still be distinguished from that stuff by a property: whether that property is available to be quantified over will depend on the quantifier-level.

We do not claim that the model theory provides a royal road to understanding stuff ontology: there will always be a gap between what we can represent using set-theoretical constructions and what the very serious stuff ontologist is committed to. But we think that the presentation given here fleshes out some of the formal behaviour that the very serious stuff ontologist should expect of her identity connective.

6. Conclusion

Our purpose in this paper has been twofold: we have offered both a critique of current work that endorses genuine Leibniz's Law violations and a positive suggestion for a metaphysical picture on which such violations make sense.

The work of BR and CGL represents the best-developed recent attempt to allow for Leibniz's Law violations in a systematic way. In particular, CGL's classical opacicism promises an appealing combination of orthodox logic and nontrivial cases of opacity. But, as we argued in §3, classical opacicism and its variants are besieged by problems: they lack a reasonable response to the metaphysical challenge of the S-argument and they leave their proponents in the untenable position of simultaneously accepting a finest-grained equivalence relation and denying that it is identity.

The proponent of the stuff ontology we sketch in §§4–5 faces none of these problems: she motivates Leibniz's Law violations by pointing to the infinite qualitative heterogeneity of stuff. On her ontological picture, the world is infinitely qualitatively heterogeneous. Our identity statements give rise to (LL) violations because our attributions of identity only hold relative to a differentiation into properties that is always subject to finer-grained division. Stuff ontology represents the best alternative of which we are aware to a fully classical, Leibniz's Law-endorsing treatment of identity.

Appendix: Stuff-Like Models

We use a language $\mathscr{L}$ containing denumerably many individual variables $\mathsf{x}_1, \mathsf{x}_2, \ldots$, forming the set Vars_0; for each $n \geq 1$, denumerably many n-ary predicate variables $\mathsf{X}^n_1, \mathsf{X}^n_2, \ldots$, forming the set Vars^n_1; the logical constants $\neg, \vee, \forall$, and $=$; parentheses and other grouping operations; at most denumerably many individual constants $\mathsf{c}_1, \mathsf{c}_2, \ldots$, forming the set Const_0; for each $n \geq 1$, at most denumerably many n-ary predicate constants $\mathsf{C}^n_1, \mathsf{C}^n_2, \ldots$, forming the set Const^n_1. The language $\mathscr{L}^+$ differs from $\mathscr{L}$ in including the variable-binding operator λ.

We use boldface for metasyntactic variables: substituends for $\mathbf{x}, \mathbf{X}, \mathbf{c}, \mathbf{C}, \mathbf{t}, \mathbf{T}, \boldsymbol{\phi}$ (and obvious variants, with or without adicity indications) are individual variables, predicate variables, individual predicate constants, individual terms, predicate terms, and sentences, respectively. Unless otherwise noted, i, j, k, m, n, and variants in the metalanguage range over natural numbers.

Formation rules for both $\mathscr{L}$ and $\mathscr{L}^+$ are standard: in $\mathscr{L}^+$, we allow λ to bind any finite sequence of object variables (but not predicate variables): thus $[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \boldsymbol{\phi}]$ is an n-ary predicate term.

We set $\mathsf{Vars}_1 = \bigcup_{n \geq 1} \mathsf{Vars}^n_1$; $\mathsf{Const}_1 = \bigcup_{n \geq 1} \mathsf{Const}^n_1$; $\mathsf{Vars} = \mathsf{Vars}_0 \cup \mathsf{Vars}_1$; $\mathsf{Const} = \mathsf{Const}_0 \cup \mathsf{Const}_1$. We use Term_0, Term^n_1, Fmla, and Sent for the sets of first-order terms, n-ary second-order terms, formulas, and sentences.

We first define some operations that will be useful in the following construction.

Definition 6.1. For infinite $A \subseteq \omega$, we define the *enumeration function* on A, $\mathsf{enum}_A : A \to \omega$, proceeding from the minimal element onward, as follows. We start with the minimal element: $v_0 = \min(A)$. We then set $v_{n+1} = \min(A \smallsetminus \{v_i : i \leq n\})$. We set $\mathsf{enum}_A(x) = n$ if and only if $v_n(A) = x$. We define $\mathsf{ev}(A) = \{x \in A : \mathsf{enum}_A(x) \text{ is even}\}$; $\mathsf{od}(A) = \{x \in A : \mathsf{enum}_A(x) \text{ is odd}\}$.

Definition 6.2. For a set $A \subseteq \mathscr{P}(B)$, we define the *finitary closure* of A (relative to B), $\mathsf{fcl}(A)$, as the smallest set with the following properties: if $x \in A$, then $x \in \mathsf{fcl}(A)$; if $x \in \mathsf{fcl}(A)$, then $B \smallsetminus x \in \mathsf{fcl}(A)$; if $x_1, \ldots, x_n \in \mathsf{fcl}(A)$, then $\bigcup_{1 \leq i \leq n} x_i \in \mathsf{fcl}(A)$ and $\bigcap_{1 \leq i \leq n} x_i \in \mathsf{fcl}(A)$.

Definition 6.3. For any $n \in \omega$ and sets $A_1, \ldots, A_n$, we define the *relational projection* of $A_1, \ldots, A_n$ as follows: $\mathsf{RP}(A_1, \ldots, A_n) = \{\langle x_1, \ldots, x_n \rangle : x_1 \in A_1 \wedge \cdots \wedge x_n \in A_n\}$.

We now proceed to the construction of a frame $\mathscr{F}$ for $\mathscr{L}$ as follows.

Definition 6.4. For the first-order domain of the frame, we use the natural numbers: $D_0 = \omega$.

For the singular second-order domain, we define a hierarchy of levels based on using ev and od to repeatedly partition ω.

Definition 6.5. $S_0 = \{\omega\}; S_{n+1} = \{\mathsf{ev}(x) : x \in S_n\} \cup \{\mathsf{od}(x) : x \in S_n\}$. We set $S_\omega = \bigcup_{n\in\omega} S_n$ and define $D_1^1 = \mathsf{fcl}(S_\omega)$. We define the *level-fragments* of D_1^1 as follows: $D_1^{1(\mathsf{level}\leq m)} = \mathsf{fcl}(\bigcup_{i\leq n} S_m)$.

Using relational projections, we define the polyadic second-order domains:

Definition 6.6. For $n > 1$, let $D_1^n = \mathsf{fcl}(\{\mathsf{RP}(A_1, \ldots, A_n) : A_1, \ldots, A_n \in D_1\})$. We define the level-fragments of the polyadic domains as follows: for $n > 1$, $D_1^{n(\mathsf{level}\leq m)} = \mathsf{fcl}(\{\mathsf{RP}(A_1, \ldots, A_n) : A_1, \ldots, A_n \in D_1^{1(\mathsf{level}\leq m)}\})$.

Finally, we define the entire frame and models over it:

Definition 6.7. Let $\mathscr{F} = \langle D_0, D_1^1, D_1^2, \ldots\rangle$. We set $D = D_0 \cup \bigcup_{i\in\omega} D_i^1$. An *interpretation* of $\mathscr{L}$ or $\mathscr{L}^+$ over $\mathscr{F}$ is a function $I : \mathsf{Const} \to D$ with the constraint that $I(\mathbf{c}) \in D_0$ and $I(\mathbf{C}^n) \in D_1^n$. A *model* for $\mathscr{L}$ or $\mathscr{L}^+$ is a pair $M = \langle \mathscr{F}, I\rangle$, where I is an interpretation of $\mathscr{L}$ or $\mathscr{L}^+$ over $\mathscr{F}$. We write $\mathfrak{Mod}$ for the class of models.

We introduce a special abbreviation for the finest-grained equivalence relation at a level:

Definition 6.8. Let $\mathfrak{E}^{(\mathsf{level}=m)} = \bigcup\{\langle x, y\rangle \in A \times A \ : A \in S_m\}$.

Clearly, $\mathfrak{E}^{(\mathsf{level}=m)} \in D_1^{2(\mathsf{level}\leq m)}$.

Variable assignments follow the standard definition. In particular, we allow the value of an n-ary second-order variable to occur anywhere in D_1^n: the values of variables on an assignment are not bounded at any specific level.

Definition 6.9. An assignment is a function α : $\mathsf{Vars} \to D$ with the constraint that $\alpha(\mathbf{x}) \in D_0$ and $\alpha(\mathbf{X}^n) \in D_1^n$. We write $\alpha \overset{\mathbf{x}}{\sim} \beta$ (resp. $\alpha \overset{\mathbf{X}}{\sim} \beta$) to indicate that α and β agree on all variables except $\mathbf{x}$ (resp. $\mathbf{X}$).

We write $\mathfrak{Assn}$ for the class of assignments.

We now define semantic notions for $\mathscr{L}$.

Definition 6.10. The semantic value of an expression ξ on model M, assignment α, identity level ℓ_1 and quantifier level ℓ_2, $[\![\xi]\!]^{\ell_1,\ell_2}_{M,\alpha}$, is defined recursively as follows. For formulas, here and throughout, $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 0$ just in case $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} \neq 1$.

- $[\![\mathbf{x}]\!]^{\ell_1,\ell_2}_{M,\alpha} = \alpha(\mathbf{x})$;
- $[\![\mathbf{X}^n]\!]^{\ell_1,\ell_2}_{M,\alpha} = \alpha(\mathbf{X}^n)$;
- $[\![\mathbf{c}]\!]^{\ell_1,\ell_2}_{M,\alpha} = I(\mathbf{c})$;
- $[\![\mathbf{C}^n]\!]^{\ell_1,\ell_2}_{M,\alpha} = I(\mathbf{C}^n)$;
- $[\![\mathbf{T}^n\mathbf{t}_1 \cdots \mathbf{t}_n]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff $\langle [\![\mathbf{t}_1]\!]^{\ell_1,\ell_2}_{M,\alpha}, \ldots, [\![\mathbf{t}_n]\!]^{\ell_1,\ell_2}_{M,\alpha}\rangle \in [\![\mathbf{T}^n]\!]^{\ell_1,\ell_2}_{M,\alpha}$;
- $[\![\mathbf{t}_1 = \mathbf{t}_2]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff $\langle [\![\mathbf{t}_1]\!]^{\ell_1,\ell_2}_{M,\alpha}, [\![\mathbf{t}_2]\!]^{\ell_1,\ell_2}_{M,\alpha}\rangle \in \mathfrak{E}^{(\mathsf{level}=\ell_1)}$;
- $[\![\neg\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 0$;
- $[\![\phi \vee \psi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ or $[\![\psi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$;
- $[\![\forall\mathbf{x}\, \phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff, for every $\beta \overset{\mathbf{x}}{\sim} \alpha$, $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$;
- $[\![\forall\mathbf{X}^n\, \phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff, for every $\beta \overset{\mathbf{X}^n}{\sim} \alpha$ such that $\beta(\mathbf{X}^n) \in D_1^{n(\mathsf{level}\leq\ell_2)}$, $[\![\phi]\!]^{\ell_1,\ell_2}_{M,\beta} = 1$;

We now define several level-relative notions of truth-in-a-model and validity.

Definition 6.11. We say:

- ϕ **holds in** M ($M \vDash \phi$) $\iff \forall\ell_1\forall\ell_2\forall\alpha[\![\phi]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$;
- ϕ **holds on pairs in** M ($M \vDash^* \phi$) $\iff \forall\ell\forall\alpha[\![\phi]\!]^{\ell,\ell}_{M,\alpha} = 1$;
- ϕ **holds pointwise in the limit in** M ($M \vDash^\infty \phi$) $\iff \forall\alpha\exists\ell_1\exists\ell_2\forall\ell_3 > \ell_1\forall\ell_4 > \ell_2[\![\phi]\!]^{\ell_3,\ell_4}_{M,\alpha} = 1$;
- ϕ **holds uniformly in the limit in** M ($M \vDash^\omega \phi$) $\iff \exists\ell_1\exists\ell_2\forall\ell_3 > \ell_1\forall\ell_4 > \ell_2\forall\alpha[\![\phi]\!]^{\ell_3,\ell_4}_{M,\alpha} = 1$;
- ϕ **holds pointwise on pairs in the limit in** M ($M \vDash^{*\infty} \phi$) $\iff \forall\alpha\exists\ell_1\forall\ell_2 > \ell_1[\![\phi]\!]^{\ell_2,\ell_2}_{M,\alpha} = 1$;
- ϕ **holds uniformly on pairs in the limit in** M ($M \vDash^{*\omega} \phi$) $\iff \exists\ell_1 \forall\ell_2 > \ell_1\ \forall\alpha[\![\phi]\!]^{\ell_2,\ell_2}_{M,\alpha} = 1$;
- ϕ **is valid** $\iff \forall M\ M \vDash \phi$;

- ϕ **is valid on pairs** $\iff \forall M\ M \vDash^* \phi$;
- ϕ **is valid pointwise in the limit** $\iff \forall M\ M \vDash^\infty \phi$;
- ϕ **is valid uniformly in the limit** $\iff \forall M\ M \vDash^\omega \phi$;
- ϕ **is valid pointwise on pairs in the limit** $\iff \forall M\ M \vDash^{*\infty} \phi$;
- ϕ **is valid uniformly on pairs in the limit** $\iff \forall M\ M \vDash^{*\omega} \phi$.

We use $\Vdash$ to range over $\vDash, \vDash^*, \vDash^\infty, \vDash^\omega, \vDash^{*\infty}, \vDash^{*\omega}$.

We now prove some results:

Theorem 6.12. $\Vdash \mathbf{t}_1 = \mathbf{t}_2 \to (\phi(\mathbf{t}_1) \leftrightarrow \phi(\mathbf{t}_2))$ *just in case* $\Vdash$ *is one of* $\vDash^\infty, \vDash^{*\infty}$.

Proof. (Sketch.) $\Rightarrow$ Let γ abbreviate the schema instance $\mathsf{x}_1 = \mathsf{x}_2 \to (\mathsf{X}\mathsf{x}_1 \leftrightarrow \mathsf{X}\mathsf{x}_2)$; since γ contains only variables, we can take M to be an arbitrary model, and it suffices to find appropriate assignments on which γ receives value 0. We choose α such that $\alpha(\mathsf{x}_1) \neq \alpha(\mathsf{x}_2)$ and $\alpha(\mathsf{x}_1), \alpha(\mathsf{x}_2) \in \mathfrak{E}^{\mathsf{level}=\ell_1}$ for some ℓ_1. (There will always be some such ℓ_1.)

There will be an $\ell > \ell_1$ such that there exists some $A \in S_\ell$ with $\alpha(\mathsf{x}_1) \in A$ and $\alpha(\mathsf{x}_2) \notin A$. Set $\alpha(\mathsf{X}) = A$. It is easy to verify that, for every quantifier-level n, $[\![\gamma]\!]^{\ell,n}_{M,\alpha} = 0$ (since γ contains no quantifiers), so $\nvDash \gamma$ and $\nvDash^* \gamma$. Because we can choose ℓ_1 such that ℓ is arbitrarily high, we also have $\nvDash^\omega \gamma$ and $\nvDash^{*\omega} \gamma$.

$\Leftarrow$ Assume $\nvDash^{*\infty}$ $\mathbf{t}_1 = \mathbf{t}_2 \to (\phi(\mathbf{t}_1) \leftrightarrow \phi(\mathbf{t}_2))$. So for some instance θ of the schema, there exists some α such that, for arbitrarily high ℓ, $[\![\theta]\!]^{\ell,\ell}_{M,\alpha} = 0$. But the antecedent of θ holds at arbitrarily high ℓ only if $\alpha(\mathbf{t}_1) = \alpha(\mathbf{t}_2)$. But then there can be no A such that $\alpha(\mathbf{t}_1) \in A$ but $\alpha(\mathbf{t}_2) \notin A$, so the consequent also holds at arbitrarily high (ℓ, ℓ). Contradiction. A similar argument shows that $\vDash^\infty \mathbf{t}_1 = \mathbf{t}_2 \to (\phi(\mathbf{t}_1) \leftrightarrow \phi(\mathbf{t}_2))$. □

Theorem 6.13. $\Vdash \mathbf{t}_1 = \mathbf{t}_2 \to \forall \mathbf{X}(\mathbf{X}\mathbf{t}_1 \leftrightarrow \mathbf{X}\mathbf{t}_2)$ *just in case* $\Vdash$ *is one of* $\vDash^*, \vDash^\infty, \vDash^{*\omega}, \vDash^{*\infty}$.

Proof. (Sketch.) $\Rightarrow$ Let δ abbreviate the schema instance $\mathsf{x}_1 = \mathsf{x}_2 \to \forall X(\mathsf{X}\mathsf{x}_1 \leftrightarrow \mathsf{X}\mathsf{x}_2)$. We reuse the construction in Theorem 6.12, with ℓ defined as therein; for any such ℓ, by the reasoning above, there is some quantifier level $\ell' > \ell$ such that $[\![\delta]\!]^{\ell,\ell'}_{M,\alpha} = 0$. So $\nvDash \delta$.

Because we can choose ℓ_1 such that ℓ is arbitrarily high, there can be no ℓ_a and ℓ'_a such that for all α, $\ell^\circ > \ell_a$, and $\ell^\bullet > \ell'_a$, $[\![\delta]\!]^{\ell^\circ,\ell^\bullet}_{M,\alpha} = 1$, so $\nvDash^\omega \delta$.

$\Leftarrow$ For $\vDash^\infty \delta$, $\vDash^{*\infty} \delta$, the reasoning of the second half of Theorem 6.12 applies. For $\vDash^* \delta$ and $\vDash^{*\omega} \delta$, assume that $\Vdash \mathbf{t}_1 = \mathbf{t}_2 \to \forall \mathbf{X}(\mathbf{X}\mathbf{t}_1 \leftrightarrow \mathbf{X}\mathbf{t}_2)$ fails for some instance θ. But then $[\![\theta]\!]^{\ell,\ell}_{M,\alpha} = 0$ for some α and ℓ. By the clauses for the quantifier and equality, however, if $\ell_1 = \ell_2$, then $[\![\mathbf{t}_1 = \mathbf{t}_2]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ just in case $[\![\forall \mathbf{X}(\mathbf{X}\mathbf{t}_1 \leftrightarrow X\mathbf{t}_2)]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$; thus, if the antecedent of θ holds for some α and ℓ, so does the consequent. □

Theorem 6.14. $\Vdash \forall \mathbf{X}\Phi(\mathbf{X}) \to \Phi(\mathbf{T})$ *just in case* $\Vdash$ *is one of* $\vDash^\infty, \vDash^{*\infty}$.

Proof. (Sketch.) We show the result using the dual form $\Phi(\mathbf{T}) \to \exists \mathbf{X}\, \Phi(\mathbf{X})$.

$\Rightarrow$ Let η abbreviate the schema instance $\forall \mathsf{x}(\mathsf{Y}\mathsf{x} \leftrightarrow \mathsf{Y}\mathsf{x}) \to \exists \mathsf{X}\forall \mathsf{x}(\mathsf{X}\mathsf{x} \leftrightarrow \mathsf{Y}\mathsf{x})$. Choose some α and ℓ such that $\alpha(\mathsf{Y}) \notin D_1^{1(\mathsf{level} \leq \ell)}$. Since first-order quantifiers are not evaluated level-relatively, $[\![\forall \mathsf{x}(\mathsf{Y}\mathsf{x} \leftrightarrow \mathsf{Y}\mathsf{x})]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ for all ℓ_1, ℓ_2; by construction, however, $[\![\forall \mathsf{x}(\mathsf{Y}\mathsf{x} \leftrightarrow \mathsf{Y}\mathsf{x})]\!]^{n,\ell}_{M,\alpha} = 0$ for all n. So So $\nvDash \eta$, $\nvDash^* \eta$. Since ℓ can be made arbitrarily large, $\nvDash^\omega \eta$, $\nvDash^{*\omega} \eta$.

$\Leftarrow$ Assume that there is some α such that, for all sufficiently large ℓ, $[\![\theta]\!]^{\ell,\ell}_{M,\alpha} = 0$ for some schema instance θ of $\Phi(\mathbf{T}) \to \exists \mathbf{X}\Phi(\mathbf{X})$. Since $\alpha(\mathbf{T}) \in D_1^{1(\mathsf{level} \leq \ell')}$ for some ℓ', $[\![\exists\, \mathbf{X}\Phi(\mathbf{X})]\!]^{\ell',\ell^*}_{M,\alpha} = 1$ for $\ell^* \geq \ell'$. Contradiction. So $\vDash^{*\infty} \theta$. A similar argument holds for $\vDash^\infty$. □

We now develop a semantics for $\mathscr{L}^+$.

Definition 6.15. The semantic clauses for $\mathscr{L}$ remain the same as those for $\mathscr{L}$ in Definition 6.10, except for atomic formulas other than equalities. The clause for such formulas becomes

- $[\![\mathbf{T}^n \mathbf{t}_1 \cdots \mathbf{t}_0]\!]^{\ell_1,\ell_2}_{M,\alpha} = 1$ iff $[\![\mathbf{T}^n]\!]^{\ell_1,\ell_2}_{M,\alpha}$ is defined and $\langle [\![\mathbf{t}_1]\!]^{\ell_1,\ell_2}_{M,\alpha}, \ldots, [\![\mathbf{t}_n]\!]^{\ell_1,\ell_2}_{M,\alpha} \rangle \in [\![\mathbf{T}^n]\!]^{\ell_1,\ell_2}_{M,\alpha}$;

For λ-terms, we first introduce the notation $\alpha \overset{\mathbf{x}_1,\ldots,\mathbf{x}_n}{\sim} \beta$ to indicate that α varies from β at most in $\mathbf{x}_1, \ldots, \mathbf{x}_n$. We now define

- $\|[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \phi]\|^{\ell_1,\ell_2}_{M,\alpha} = \{\langle a_1, \ldots, a_n \rangle : \exists \alpha \in \mathfrak{Assn}\ (\beta \overset{\mathbf{x}_1,\ldots,\mathbf{x}_n}{\sim} \alpha \wedge \beta(\mathbf{x}_1) = a_1 \wedge \cdots \wedge \beta(\mathbf{x}_n) = a_n \wedge [\![\phi]\!]^{\ell_1,\ell_2}_{M,\beta} = 1)\}$.

Finally, we set

- $[\![[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \phi]]\!]^{\ell_1,\ell_2}_{M,\alpha} = \|[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \phi]\|^{\ell_1,\ell_2}_{M,\alpha}$ if $\|[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \phi]\|^{\ell_1,\ell_2}_{M,\alpha} \in D_1^{n(\mathsf{level} \leq \ell_2)}$; $[\![[\lambda \mathbf{x}_1 \cdots \mathbf{x}_n \,.\, \phi]]\!]^{\ell_1,\ell_2}_{M,\alpha}$ is undefined otherwise;

The notions of truth and validity in Definition 6.11 remain the same for $\mathscr{L}^+$.

Acknowledgements

We thank Michael Caie, Myshkin Tendulkat, Tushar Menon, Sam Roberts, Juhani Yli-Vakkuri, and the members of the colloquia of Barbara Vetter (at the Freie Universität Berlin) and Carolin Antos and Leon Horsten (at the University of Konstanz). Thanks as well to the members of the Being without Foundations colloquium, in particular, Philipp Blum, Stephan Leuenberger, Stephanie Rennick, Mario Schaerli, and Lisa Vogt, as well as to the participants at the Fragmentalism in Mind and Reality workshop, Ousia, Centre for the Philosophy and Theology of Being, University of Lucerne, Ligerz, Switzerland, and the participants in the Indiscernibility-of-Identicals Arguments and Non-Standard Ontologies workshop, Department of Philosophy, University of Geneva, Ligerz, Switzerland, where versions of parts of this research were presented. Catharine Diehl's research was funded in part by the Swiss National Science Foundation, project number 182147.

References

Ariew, R. and D. Garber (Eds.) (1989). *G.W. Leibniz: Philosophical Essays*. Indianapolis: Hackett.

Armstrong, D. (1989). *A Combinatorial Theory of Possibility*. Cambridge: Cambridge University Press.

Arntzenius, F. (2003). Is quantum mechanics pointless? *Philosophy of Science 70*, 1447–57.

Arntzenius, F. (2008). Gunk, topology, and measure. *Oxford Studies in Metaphysics 4*, 225–47.

Bacon, A. and J. S. Russell (2019). The logic of opacity. *Philosophy and Phenomenological Research 99*, 81–114.

Baker, D. J. (2009). Against field interpretations of quantum field theory. *British Journal for the Philosophy of Science 60*, 585–609.

Baker, D. J. (2016). The philosophy of quantum field theory. In *Oxford Handbooks Online*. Available at DOI: 10.1093/oxfordhb/9780199935314.013.33.

Baxter, D. L. (2018). Self-differing, aspects, and Leibniz's law. *Noûs 52*, 900–20.

Bunt, H. C. (1985). *Mass Terms and Model-Theoretic Semantics*. Cambridge: Cambridge University Press.

Burge, T. (1972). Truth and mass terms. *Journal of Philosophy 69*, 263–82.

Caie, M., J. Goodman, and H. Lederman (2020). Classical opacity. *Philosophy and Phenomenological Research 101*, 524–66.

Carnap, R. (1956). *Meaning and Necessity: A Study in Semantics and Modal Logic* (2nd ed.). Chicago: University of Chicago Press.

Church, A. (1951a). The logic of sense and denotation. In P. Henle, H. M. Kallen, and S. K. Langer (Eds.), *Structure, Method, and Meaning: Essays in Honor of Henry M. Sheffer*, pp. 3–24. New York: Liberal Arts Press.

Church, A. (1951b). The need for abstract entities in semantic analysis. *Proceedings of the American Academy of Arts and Sciences 80*, 100–12.

Church, A. (1956). *Introduction to Mathematical Logic* (2nd ed.). Princeton: Princeton University Press.

Clifton, R. and H. Halvorson (2001). Are Rindler quanta real? Inequivalent particle concepts in quantum field theory. *British Journal for the Philosophy of Science 52*, 417–70.

Diels, H. and W. Kranz (1951). *Die Fragmente der Vorsokratiker, griechisch und deutsch* (6th ed.). Berlin: Weidemann.

Dummett, M. (1963). The philosophical significance of Gödel's theorem. *Ratio 5*, 140–55.

Dummett, M. (1991). *Frege's Philosophy of Mathematics*. Cambridge, Mass.: Harvard University Press.

Fine, K. (1982). Acts, events, and things. In W. Leinfellner, E. Kraemer, and J. Schank (Eds.), *Language and Ontology: Proceedings of the 6th International Wittgenstein Symposium, 23rd to 30th August 1981, Kirchberg Wechsel, Austria*, pp. 97–105. Vienna: Hölder-Pichler-Tempsky.

Fine, K. (2003). The non-identity of a material thing and its matter. *Mind 112*, 195–234.

Fraser, D. (2008). The fate of "particles" in quantum field theories with interactions. *Studies in History and Philosophy of Modern Physics 39*, 841–59.

Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik 100*, 25–50.

Gibbard, A. (1975). Contingent identity. *Journal of Philosophical Logic 4*, 187–221.

Halvorson, H. and R. Clifton (2002). No place for particles in relativistic quantum theories? *Philosophy of Science 69*, 1–28.

Hawthorne, J. (2003). Identity. In M. Loux and D. Zimmerman (Eds.), *The Oxford Handbook of Metaphysics*, pp. 99–130. Oxford: Oxford University Press.

Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*, pp. 481–563. Oxford: Oxford University Press.

Kiss, T., F. J. Pelletier, and H. Husić (Eds.) (2021). *Things and Stuff: The Semantics of the Count-Mass Distinction*. Cambridge: Cambridge University Press.

Kleinschmidt, S. (2007). Some things about stuff. *Philosophical Studies 135*, 407–23.

Kripke, S. (1971). Identity and necessity. In M. K. Munitz (Ed.), *Identity and Individuation*, pp. 135–64. New York: New York University Press.

Landman, F. (1989a). Groups I. *Linguistics and Philosophy 12*, 559–605.

Landman, F. (1989b). Groups II. *Linguistics and Philosophy 12*, 723–44.

Laycock, H. (2006). *Words without Objects: Semantics, Ontology, and Logic for Non-Singularity*. Oxford: Clarendon Press.

Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. In R. Bäuerle, C. Schwarze, and A. von Stechow (Eds.), *Meaning, Use, and Interpretation of Language*, pp. 302–23. Berlin: de Gruyter.

Loets, A. (2021). Qua objects and their limits. *Mind 130*, 617–38.

Magidor, O. (2011). Arguments by Leibniz's law in metaphysics. *Philosophy Compass 6*, 180–95.

Malament, D. (1996). In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In R. Clifton (Ed.), *Perspectives on Quantum Reality: Non-Relativistic, Relativistic, and Field-Theoretic*, pp. 1–10. Dordrecht: Kluwer.

Markosian, N. (2004). Simples, stuff, and simple people. *Monist 87*, 405–28.

Markosian, N. (2015). The right stuff. *Australasian Journal of Philosophy 93*, 665–87.

Marmodoro, A. (2017). *Everything in Everything: Anaxagoras's Metaphysics*. Oxford: Oxford University Press.

McGee, V. (1996). Logical operations. *Journal of Philosophical Logic 25*, 567–80.

McKay, T. J. (2015). Stuff and coincidence. *Philosophical Studies 172*, 3081–100.

Moltmann, F. (1997). *Parts and Wholes in Semantics*. Oxford: Oxford University Press.

Montague, R. (1973). Comments on Moravcsik's paper. In J. Hintikka, J. Moravcsik, and P. Suppes (Eds.), *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, pp. 289–94. Dordrecht: Reidel.

Moravcsik, J. (1973). Mass terms in english. In J. Hintikka, J. Moravcsik, and P. Suppes (Eds.), *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, pp. 263–85. Dordrecht: Reidel.

Quine, W. (1961). Reference and modality. In *From a Logical Point of View: Nine Logico-Philosophical Essays*, pp. 139–69. New York: Harper and Row.

Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford: Oxford

University Press.

Russell, B. (1918). The philosophy of logical atomism [I]. *Monist 28*, 32–63.

Russell, B. (1919a). The philosophy of logical atomism [II]. *Monist 29*, 190–222.

Russell, B. (1919b). The philosophy of logical atomism [III]. *Monist 29*, 345–80.

Russell, B. (1919c). The philosophy of logical atomism [IV]. *Monist 29*, 495–527.

Salmon, N. (1986). *Frege's Puzzle*. Cambridge, Mass.: MIT Press.

Saul, J. (2007). *Simple Sentences, Substitution, and Intuitions*. Oxford: Oxford University Press.

Schaffer, J. (2010). Monism: The priority of the whole. *Philosophical Review 119*, 31–76.

Sher, G. (1991). *The Bounds of Logic: A Generalized Viewpoint*. Cambridge, Mass.: MIT Press.

Soames, S. (1989). Direct reference and propositional attitudes. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*, pp. 393–419. Oxford: Oxford University Press.

Szabó, Z. G. (2003). On qualification. *Philosophical Perspectives 17*, 385–414.

Tarski, A. (1986). What are logical notions? *History and Philosophy of Logic 7*, 143–54.

Turner, J. (2011). Ontological nihilism. *Oxford Studies in Metaphysics 6*, 3–54.

Wallace, D. (2011). Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Modern Physics 42*, 116–25.

Wiggins, D. (1967). *Identity and Spatio-Temporal Continuity*. Oxford: Blackwell.

Williamson, T. (2002). Vagueness, identity and Leibniz's law. In A. Bottani, M. Carrara, and P. Giaretta (Eds.), *Individuals, Essence and Identity: Themes of Analytic Metaphysics*, pp. 273–303. Dordrecht: Kluwer.

Zimmerman, D. (1997). Coincident objects: Could a "stuff ontology" help? *Analysis 57*, 19–27.