Logical Atomism

Verónica Gómez & Ezra Rubenstein

1. Motivations

'Logical atomism' is a vague term for an interconnected set of views, ranging across many areas of philosophy. But its core, as we understand it, is the metaphysical thesis that fundamental reality is logically simple. On this view, reality consists fundamentally in the instantiation of fundamental properties and relations by other fundamental entities. Perhaps there are brute and irreducible 'atomic' facts like *Sparky is negatively charged* and *Sparky occupies Pointy*. But logically complex facts, like *Sparky is either negatively charged or positively charged*, or *All electrons are negatively charged*, are not brute or irreducible.[1] Relatedly, if *being negatively charged* is a fundamental property, then it is itself simple: it cannot be analyzed in terms of simpler properties, via logical operations like conjunction and disjunction. All of reality is in some sense 'constructed from' logically simple ingredients.[2]

This is only an inchoate picture, and we will consider various ways of developing it in due course. But the picture is undeniably attractive, as reflected by its centrality to the analytic tradition. It featured prominently in the tradition's origins, especially through Russell (1918, 1924) and Wittgenstein (1922). Indeed, it may lie behind the very term 'analytic':

> … the philosophy I espouse is analytic, because it claims that one must discover the simple elements of which complexes are composed, and that complexes presuppose simples, whereas simples do not presuppose complexes … You will note that this philosophy is the philosophy of logical atomism. (Russell 1911: 94)

[1] Perhaps (contra Russell 1924: 377-8) reality has fundamental 'higher-order' aspects, taking forms like 'N (F, G)', or 'F is more similar to G than to H': this is consistent with logical atomism as we understand it.

[2] The principle that all fundamental entities are logically simple is naturally paired with the converse principle that all logically simple entities are fundamental. On the resulting view, all 'grounding' involves the reduction of logically complex entities to their simple constituents. For example, there are no atomic truths involving 'composite entities', like mereological fusions or sets, which are grounded in more fundamental atomic truths involving their parts or elements. This is our preferred picture, but it is not an official component of logical atomism as we understand it here.

In the middle of the twentieth century, logical atomism continued to exert a powerful methodological influence through 'logical positivists' such as Carnap (1928), whilst being stripped of its metaphysical significance. As positivism waned, the metaphysics returned: the overall world-views of Lewis (1986) and Armstrong (1989, 1997, 2004), for example, were undoubtedly inspired by logical atomism. More recently, a revival of interest in the view has been fueled by Fine's (2012) work on grounding, and the 'higher-order' movement in metaphysics (Dorr 2016, Bacon 2020, Fritz & Jones 2024).[3]

Logical atomism's central place in analytic philosophy is well-deserved. Most straightforwardly, complexity cries out for analysis:

> I confess it seems obvious to me (as it did to Leibniz) that what is complex must be composed of simples… (Russell 1924: 377)

Just as mereological complexity indicates construction from parts, and set-theoretic complexity indicates construction from elements, logical complexity indicates construction from atomic propositions.[4] Analysis by decomposition is a powerful tool. The natures of the various chemical elements, and their inter-relations, are elegantly explained via their construction from simpler constituents. Logic seems to similarly explain the natures of propositions, and their inter-relations.

Another intuitive motivation for logical atomism is the admittedly vague conviction that logic is 'representational' rather than 'worldly': it reflects something about the nature of the representation relation, not the reality represented.[5] When one learns logical principles, such as the introduction and elimination rules for conjunction and disjunction, they do not seem akin to laws of physics, as

---

[3] Recent relevant discussion includes Dorr 2002, Correia 2010, Plate 2016, McSweeney 2019a, 2020, Goodman 2023, Jackson 2024, and Fritz & Bacon ms.

[4] Not that logical atomists must reject fundamental fusions and sets, or assimilate these kinds of complexity to logical complexity; there may of course be important differences between them.

[5] Raven (2020) provides a characterization of logic's 'unworldliness' quite different from that developed below.

one might expect if the logical notions they concern are fundamental. Rather, they feel more like semantic rules, describing how logical language works.[6]

This is reminiscent of 'conventionalism': the once-popular view that logical truths are somehow a matter of 'our determination to use symbols in a certain fashion' (Ayer 1936: 31). It shares with conventionalism the sense that reality makes no distinctively logical contribution to the truth of logically complex sentences. However, on our preferred understanding—developed below—logical truths are nonetheless made true by reality, are not 'about' representation, and may be deep, substantive, and objective. Moreover, our view has no distinctive normative implications: there may be a uniquely correct way to reason using logical notions. Logic is 'just' a way of representing reality, but the nature of representation makes it an especially natural way of doing so. *Pace* conventionalists, we see no reason to think that we control the representation relation itself, as opposed to the linguistic entities that we tie to one end of it. We manufacture and launch the anchor; the meta-semantic currents decide where in the seabed it sticks.

Further motivations for logical atomism are negative in character: they derive from challenges to the opposed view of 'Logical Realism'. According to this view, prominently defended by Sider (2011), (i) some truth-functionally complete family of logical connectives (say, ∧ and ¬) and some first-order quantifier (say, ∃) are fundamental notions, and (ii) any truth that can be stated using these and other fundamental notions is a fundamental truth. (We say more about 'fundamental' truths/notions below; a rough understanding suffices for now.)

As Sider (2011: §10.2) and McSweeney (2019a) discuss, *Logical Realism* faces an arbitrariness worry. Among the truth-functional connectives, should we favor conjunction, disjunction or the conditional? And among the quantifiers, should we favor the universal or the existential? Of the many families of notions from which the rest can be defined, none seem to be privileged.

[6] This motivation may be in tension with the intuition that logical complexity, like mereological complexity, is built, insofar as the latter suggests that it is worldly. The various versions of logical atomism discussed below each seem to vindicate one of these guiding intuitions more clearly than the other.

Arbitrariness worries take many forms. Ours isn't based on any 'principle of indifference': logical realists may divide their credence evenly between 'symmetric' theories of logical structure. Nor is it based on optimism that metaphysical questions are in principle answerable: reality's logical structure may be an unsolvable mystery. Nor is it based on some iron-clad 'principle of sufficient reason': reality's logical structure may simply be brute (McSweeney 2019a). But we do think that arbitrariness is a defeasible guide to fundamental structure: theories that minimize it are more plausible, all else equal. When we face many options, and none seem privileged, this indicates that some deeper structure may underlie each of them.

As Sider notes, logical realists may avoid arbitrariness by counting all truth-functional connectives and quantifiers as fundamental. But why stop at one- and two-place connectives? What about three-place connectives, or even infinitary connectives? And why stop at ∃ and ∀? What about 'nothing' (as in, *Nothing is positively charged*), or quantifiers binding multiple variables (as in, *Some things are related by occupation*)? And, as Donaldson (2015) asks, what about other kinds of logical notions, like predicate functors? Full egalitarianism involves enormous ideological redundancy: expressing all truths would only require a tiny fraction of the fundamental notions. Like arbitrariness, ideological redundancy is a defeasible guide to fundamental structure. Perhaps some redundancy is unavoidable, but massive redundancy would be surprising. The vices of arbitrariness and redundancy trade off. Logical realists may have to live with a significant amount of both.

A second challenge involves redundancy of truths, rather than notions. Given *Logical Realism*, some fundamental truths are truth-functionally constructed, and thus metaphysically necessitated by the truths they are constructed from.[7] For example, the fundamental truths include conjunctive and disjunctive truths (or at least, truth-conditional equivalents), which are necessitated by their conjuncts/disjuncts. We take this to be undesirable excess: there are more fundamental truths than needed to fully describe reality.

[7] In particular: for any n-ary truth-function f which outputs a truth when all inputs are truths, and any n fundamental truths pp, some fundamental truth is truth-conditionally equivalent to f(pp).

This challenge needn't be tied to the controversial 'combinatorialist' thesis that no fundamental aspects of reality are necessarily connected.[8] For well-known reasons, this is hard to maintain in full generality. For example, it would be violated by fundamental truths involving necessarily transitive relations like *at-least-as-massive-as*, or by any necessary fundamental truths, such as those concerning the higher-order structure of physical quantities (Sider 2011: 218, Eddon 2013). Instead, like ideological redundancy, we treat redundancy among truths as a defeasible epistemic guide. Redundancy is sometimes tolerable, especially when sleeker foundations seem arbitrary (as with the *at-least-as-massive-as* truths). But in this case, the atomic truths provide a natural foundation.

A third challenge for *Logical Realism* derives from the idea that fundamental truths are structured: roughly, any fundamental truth is built from fundamental ingredients according to a unique 'recipe'. If *Mass* and *Charge* are fundamental properties, the result of applying *Mass* to some fundamental individual is distinct from the result of applying *Charge* to any fundamental individual. Likewise, if *Sparky* and *Pointy* are fundamental individuals, the result of applying some fundamental property to *Sparky* is distinct from the result of applying any fundamental property to *Pointy*.[9]

More precisely, let *Fun*(*X*) express that *X* is a fundamental entity. We define the notion '*Pure*' recursively via the following schemas:

$$Fun(X) \rightarrow Pure(X)$$

$$Pure(X) \wedge Pure(\hat{x}) \rightarrow Pure(X(\hat{x}))$$

where $\hat{x}$ is a sequence, which is *Pure* just in case all of the entities in it are *Pure*. We can then state the idea schematically as follows:

[8] For development and discussion of this thesis, see e.g. Armstrong 1989, Wilson 2010, Wang 2013, Russell & Hawthorne 2018, Bacon 2020.

[9] If the fundamental truths involve order-sensitive relations, then this seems to require privileging some relation over its permutations (presumably, for example, *Before*(*a*, *b*) = *After*(*b*, *a*).) We take this to be an important motivation for pursuing approaches on which there are only fundamental plural properties (Dorr 2004: 190) or 'neutral' relations (Fine 2000, Gilmore 2013).

Fundamental Structure:

$Pure(X) \wedge Pure(\hat{x}) \wedge Pure(Y) \wedge Pure(\hat{y}) \wedge X(\hat{x}) = Y(\hat{y}) \rightarrow X = Y \wedge \hat{x} = \hat{y}$

where '$\hat{x} = \hat{y}$' abbreviates the conjunction $x_1 = y_1 \wedge \ldots \wedge x_n = y_n$ when $\hat{x}$ and $\hat{y}$ are sequences corresponding to the same sequence of types, and a contradiction otherwise.

Given the way lambda-terms reduce, *Fundamental Structure* entails that lambda-abstraction does not preserve purity in the relevant sense. For example, by 'Immediate Beta-equivalence' (Dorr 2016: 52), $\lambda x.Rax(a) = \lambda x.Rxx(a)$ for any binary relation R. However, we take this consequence to be natural (even for logical realists). The guiding idea is that there is a unique way to build a fundamental proposition from fundamental ingredients through application, not that there is a unique way of 'decomposing' that proposition by abstracting 'ingredients' out of it. The observation that $\Diamond Fa$ may be formed either by applying possibility to Fa or by applying possible F-ness to a does not strike us as violating the core intuition that $\Diamond Fa$ is structured: these seem to be two descriptions of a single recipe.

Given *Fundamental Structure*, *Logical Realism* conflicts with the observation that propositions lack unique logical forms. For example, the following identifications all seem fairly plausible:

1. $\text{Fa} \wedge \text{Ga} = \text{Ga} \wedge \text{Fa}$[10]
2. $\text{Fa} = \text{Fa} \vee \text{Fa}$
3. $(\text{Fa} \wedge \text{Ga}) \wedge \text{Gb} = \text{Fa} \wedge (\text{Ga} \wedge \text{Gb})$
4. $\neg\exists x\, \text{F}x = \forall x\, \neg\text{F}x$
5. $\text{Fa} = \exists x\, \text{F}x \wedge \text{Fa}$
6. $\text{Fa} = \neg\neg\text{Fa}$

[10] This parallels failures of *Fundamental Structure* involving symmetric relations (like *a is next to b* = *b is next to a*). Note, however, that a natural solution—taking conjunction to be a plural property of propositions—is inconsistent with *Fundamental Structure* for Cantorian reasons: it would generate an injection from pluralities of propositions to propositions.

These identifications suggest that the way logically complex truths are constructed is very differently from the way the atomics are (Wittgenstein 1922: §4.0621, §5.44; Ramsey 1927: 162). Logical notions do not behave like fundamental properties and individuals: they are more like eerie ghosts than solid building blocks, disconcertingly 'dissolving into' propositions rather than leaving their mark. They can add nothing (as in 2 and 5) or even self-annihilate (as in 6).[11]

Unlike those based on redundancy and arbitrariness, this challenge is 'metaphysical'. *Fundamental Structure* is intended to describe the nature of fundamentality, not merely to defeasibly constrain fundamental theorizing. It might be seen as a successor to combinatorialism: perhaps different constructions from fundamental ingredients can be necessarily connected, but they cannot be identical.

Of course, logical atomism is not the only alternative to *Logical Realism*, but it is the most natural alternative which avoids the problems described above. First, as McSweeney (2019a) suggests, one might posit some fundamental logical notions other than the familiar truth-functional connectives and quantifiers (perhaps they are 'altogether alien', in Eklund's (2024) sense).[12] This is hard to assess in the abstract, but insofar as these unfamiliar notions are properly counted as 'logical' we would expect the same problems to arise. Second, one might allow some special logically complex truths to be fundamental, such as totality truths, laws, or certain basic negative truths, whilst avoiding enormous redundancy.[13] But this is at best a fallback position, in case logical atomism cannot be had in full generality. Third, one might posit logical 'modes of combination' alongside application, instead of fundamental logical notions.[14] However, this still yields

[11] A version of the Russell-Myhill paradox follows from *Fundamental Structure* together with an extension of *Logical Realism* to higher-order quantifiers and identity, assuming that i) lambda-abstraction preserves purity, and ii) purity is itself pure. We don't see this as a motivation for logical atomism, however. We don't find either i) or ii) attractive, and besides, *Logical Realism* needn't extend to higher orders (Sider 2011: §9.15).

[12] As McSweeney (2019b:7) notes, object-predicate structure itself might be regarded as 'logical'. However, the challenges described above do not seem to motivate the further step of removing even this structure from fundamental reality.

[13] Russell (1918) ultimately admitted both negative and general facts, and similar views have more recently been defended by Fine (2012) and Jackson (2024).

[14] This seems to be Bacon's (2020) picture, and it arguably captures an important strand of Russell's thinking.

redundancy among the fundamental truths, and perhaps it faces a version of the other two problems.

One might also 'reject' *Logical Realism* without endorsing logical atomism by adopting the meta-metaphysical view that the debate is verbal, meaningless, or somehow defective. As Sider (2011: 217) admits, 'You don't have to be a logical positivist to feel that something is wrong with these questions.' However, as he goes on to observe, it is the question of *which* logical notions are fundamental that is especially suspect. Logical atomism vindicates this reaction by rejecting the question's presupposition. More generally, the arguments presented above suggest that the question of *Logical Realism* should be answered negatively, not that it shouldn't be asked. In any case, we adopt the working assumption that the notion of fundamentality applies meaningfully to logical notions and truths, and the next section considers different ways of fleshing out this idea.

2. Variations

2.1 Generative atomism

We started with the slogan that fundamental reality is logically simple. A standard way of theorizing fundamentality is in terms of ground: the fundamental facts are the ungrounded facts. Thus, the slogan might be interpreted as the view that all logically complex facts are grounded. Given a non-factive notion of ground, this view naturally extends: all logically complex propositions (or states of affairs) have (non-factive) grounds. Adding that grounding is necessarily well-founded, this entails that, necessarily, all logically complex truths are fully grounded in atomic truths. Call this 'grounding atomism'.

The term 'ground' is used flexibly in contemporary metaphysics. This section focuses on a 'generative' conception, on which: i) grounding resembles causation in its structure and conceptual role — the grounds generate or produce the grounded, much as causes generate or produce their effects; ii) grounding is worldly — it relates equally real and genuinely distinct facts or entities; and iii) talk of ground is joint-carving — it limns the world's mind-independent and objectively privileged structure. Perhaps no philosopher embraces all these ideas, but many grounding enthusiasts are disposed to apply the concept in ways that are naturally justified by some substantial portion of them.[15]

Applying this conception to grounding atomism yields 'Generative Atomism': the view that, necessarily, all logically complex truths are generated by atomics. The principles governing logical generation are familiar from Fine (2012; see also Correia 2010, Rosen 2010): conjunctions are generated by their conjuncts, disjunctions by their (true) disjuncts, existentials by their (true) instances, and so on.[16] Plenty of work would be needed to develop a comprehensive and consistent view along these lines (Fine 2010, Fritz 2021). But, even before climbing into the trenches, we doubt that victory would be worth the war: logical atomism's initial attractions do not motivate *Generative Atomism*. (The view may, of course, have motivations other than those described above.)

Importantly, generated logically complex facts are not thereby second-rate or unreal. Just as later events may be as real as the earlier events which cause or otherwise determine them 'through time', derivative facts may be as real as the prior facts which generate them 'up the levels'. The differences between generation and causation seem irrelevant in this regard. All *Generative Atomism* delivers is that logically complex facts never lie at the foundation of a certain dependence structure. If anything, this dependence structure—containing, as it does, worldly, logically complex entities—seems to constitute irreducible logical structure. This may be reconciled with the letter of *Generative Atomism*: perhaps the generative connections are themselves somehow generated by atomic facts. But its tension with logical atomism's spirit can be appreciated by considering the view's initial motivations.

First, the arbitrariness dilemma. Generative atomists face the question: which logical structures do the generated logically complex facts have?[17] To illustrate, consider the proposition $p \vee q$. On one view, this proposition is disjunctive, and so is (non-factively) immediately generated by p and by

[15] We take this generative conception to be suggested by the work of Fine (2012), Bennett (2017), Rosen (2010), Schaffer (2016) and Skow (2016), amongst others.
[16] For skepticism, see Turner 2016 and McSweeney 2020.
[17] Cf. McSweeney 2019a: 134.

q. On another view, this proposition is 'really' the negative proposition $\neg(\neg p \wedge \neg q)$, and so is immediately generated by $\neg\neg p$ and by $\neg\neg q$ (which are in turn generated by p and by q). And one can imagine many more views besides.[18] Thus, the logical realist's question of which notions are fundamental becomes the question of how logically complex propositions are generated. On one hand, generative atomists might recognize many distinct yet logically equivalent propositions, all generated in different ways; on the other, they might arbitrarily collapse some propositions into others.

This point might be made using the idea that generative connections obey 'metaphysical laws', (Schaffer 2017). *Generative Atomism* then presumably entails that there are certain irreducible laws of logical generation, and the question arises of which logical notions these laws involve. For example, is there an irreducible law for conjunction, and another for disjunction? If only a few special notions are involved, these laws seem arbitrary; if many are involved, they seem redundant.

Second, the redundancy of truths. There are no violations of combinatorialism—the generated facts are non-fundamental—and the necessary connections are not brute, but explained by underlying generative connections. However, the underlying concern remains. In addition to the atomics, this view posits an entire hierarchy of logically complex facts. The resulting vision strikes us as uneconomical. It fails to vindicate the natural thought that the atomics already provide a 'complete inventory' of reality.

Arguably, the idea that non-fundamentalia are 'no addition to being' (Armstrong 1997: 12) doesn't fit the generative conception of ground (Rubenstein 2024: §3.2). Nonetheless, grant that generated entities do not count against a theory's ontological parsimony: when tallying *Generative Atomism*'s ontological commitments, the logically complex facts go ignored (Schaffer 2015, Bennett 2017: §8.2.2). And let the generative connections themselves unfurl upwards from the atomics (Bennett 2011), or spring from autonomous essences (Dasgupta 2014), or even be created *ex nihilo* (Litland

[18] These alternatives collapse if propositions just are their double-negations. But generative atomists naturally view p as generating, and so distinct from, $\neg\neg p$ (Fine 2012: 63). Besides, it is far from a general remedy: the problem arises for all logical forms, including negation-free forms such as $p \wedge (q \vee r)$ and $(p \wedge q) \vee (p \wedge r)$.

2017). Still, there surely remains a sense in which the theory is less economical than one which simply excises the hierarchy altogether.[19] If *Generative Atomism* is preferable to such an austere picture, it must be for other reasons: presumably, that something indispensable has been lost.

Third, *Fundamental Structure*. Generative atomists allow ungenerated propositions to be structured, since they are all atomic. But their view suggests that non-fundamental propositions are structured too. On a natural interpretation, these propositions are generated by successively applying 'construction-operations' to the atomics, where these operations 'leave their mark' in the resulting propositions, just as the fundamental ingredients do. Thus, each generated proposition has a unique recipe, describing how it is constructed by applying operations to atomics. (Compare the idea that a and b's mereological fusion and their doubleton-set are distinct, since they result from different construction-operations.) This is in tension with identifications like those listed above, in which the fact on the left seems either constructed differently from the fact on the right (e.g. $Fa \wedge (Ga \vee Gb) = (Fa \wedge Ga) \vee (Fa \wedge Gb)$), or not constructed at all (e.g. $Fa = Fa \wedge Fa$).[20]

Finally, the (admittedly elusive) conviction that logic is 'representational' rather than 'worldly' is perhaps most strikingly neglected. If anything, treating logical complexity as generated embeds logic deeply into reality's fabric. Generative atomists do not regard truth-tables as encoding representational rules, but rather as encoding the productive principles by which certain worldly entities give rise to others. An alternative interpretation of logical atomism would be welcome.

2.2 Structural Atomism

Whilst *Generative Atomism* invites the idea that some logical notions carve reality perfectly at the joints, we deny that any do. No logical connective or quantifier is joint-carving, and no properties,

[19] Cf. Turner 2016: §5. The remaining complexity might be characterized in terms of the laws of logical generation, which count against the theory's 'explanatory economy' (even if they are not themselves fundamental). Compare Rubenstein forthcoming-a: §1.3.

[20] *Generative Atomism* doesn't entail the structured view: perhaps the operations themselves are unstructured (Bacon 2020), or perhaps some propositions are 'overdetermined' by distinct but converging operations.

relations, or propositions essentially involving such notions are joint-carving. Call this view 'Structural Atomism'.

We regiment 'joint-carving' talk in a relationally typed language, where <> is the type of propositions, *e* is the type of objects, and for any types $\sigma_1, \ldots, \sigma_n$, there is a type $<\sigma_1, \ldots, \sigma_n>$. Entities of type $<\sigma_1, \ldots, \sigma_n>$ apply to entities of type $\sigma_1, \ldots, \sigma_n$ (in that order) to yield a proposition. For every type τ, there is a predicate $Structural_\tau$ that combines with terms of type τ to yield a sentence whose intuitive meaning is that the entity in question belongs to reality's fundamental structure. We also employ a primitive notion of logicality for each type. *Structural Atomism* can then be expressed as the conjunction of all instances of the following schema:

$$Structural_\tau\,(X^\tau) \rightarrow \neg Logical_\tau\,(X^\tau)$$

Logicality can be characterized inductively, starting with a primitive notion of 'simplicity', and a primitive division of the simple entities into logical and non-logical.[21] The simple entities correspond to the constants of Russell's (1918) 'logically perfect' language. We take the logical constants to include truth-functional connectives, quantifiers and identity relations for each type, and any notion as similar to these as they are to each other. (Certain modal operators might be logical too, though we doubt that all modal notions are.) The non-logical entities are those which result from applying non-logical entities to non-logical entities. Since application always yields a proposition, this entails that all non-logical entities are either simple or else atomic propositions (understood to include any propositions formed from atomics through application, such as *Believes* (*a*, *Fa*).). More formally:

$$\neg\, Logical_{<>}\,(p^{<>}) \text{ iff for some types } <\sigma_1, \ldots, \sigma_n>, \sigma_1, \ldots, \sigma_n, \exists_{<\sigma 1, \ldots, \sigma n>} Y\, \exists_{\sigma 1}\, x_1 \ldots \exists_{\sigma n}\, x_n$$
$$\neg\, Logical_{<\sigma 1, \ldots, \sigma n>}\,(Y) \wedge \neg\, Logical_{\sigma 1}\,(x_1) \wedge \ldots \wedge \neg Logical_{\sigma n}\,(x_n) \wedge p = Y\,(x_1, \ldots, x_n).$$

[21] Depending on one's view of grain, one might define simplicity in terms of identification, along the lines of Dorr 2016: §9. A natural idea is that X is simple iff it is not identical to anything not 'involving' it.

For all types τ other than <>, $\neg\, Logical_{\tau}(X^{\tau}) \rightarrow Simple_{\tau}(X^{\tau})$.[22]

(We're quantifying over types here. While there are higher-order languages which allow this, our purposes don't require them. Instead, we may be interpreted as using a 'truncated' higher-order language with finitely many types, within which we can expect all structural notions to be found. In such a language, quantification into type position can be understood substitutionally.)[23]

The central notion of 'structure' is sometimes deemed obscure. We don't entirely share this concern. We think that Lewis (1983) and Sider (2011) have done much to clarify the notion, by laying out the crucial role it plays in metaphysics and in science. However, the notion is associated with many different ideas, and we needn't adopt all of them (Dorr & Hawthorne 2013). For us, the core theoretical role clusters around *simplicity*, where this has both an epistemic aspect, concerning induction, and a metaphysical aspect, concerning reduction.

The epistemic aspect is the inductive importance of 'simple' patterns. Observing various green emeralds supports the hypothesis that all emeralds are green much more than the hypothesis that all emeralds are grue. More generally, inductive learning requires treating simpler patterns as more likely, where simplicity is only revealed when these patterns are represented in structural terms. Note that, given *Structural Atomism*, a single, logically complex sentence cannot structurally represent a pattern—instead, it must somehow be represented by many atomic sentences. Nonetheless, we think that the notion of simplicity meaningfully extends to these 'collective' representations—for example, simpler representations will exhibit certain kinds of symmetries.[24]

---

[22] This notion is transparent, meaning that certain logically complex terms, such as Fa ∧ Fa, or λx Scarlet(x) ∧ Red(x), may (depending on one's view of grain) denote non-logical entities.

[23] The above principle entails that all (irreducible) lambda-terms, such as λx Rax, are logical. At least some of the motivations for denying that paradigmatic logical terms are structural extends to these terms also. The principles governing lambda-terms (such as Immediate Beta-equivalence) seem to be representational rather than worldly, adding them seems to yield ideological redundancy, and (as noted above) pure lambda-terms would violate *Fundamental Structure*.

[24] Making this idea precise is an important project which we cannot carry out here; see Gómez (ms) for further discussion.

The metaphysical aspect of simplicity is the idea that structural entities provide the basic building blocks from which all else is constructed. For example, the property of being a water molecule can be analyzed as the property of being a fusion of suitably bonded hydrogen and oxygen atoms (and the latter can in turn be reduced, to yet more fundamental entities). Thus, *Structural Atomism* fits well with Russell's (1918: 270) characterization of logical atomism as the view that 'you can get down in theory, if not in practice, to ultimate simples, out of which the world is built, and that those simples have a kind of reality not belonging to anything else'.

Without constraining 'construction', the metaphysical aspect of simplicity is not especially interesting. Non-reductive physicalists could recognize *sui generis* mind-building operations which 'construct' novel mental properties from physical properties. On the other hand, given a suitably constrained notion of construction, it is difficult to see how logical entities can be constructed without logical building blocks. One approach, alluded to above, holds that some construction is itself logical. But this isn't in the spirit of *Structural Atomism*, since there is an important sense in which it outfits reality with irreducible logical structure.

We are interested instead in an ambitious version of the reductive idea, on which all construction is 'application', yielding:

> *Structural Completeness*: All entities can be constructed by (iteratively) applying structural entities to structural entities.[25]

That is, using the notion of purity defined above:

$$\forall X^{\tau}\, Pure(X^{\tau}).$$

[25] We adapt this terminology from Bacon (2020: 566). Abstraction (or Bacon's (2023: 208) 'complication', in the setting of a functionally typed language) is naturally regarded as another kind of construction. As mentioned above, however, we regard abstraction as logical, so this is less natural from an atomist perspective.

Applying non-logical entities to non-logical entities never results in logical entities. Thus, given *Structural Completeness*, *Structural Atomism* entails 'Eliminative Atomism': the surprising thesis that there are no logical entities. We embrace this result, and turn now to developing it.

2.3 Eliminative Atomism

Not only do we maintain that logical entities aren't structural, but that they aren't anything at all. There is no conjunction, or disjunction, or the conjunctive property of being *F* and *G*, or the disjunction of *Fa* and *Gb.* Every entity of every type is non-logical.

This isn't entailed by 'first-order' nominalism: the view that there are no abstract objects. We are thinking in higher-order terms. We treat conjunction, for example, as a (purported) higher-order relation between propositions, for which the question of first-order existence can't coherently be raised. Moreover, following Prior (1971: ch.3) and Williamson (2003: §IX), amongst others, we embrace higher-order quantification. We take the inference from *Sparky is negatively charged* to *Sparky is somehow* to be valid, because *being negatively charged* is a possible value for a variable appearing in predicate position. It is 'something' in this sense, though not in the more ordinary first-order sense. Likewise for propositions: the inference from *Sparky is negatively charged* to *Something is the case* is valid because the former can be the value of a variable in sentence position. These values are what we're calling 'propositions'.

Eliminative Atomism is unsurprising as a first-order claim: of course, conjunction isn't a thing in the way a table is! But it is quite surprising when understood in higher-order terms. What we are proposing is that there are (in a higher-order sense) properties like *being negatively charged*, and propositions like *Sparky is negatively charged*, but not conjunctive properties like *being negatively charged and massive*, or existential propositions like *Something is negatively charged*. Similarly, we want to allow for properties of and relations between propositions—probability and causation, for example—but we deny that negation and conjunction are among them.

Russell's (1918: 39–40) discussion of disjunction suggests a view along these lines:

> I do not suppose there is in the world a single disjunctive fact corresponding to "*p* or *q*". […] You must not look about the real world for an object you can call "or", and say, "Now, look at this. This is 'or'." There is no such thing, and if you try to analyse "p or q" in that way you will get into trouble.[26]

For reasons discussed in §3, Russell ends up conceding that negative and general facts are 'in the world'. Similarly, Armstrong (1997) rejects disjunctive states of affairs, but recognizes states of affairs involving logically complex properties. *Eliminative Atomism* extends Russell's and Armstrong's conception of disjunction to all logical terms.[27]

Eliminative Atomism can be expressed schematically as follows:

$$\neg\exists_\tau X\, Logical_\tau(X)$$

where, as before, the only complex non-logical entities are atomic propositions, formed by application of some structural relation. Just as mereological nihilists are naturally taken to reject both mereologically complex entities and the relation of parthood (or fusion) itself (Dorr 2005: §17), so eliminative atomists reject both logically complex entities and the logical relations by which they are formed.

[26] Rhetoric aside, we take Russell to be proposing a stronger view than the commonsensical claim that disjunction is not a (visible) object.

[27] However, unlike Armstrong, we don't treat states of affairs as first-order entities, and unlike Russell, we don't reject propositions in favor of facts.

We arrived at *Eliminative Atomism* from *Structural Atomism* via *Structural Completeness*. But the reverse inference is also natural, via the principle that being structural entails existence.[28] Even without this principle, the entailment would be secured by the idea that any structural entity is indispensable, in that some propositions cannot be constructed without it. This entails that, if *Structural Atomism* is false, then there are some logical propositions. We tentatively suggest that *Structural Atomism* and *Eliminative Atomism* are equivalent. But in any case, we find each view attractive in its own right, and we want to explore a theory that validates both.

In standard higher-order logic, *Eliminative Atomism* is inconsistent. For example:

> *Logical*(∧)
> Therefore, $\exists_{<<>,<>>} X$ *Logical*(X)

> *Logical*(Fa ∧ Gb)
> Therefore, $\exists_{<>} p$ *Logical*(p)

Rejecting the meaningfulness of the logical terms in such proofs would be both absurd and self-defeating: surely, many logically complex sentences are true—including, on our own view, the negations of the conclusions above. We must therefore reject existential introduction itself.

For this reason, we adopt a version of higher-order free logic. Rather than requiring an existence premise, we employ a syntactic distinction between the 'real' terms and the rest, recursively defined as follows:

> Any non-logical constant M^{σ} is a real term.
> For any real terms $M^{<\sigma 1, \ldots, \sigma n>}$, $N_1^{\sigma 1}$, …, $N_n^{\sigma 1}$, $M(N_1, \ldots, N_n)^{<>}$ is a real term.
> Nothing else is a real term.

[28] This principle isn't completely obvious to us: we are tempted to regard application (more generally, 'modes of combination') as structural but non-existent.

We restrict existential introduction (and universal elimination) to real terms. This blocks the problematic inferences above: neither '∧' nor 'Fa ∧ Gb' is real.

Inconsistency returns if we introduce primitive non-logical predicates for logically complex properties. If *Grue* = $\lambda x.\varphi(x)$ (where the latter is complex), we can infer that *Logical*(*Grue*), and hence that $\exists X$ *Logical*(X). But we don't see this as a problem: our existential introduction rule is intended to operate in a logically perfect language, just as classical logic operates in an 'ontologically perfect' language without empty terms.[29]

Conservatives will dislike this abandonment of classical logic. We don't share the sense that exploring metaphysical theories that go beyond classical logic is misguided. Moreover, assuming that there are no first-order logical entities, our theory recovers classical first-order logic—arguably the part that receives most direct abductive support. From a metaphysical perspective, the present approach strikes us as a natural extension of classical first-order logic to higher orders.

An important concern is that our grip on the higher-order quantifiers is lost once the classical rules are rejected. Whereas all quantifiers obeying the classical rules 'collapse', there are many ways of restricting existential introduction to some of the terms in a given language. This opens the door to a higher-order analogue of Hirsch/Putnam-style quantifier variance (Hirsch 2011), on which eliminative atomists —or 'logical nihilists'—and their opponents use different but equally legitimate higher-order quantifiers, making the debate between them verbal. (This would be a meta-metaphysical 'rejection' of logical atomism, as mentioned in §1.)

We cannot properly chart this difficult territory here, but we will briefly describe two responses. A 'hard-line' response holds that there are no legitimate alternative (unrestricted) quantifiers. In the first-order case, mereological nihilists might insist that there are good arguments that e.g. Tibbles does not exist, where this does not entail that something does not exist. Similarly, we think that the motivations described in §1 support the claim that e.g. conjunction does not exist, where this

[29] A more flexible alternative approach states the inference rules using the logicality predicate itself. The syntactic approach has the advantage of revealing the continuity with first-order classical logic.

violates classical higher-order logic. In both cases, the idea is that certain terms are meaningful but empty, where the relevant sense of emptiness eludes expression with a classical quantifier.

A more concessive reply allows for alternative quantifiers but denies that they are equally joint-carving. Now, given *Structural Atomism*, no quantificational talk is perfectly joint-carving. But some quantifiers can still be relatively joint-carving (for example, some have simpler reductions). If so, then in contexts where the standards of 'perspicuity' have been raised—such as when theorizing the metaphysics of logic—we expect *Eliminative Atomism* to be true, and the claims of classical higher-order logic to be false.[30]

Both replies share the underlying idea that our grip on the relevant quantifier need not depend on logical rules, as opposed to some more primitive sense of 'reality'. Given *Structural Completeness*, this could be understood as the domain of pure entities. (Of course, this won't convince those sceptics who reject 'joint-carving' talk altogether: on their view, it is not that the various quantifiers are equally non-perspicuous, but that the question of their perspicuity cannot even be raised.)

Structural Atomism and *Eliminative Atomism* together fare much better than *Generative Atomism* in avoiding the problems that arise for *Logical Realism*.

First, we avoid any arbitrary choice of logical structure by eschewing it altogether. Mereological nihilism answers the 'special composition question' non-arbitrarily: things never compose. Likewise, *Eliminative Atomism* non-arbitrarily answers the 'special complication question': entities never complicate.

Second, *Structural Atomism* avoids logical redundancy among the 'real' truths, in that no logically complex propositions are pure. *Eliminative Atomism* arguably does even better: there are no logically complex propositions!

[30] Whether to take the hard line or the soft line turns on the meta-semantic trade-off between fit with use and fit with reality's structure (cf. Sider 2013: §4). When the latter matters more, we think that *Eliminative Atomism* prevails—but we do not know how pervasive such contexts are.

Third, structural atomists can maintain *Fundamental Structure* whilst avoiding an implausibly fine-grained conception of logically complex propositions. Given *Structural Completeness*, *Fundamental Structure* entails:

Quantified Structure: $\forall X \forall Y \forall \hat{x}\ \forall \hat{y}\ X(\hat{x}) = Y(\hat{y}) \rightarrow X = Y \wedge \hat{x} = \hat{y}$.

In our preferred logic, *Quantified Structure* is consistent.[31] Of course, we don't accept all its 'instances': for example, we reject $(p \wedge p) = (p \vee p) \rightarrow \wedge = \vee$. But this respects *Quantified Structure*, since logical terms aren't real.

Returning to logical atomism's positive motivations, *Eliminative Atomism* renders the intuition that all logical complexity is built vacuously true. But the next section suggests a non-vacuous sense in which logical complexity is built: the ways in which logically complex language represents the world can themselves be constructed from simple ingredients. Finally, the inchoate idea that logic is not 'worldly' is captured well by *Structural Atomism* (disjunction itself is like 'grue'), and even better by *Eliminative Atomism*. And the next section can be understood as developing the positive suggestion that logic is 'representational'.

3. Logically complex truths

3.1 The problem

Without logically complex propositions to express, how can logically complex sentences be meaningful and true? Consider:

i) 'Snow is white and grass is green' is true.
ii) A sentence S is true if and only if, for some proposition *p*, S expresses *p* and *p*.
iii) If 'Snow is white and grass is green' expresses a proposition, it expresses a logically complex proposition.

[31] This could be shown using the toy model described in §3.

Therefore, there are logically complex propositions.[32]

We don't think eliminative atomists should deny (i): we have been sincerely asserting, and will continue to sincerely assert, logically complex sentences. And (iii) is very plausible, at least given a standard conception of 'expressing' a proposition.

We reject (ii). Our idea is that each logically complex sentence bears representation relations to pluralities of atomics, and is true if and only if it represents some plurality *pp*, and all the *pp* obtain. In this sense, a sentence represents its 'potential truth-makers'. Our proposal may be understood through a creation story:

> On the first day, God created two objects, and denoted them a and b. On the second, day, God created three alternative properties, and denoted them F, F^* and F^{**}. On the third day, God created two alternative two-place relations, and denoted them R and R^*. On the fourth day, God made 14 atomic propositions, which She denoted with atomic sentences Fa, Fb, F^*a, F^*b, $F^{**}a$, $F^{**}b$, Raa, Rbb, Rab, Rba, R^*aa, R^*bb, R^*ab, R^*ba. On the fifth day, She distributed the properties over the objects, thereby making Fa and F^*b true. On the sixth day, She distributed relations over the objects, thereby making Rab, Rba, R^*aa, and R^*bb true… And, on the seventh day, while restfully contemplating her creation, God had her first logically complex thought: $Fa \wedge F^*b$.

Assuming no facts were added to the world after the sixth day, did God think truly on the seventh day? Or did She think falsely or meaninglessly? If you share our intuition that logic is 'representational', you may share our sense that God thought truly, with no further logical relation or fact required. But what then does the truth of her thought consist in?

A natural idea is that, while the sentence $Fa \wedge F^*b$ denotes no single aspect of reality, it can still represent truly by representing multiple aspects of reality. Thus, in our story, $Fa \wedge F^*b$ is true

[32] Turning this argument against *Structural Atomism*, by adding *Structural Completeness*, captures the idea that reality provides the ingredients for anything that can be said.

because it represents the plurality of atomic propositions *Fa*, F^*b, and these propositions obtain. The rest of this section shows how this basic idea can be generalized to other forms of logical complexity, building on a version of 'truth-maker semantics' (van Fraassen 1969, Fine 2017a, 2017b). More specifically, our aim is to show how, given a simple relation of denotation connecting God's atomic sentences to atomic propositions, we can define a more expansive relation of 'representation' that gives complex sentences truth-conditions.

For concreteness and simplicity, we focus on the logically perfect language of our fictional God: the 'ur-language'. The resulting semantics is not completely satisfying, for two reasons. First, it is not obvious that the same kind of approach could be applied to our own language, which presumably is not logically perfect. Second, we are pretending that denotation itself is simple. We discuss these issues in §3.4.

3.2 Boolean connectives

Our starting point is the idea that the conjunction $\mathrm{Fa} \wedge \mathrm{F^*b}$ represents the propositions *Fa*, F^*b. What about disjunctions like $\mathrm{Fa} \vee \mathrm{F^*b}$?

One idea is to distinguish 'disjunctive representation' and 'conjunctive representation': a disjunctive sentence disjunctively represents its disjuncts, whereas a conjunctive sentence conjunctively represents its conjuncts. However, this approach does not generalize to more complex sentences. Take $(\mathrm{Fa} \wedge \mathrm{F^*a}) \vee \mathrm{F^*b}$. It can't disjunctively represent $Fa \wedge F^*a$, since no such proposition exists!

Alternatively, we might introduce a new kind of composite entity: a fusion of the atomics *Fa*, F^*b (or their set). Conjunctions could then represent fusions, whilst disjunctions represent corresponding pluralities. But these composite entities return the very problems which logical atomism was supposed to avoid. Their obtaining seems objectionably redundant and, moreover, arbitrary, insofar as there are many candidate composition operations.[33]

[33] Why not, for example, posit 'disjunctive fusions', which obtain so long as one of their constituent propositions obtains? Compare Fine 2017a: §4.

Instead, we suggest that representation is not only plural, but one-many: a single sentence can represent many pluralities of atomics. In particular, $Fa \lor F^{*}b$ represents the single-membered plurality *Fa*, and (separately) the single-membered plurality $F^{*}b$. We also take it to represent the plurality *Fa*, $F^{*}b$, since this will help to ensure plausible equivalences, such as $Fa \lor F^{*}b \approx Fa \lor F^{*}b \lor (Fa \land F^{*}b)$. Since truth amounts to representing some plurality of atomics all of which obtain, this ensures that the sentence $Fa \lor F^{*}b$ is true if and only if *Fa* obtains or *F*b* obtains.

Negation is trickier.[34] Prima facie, it is hard to see what atomics could make ¬Fb true. According to one view, the fact that *Fa*, $F^{*}b$, ... are all the true atomics makes it true. There are at least two problems with this proposal. Firstly, this fact is arguably logically complex, and secondly, many of the positive atomics are intuitively irrelevant to the truth of ¬Fb. According to a second view, atomists should make an exception for negative atomics (Russell 1918: 211). Given *Fundamental Structure*, the idea that negation itself is a fundamental entity would violate the attractive idea that a proposition is equivalent to its double-negation. Perhaps this can be avoided by weakening *Structural Completeness*, to allow propositions to be formed through 'negative application' as well as application itself (Jago & Barker 2012). But this would still compromise the attractively pure vision of reality consisting in the instantiation of simple properties and relations.[35]

Our preferred view takes properties and relations to come in families of mutually exclusive and jointly exhaustive 'alternatives', such as, for example, positive, negative, and neutral charge. We represent families of alternatives with families of syntactically related predicates. Thus, God's predicates F, F^{*} and F^{**} aptly represent a family of three alternative properties, and R and R^{*} aptly represent a family of two alternative relations.[36]

---

[34] Witness the large contemporary literature on the problem of negative truths e.g. Molnar 2000, Lewis 2001, and Armstrong 2004: ch.5.

[35] Russell reports that his endorsement of negative facts 'nearly produced a riot' at Harvard. The students' representative, Demos, claimed (1917: 189) that he 'once undertook a fairly systematic interrogation on the matter among intelligent acquaintances… and they were practically unanimous in their testimony that they had never encountered a negative fact'. (Admittedly, though, his survey's results may reflect the sheltered life of Harvard students.)

[36] A further question is whether this structure consists in the holding of some simple relation between properties and relations (Wang 2013), or whether it is better understood as akin to the 'grammatical' division

Given this structure, we say that two atomic propositions are alternatives when they attribute alternative properties or relations to the same entities (taken in the same order). Thus, atomics will themselves divide into families, with exactly one proposition from each family obtaining. We can then think of ¬Fb as representing each alternative to *Fb*—$F^{*}b$ and $F^{**}b$—and so true when either of these alternatives obtains.[37] Historically, of course, logical atomists have insisted that atomic propositions are 'freely recombinable' (Wittgenstein 1922: §4.211). However, we do not share this commitment to combinatorialism, and we do not think it is a core motivation for logical atomism.

This treatment of negation can be generalized in an elegant way by defining a relation of 'negative representation', holding between a sentence and its 'falsifiers', as suggested by Fine's (2017a) 'bilateral' truthmaker semantics.[38] The idea is that a negative sentence ¬p positively represents what p negatively represents (and vice versa). For example, Fb positively represents *Fb* whilst negatively representing $F^{*}b$ and $F^{**}b$, and negating it swaps its positive and negative content.

This approach may seem to place implausibly strong constraints on the world's structure. It requires, for example, that any object lacking a given mass (say, 1kg) instantiate some positive alternative of that mass. But it seems that certain kinds of objects, like space-time points, could simply lack mass altogether, where this does not amount to having some alternative positive property instead. There is, however, a way to accommodate this intuitive possibility within this framework, if we allow for finer type-distinctions than a standard higher-order language does. We find it natural to represent points and particles with terms of different types, and to forbid the mass predicate from syntactically combining with terms of point-type. (Without such type-distinctions,

---

of entities into types. If there is indeed a relation, the question arises as to what its own alternatives are, and a natural view is that it has none (Gómez ms). For simplicity, we include no such relation in our fictional model.

[37] Demos (1917) defended a view along these lines. Russell (1918: 214) complained that there are no propositions (only facts).

[38] This allows us to preserve the equivalence between propositions and their double-negations. Without negative representation, ¬Fb threatens to be synonymous with $F^{*}b \lor F^{**}b$, whose negation is more complex than Fb (Elgin 2023: 618).

a natural model for the approach consists of a single connected spacetime, in which every point has exactly one value of each fundamental field.)

Putting together the above ideas, we can define positive and negative representation for a fragment of the ur-language, consisting of God's atomics and their Boolean combinations. We do this in two stages. First, we define positive and negative representation for atomic sentences (of complexity 0), as follows:

p represents_0^+ *pp* $=_{\text{def}}$ p denotes *pp*'s single member

p represents_0^- *pp* $=_{\text{def}}$ for some *q*, p denotes *q*, and every member of *pp* is an alternative to *q* (and *pp* is non-empty).[39]

Second, we define representation for sentences of complexity n+1, in terms of representation for sentences of complexity n:[40]

p $\text{represents}_{n+1}^+$ *pp* $=_{\text{def}}$

(i) p represents_n^+ *pp*, or

(ii) for some q, p = ¬q and q represents_n^- *pp*, or

(iii) for some q, r, p = q ∧ r and, for some *qq*, *rr*, *pp* = *qq*, *rr*, q represents_n^+ *qq*, and r represents_n^+ *rr*.

p $\text{represents}_{n+1}^-$ *pp* $=_{\text{def}}$

(i) p represents_n^- *pp*, or

(ii) for some q, p = ¬q and q represents_n^+ *pp*, or

[39] This clause implies that an atomic negatively represents any plurality of alternatives (not just the alternatives individually); this is akin to the inclusive treatment of disjunction described above.

[40] Of course, this process never yields a definition of representation for the entire language, though a definition could be given using infinitary disjunction.

(iii) for some q and r, p = q ∧ r and, for some qq_1, qq_2, rr_1, rr_2, q represents_n^- qq_1 and qq_2, r represents_n^- rr_1 and rr_2, and *pp* is either between qq_1 and qq_2, rr_2 or between rr_1 and qq_2, rr_2.

(where *pp* is 'between' *qq* and *rr* when it contains *qq* and is contained in *rr*). Of course, the process of defining notions of representation for increasing levels of complexity never yields a definition of representation for the entire language, though a definition could be given using infinitary disjunction.

The final clause calls for clarification. It entails that any plurality which falsifies either q or r also falsifies q ∧ r. Call these the 'basic falsifiers'. It also entails that the plurality consisting of all the basic falsifiers falsifies q ∧ r (since, by induction, a sentence's falsifiers are closed under union). Call this the 'maximal falsifier'. We can then understand the falsifiers of q ∧ r as any plurality which is between some basic falsifier and the maximal falsifier. To illustrate, take the sentence ¬(Fa ∧ Fb) ∧ Rab. The two basic falsifiers are *Fa*, *Fb* (which falsifies the left conjunct) and R^*ab (which falsifiers the right conjunct). The maximal falsifier is therefore *Fa*, *Fb*, R^*ab. The clause then ensures that ¬(Fa ∧ Fb) ∧ Rab is also falsified by *Fa*, R^*ab and by *Fb*, R^*ab (since these lie between R^*ab and the maximal falsifier).[41]

The addition of these 'intermediate falsifiers' may not be immediately intuitive. But they are needed for the principles of equivalence we like, especially the distribution of conjunction over disjunction, and vice versa. For example, if (Fa ∨ Fb) ∧ Rab is equivalent to (Fa ∧ Rab) ∨ (Fb ∧ Rab), it should be falsified by F^*a, R^*ab. This is an intermediate falsifier: it lies between the basic falsifier R^*ab and the maximal falsifier F^*a, F^*b, R^*ab.

It can be shown that these clauses provide standard truth-conditions, given the definition of truth in terms of representing some obtaining plurality of atomics:

[41] This corresponds to Fine's (2017a) 'convexity' requirement.

p ∧ q is true if and only if p is true and q is true.[42]

¬p is true if and only if p is not true.[43]

For simplicity, we haven't included clauses for disjunction. It behaves just like conjunction, but with positive and negative representation swapped (so the clause for positive representation is like conjunction's clause for negative representation, but with 'represents$_n^-$' swapped for 'represents$_n^+$', and likewise with the clause for negative representation). This ensures that p ∨ q represents in the same way as ¬(¬p ∧ ¬q).

More generally, this semantics ensures that any Boolean combination of atomics represents in the same way as an 'atomist normal form': a disjunction of conjunction of literals (atomics and their negations). A conjunction of literals represents any plurality obtained by 'replacing' each conjunct with some atomic propositions it represents and 'collecting' them. A sentence then represents any plurality represented by a disjunct in some equivalent atomist normal form.

One might wonder whether this approach encodes an arbitrary asymmetry between conjunction and disjunction, by taking representation to reflect disjunctive normal forms. But we don't take our chosen representation relation to be privileged over the parallel relation which reflects conjunctive normal forms. The latter could equally be used to state an adequate semantics, in which the roles of conjunction and disjunction would be reversed.

[42] Proof-sketch: p ∧ q is true (at level n+1) iff for some *pp*, *qq*, p represents$_n^+$ *pp*, q represents$_n^+$ *qq* and *pp*, *qq* obtain iff p is true and q is true (at level n).

[43] Proof-sketch: by induction on complexity. Base case: suppose p is atomic, representing *p*. Then ¬p is true iff p negatively represents some obtaining plurality iff some alternative to *p* obtains iff *p* does not obtain iff p is not true.

Inductive step: Suppose p is at level n but not level n-1. Then it is either (i) a negation or (ii) a conjunction of sentences at level n-1 or below. It can then be shown that the claim holds in both these cases, using the inductive hypothesis together with the relevant clauses.

3.3 Identity and quantification

Intuitively, God had no need for identity when creating the world, any more than conjunction or disjunction (Wittgenstein 1922: §5.53, Burgess 2012: 90). Identity, like the Boolean connectives, feels 'representational'. Suppose, then, that our fictional deity introduces an identity predicate for each type into her ur-language, allowing her to think that a=a, F=F, Fa=Fa, and so on. How can these thoughts be true, when God never distributed any identity relation across the entities She created?

Our favored approach builds on a suggestion of Wittgenstein's (1922: §5.5303): "Roughly speaking: to say of two things that they are identical is nonsense, and to say of one thing that it is identical with itself is to say nothing." Call a sentence 'trivially true' when it positively represents the empty plurality, and does not negatively represent anything. Intuitively, trivial truths require nothing of reality, and there is no way for reality to make them false. Conversely, 'trivially false' sentences do not positively represent anything, and negatively represent the empty plurality. We can incorporate the idea that true identities are trivially true, and false identities are trivially false, by adding clauses like the below (for each type) to our definitions of representation:[44]

> p represents$_{n+1}^{+}$ *pp* if:
> p is of the form a = b, *pp* is the empty plurality, and for some *x*, a denotes *x* and b denotes *x*.
>
> p represents$_{n+1}^{-}$ *pp* if:
> p is of the form a = b, *pp* is the empty plurality, and there is no *x* such that a denotes *x* and b denotes *x*.

(Notice that if a and b co-denote, then the definition of negative representation will entail that a = b does not negatively represent anything, and conversely, if a and b don't co-denote, the definition of positive representation will entail that a = b does not positively represent anything.)[45]

[44] Ramsey (1926) suggests another way of developing this idea.
[45] Denotation will need to extend via synonymity, to allow for truths such as Fa ∧ Fa = Fa, and λx Fx = F.

Triviality introduces a problem. Let ⊥ be a trivial falsehood. The natural extension of our semantics yields the result that Fa ∧ ⊥ is synonymous with ⊥. But this violates a plausible distribution principle, on which (Fa ∧ ⊥) ∨ Fb is synonymous with (Fa ∨ Fb) ∧ (⊥ ∨ Fb). If Fa ∧ ⊥ is synonymous with ⊥, then (Fa ∧ ⊥) ∨ Fb represents only *Fb* whereas (Fa ∨ Fb) ∧ (⊥ ∨ Fb) represents *Fa*, *Fb*.[46] This wrinkle can be ironed out by introducing two further representation relations which keep track of each sentence's positive and negative 'subject-matters'. Intuitively: the positive subject-matter is the plurality of atomics that a sentence is 'about', and the negative subject-matter is the plurality of atomics that its negation is 'about'.[47]

This semantics makes identity statements involving non-denoting expressions, such as ∧ = ∧ and (Fa ∧ Gb) = (Gb ∧ Fa), trivially false. (Inverting Quine: 'no identity without an entity'.) But we would also like to accommodate an identity-like notion of 'equivalence', which guarantees substitutability. For terms of type < >, the idea would be that an equivalence p ≈ q of complexity n+1 $\text{represents}_{n+1}^{+}$ the empty plurality iff p and q represent_n the same pluralities of propositions. This predicts, for example, that Fa ∨ Ga ≈ Fa ∨ Ga is true, whereas Fa ∨ Ga ≈ Fa ∧ Ga is not.[48]

Ideally, we would like to extend equivalence to terms of other types. For example, we would like to recognize a sense in which λx(Fx ∨ Gx) ≈ λx(Gx ∨ Fx) but λx(Fx ∨ Gx) ≉ λx (Fx ∧ Gx). This requires a semantics for logically complex λ-terms, which we won't attempt here. But rather than a denotational semantics, we envisage an 'intra-linguistic' semantics. For example, the meaning of λx(Fx ∨ Gx) might be given by 'reduction' relations to sentences relative to various arguments: e.g., relative to argument a, λx(Fx ∨ Gx) reduces to Fa ∨ Ga, and λx(Gx ∨ Fx) reduces to Ga ∨ Fa.

---

[46] See Fine 2017a: 642 for related discussion. Thanks to Ethan Russo for pointing this out to us.

[47] See Fine (2017b) for the notion of subject-matter. Unlike him, we take a sentence's positive(/negative) subject-matter to be an independent aspect of its content, rather than simply the union of its verifiers(/falsifiers). In particular, a conjunction's subject-matter is always the union of the subject-matters of its conjuncts: for example, Fa ∧ ⊥ has *Fa* as its subject-matter, despite having no verifiers. We then take positive(/negative) representation to be 'closed' with respect to positive(/negative) subject-matter: for example, since (Fa ∧ ⊥) ∨ Fb represents *Fb*, and has *Fa* among its subject-matter, it also represents *Fa*, *Fb*.

[48] This is in the spirit of Elgin 2023 (though he is not working in a logical atomist setting).

The synonymity of lambda-terms might then be understood in terms of that of the sentences they reduce to.

Finally, quantifiers. A familiar idea, going back (at least) to Wittgenstein (1922: §5.3), is that quantification can be understood truth-functionally. Thus, our fictional deity might use $\forall x\, Fx$ to abbreviate the conjunction of its instances, $Fa \wedge Fb$, and $\exists x\, Fx$ to abbreviate $Fa \vee Fb$. This generalizes across types: She might also take $\forall X\, Xa$ to abbreviate $Fa \wedge F^{*}a \wedge F^{**}a$, and $\exists X\, Xab$ to abbreviate $Rab \vee R^{*}ab$.

A worry for this Tractarian approach, going back (at least) to Russell (1918: 234-5), is that $Fa \wedge Fb$ does not seem necessarily equivalent to $\forall x\, Fx$, since it is true relative to a possible situation in which an additional object c that doesn't actually exist is not F. The Tractarian must deny that such situations are possible. In fact, the above semantics for identity forces our fictional deity to deny that any object distinct from a and b could have existed, for $\forall x\, (x{=}a \vee x{=}b)$ is synonymous with the trivially true $(a{=}a \vee a{=}b) \wedge (b{=}a \vee b{=}b)$. It follows that any non-opaque notion of necessity that God adds to her ur-language truly applies to the claim that no object distinct from a and b exists. By parallel reasoning, God could not truly assert the contingency of $\exists x(x{=}a) \wedge \exists x(x{=}b)$. If She does introduce modal talk, She must embrace the controversial doctrine of 'necessitism'.

As we see it, these controversial modal claims follow almost inevitably from *Eliminative Atomism*. Popular solutions among generative atomists don't seem to extend. Take the view that adds 'totality propositions' for each type, attributing a totality relation to all the real entities of that type.[49] Universals could then represent pluralities of instances together with corresponding totality propositions. However, this would not accommodate the possibility of $Fa \wedge Fb \wedge \neg\forall x\, Fx$, in the absence of any atomic propositions involving some further object. Another natural approach in the context of *Generative Atomism* is the idea that generated universal facts merely supervene on the atomic facts, rather than being necessitated by them (e.g. Schaffer 2016: 80). But it is difficult to see what analogous solution there could be in the current context.

[49] In the case of propositions, this would seem to require an infinite hierarchy of totality propositions (cf. Armstrong's 1989: 94 'paradox of totality').

Our view's modal upshots are admittedly counter-intuitive. We take solace in the following observations. First, they have some independent motivation, and we're not alone in endorsing them.[50] Second, contrary intuitions may be partly explained by the truth of contingentism for epistemic (opaque) modalities. Third, we needn't embrace 'spooky non-concreta' like merely possible children (as Williamson does). *Eliminative Atomism* is naturally paired with an austere ontology, such as the view that there are only spacetime points. On this view, it is impossible for something to be Wittgenstein's child, so we needn't posit any such thing as his merely possible child.[51]

To implement the Tractarian approach compositionally, we suggest treating variables as indeterminate names for entities in their domains. For any type σ, we posit a family of denotation relations d_σ. If σ has an empty domain, then d_σ contains a single empty relation. But, if σ is one of the 'real' types other than < > (in our fiction, e, <e> or <e,e>) then d_σ contains many relations, each mapping every constant and variable of syntactic type σ to a corresponding type of entity. We assume that these relations agree on constants but vary on variables like 'assignment functions' do: for each way of mapping variables of type σ to the corresponding domain, there is a corresponding denotation relation. Finally, if d_1, d_2, d_3 are denotation relations for our three real types (e, <e> and <e,e>) then there is a denotation relation $d_{1,2,3}$ for type < > 'based on' d_1, d_2, d_3 which covers all and only the atomic sentences. This relation is defined as you would expect, with the caveat that it does not relate sentences involving variables of non-real types to any atomics.[52]

---

[50] See Ramsey (1927: 170) and Williamson (2013). Williamson's case centrally involves the classical existential introduction rule, which (as eliminative atomists) we must restrict (at least, in non-objectual types). However, our preferred logic vindicates Williamson's reasoning. For example, whereas a familiar kind of opponent denies that (where a is some real term) a=a is a logical truth (or else that 'logical' truths are necessary in the relevant sense), we do not. And, whereas a familiar opponent rejects the inference from Fa ∨ ¬Fa to ∃x Fx ∨ ¬Fx, we do not.

[51] Perhaps there is also a non-perspicuous way of talking about how points are arranged, on which it is true that there are children. We suspect that, when engaged in such talk, it is true that the points could have been arranged so that Wittgenstein had a child, even though the points are actually arranged in such a way that there is nothing that could have been Wittgenstein's child. But it's not our goal here to vindicate this verdict.

[52] That is, if d_1, d_2, d_3 are denotation relations for types e, <e> and <e,e> respectively, then there is a denotation relation for type < >, $d_{1,2,3}$, such that, for any atomic sentence p, property *X*, relation *Y*, and objects *x*, *y*:

Since we now have many denotation relations rather than one, we must replace our definitions of representation with similar denotation-relative definitions. Truth simpliciter is then truth relative to every denotation relation.

To state a semantics, we need two additional pieces of ideology. First, for each variable x, we need an 'x-variant' relation. Two denotation relations for type < >, $d_{f,g,h}$ and $d_{i,j,k}$, are 'x-variants' iff x is of a real type, d_f, d_i agree on all variables except possibly x, and likewise for d_g, d_j and d_h, d_k. Second, we introduce a notion of representation relative to a plurality of denotation functions of type < >. Intuitively, a plurality *qq* is represented by p relative to *dd* when it 'collects' pluralities represented by p relative to each member of *dd.*[53]

Our clauses for the universal quantifier can then be stated as follows:

> $\forall$x p represents$_{n+1}^{+}$ *qq* relative to *d* iff:
> p represents$_n^{+}$ *qq* relative to the x-variants of *d*
>
> $\forall$x p represents$_{n+1}^{-}$ *qq* relative to *d* iff:
> for some pluralities *pp*, *rr* such that p represents$_n^{-}$ *pp* relative to some x-variant of *d*, and p represents$_n^{-}$ *rr* relative to the x-variants of *d*, *qq* is between *pp* and *rr*.

When x belongs to a non-real type, the plurality of x-variants of *d* is empty, and so $\forall$x p is trivially true relative to *d.* The clauses for the existential quantifier are the same, with positive and negative representation reversed.

If p = $\beta^{<e>}(\alpha^e)$, $d_2(\beta^{<e>}, X)$ and $d_1(\alpha^e, x)$, then $d_{1,2,3}$(p, *Xx*), and
if p = $\beta^{<e,e>}(\alpha_1^e, \alpha_2^e)$, $d_3(\beta^{<e,e>}, Y)$, $d_1(\alpha_1^e, x)$ and $d_1(\alpha_2^e, y)$, then $d_{1,2,3}$(p, *Yxy*), and
if p involves a variable of a non-real type then, for all *q*, $\neg d_{1,2,3}$(p, *q*).

[53] More precisely, p represents$_n^{+}$ *qq* relative to *dd* iff, for every member *d* of *dd*, *qq* contains a plurality represented$_n^{+}$ by p relative to *d*, and every *q* in *qq* is in some plurality represented$_n^{+}$ by p relative to some *d* in *dd*. (And likewise for representation$_n^{-}$.)

The above semantics vindicates *Eliminative Atomism*. In our creation story, there are no relations between propositions. So, $\exists X\,(X = \wedge)$ does not denote any atomics, and its negation is trivially true. The same goes for disjunction. And, if we add logicality predicates into the ur-language, the semantics will vindicate the eliminative atomist schema $\neg\exists_\tau X\ \text{Logical}_\tau(X)$. For example, $\exists_{<e,e>}X$ $\text{Logical}_{<e,e>}(X)$ is equivalent to the false disjunction $\text{Logical}_{<e,e>}(R) \vee \text{Logical}_{<e,e>}(R^*)$. And $\exists_{<<>,<>>}X\ \text{Logical}_{<<>,<>>}(X)$ is trivially false, despite the truth of $\text{Logical}_{<<>,<>>}(\wedge)$.

3.4 What has been achieved

We've been sketching a semantic theory for our fictional 'ur-language', which vindicates the intuitive verdict that atomic propositions suffice to provide meanings for its logically complex sentences. Let's take a step back and consider the significance of this vindication.

One way to approach this issue is to ask why we went to all this trouble. It would have been much simpler to give an 'ostrich atomist' semantics, along the following lines:

> If p means that *p* and q means that *q*, then p ∧ q means that *p and q*.
> If p means that *p*, then ¬p means that *not p*.
> etc.

The ostrich atomist may then add that:

> If a sentence S means that *p*, then S is true if and only if *p*.

Nothing we say is incompatible with the ostrich atomist's claims. In fact, we find them very plausible, and would be open to an atomist-friendly semantics that vindicates them. But we don't think that they address the problem of logically complex truths. Our initial puzzlement is, in part, puzzlement about how a complex sentence could mean that *p and q* without bearing a relation to a corresponding conjunctive proposition.

This puzzlement stems from something like the following (schematic) 'reality constraint':

> If p, then p is settled by the distribution of (real) properties/relations over (real) entities.[54]

The puzzlement arises because it is not obvious how the meaningfulness of logically complex sentences could be settled in accordance with this constraint, without meaning relations connecting them to logically complex propositions. Our semantics for the ur-language is intended to address this puzzlement.

The notion of 'settling' might be taken to be some form of necessitation or entailment. But there is a stronger reading of the reality constraint which we also find attractive, which understands 'settling' as metaphysical explanation. Seen through this lens, the problem of logically complex truths is an explanatory demand: how does the distribution of real properties and relations over real entities metaphysically explain the truth and meaningfulness of logically complex sentences?

We haven't been explicitly theorizing in explanatory terms. We earlier rejected a 'generative' conception of the grounding relation between logically complex truths and atomics, but our semantics is naturally regarded as supporting a more 'lightweight' conception of this notion. In particular, let $\leq$ be the relation of 'ultimate grounding', which a truth bears to its fundamental (weak, full, non-factive) grounds. Sentences of the form '$p_1, \ldots, p_n \leq q$', where $p_1, \ldots, p_n$ are atomic sentences in the ur-language and q is a logically complex sentence, can be thought of as trivial truths when the plurality of atomic propositions which $p_1, \ldots, p_n$ collectively represent are represented by q.[55] This recovers a number of natural logical principles concerning ultimate grounding. And it follows that for any true sentence q in the ur-language, there is a corresponding true sentence of the form $p_1, \ldots, p_n \leq q$, where $p_1, \ldots, p_n$ are true atomics (assuming that the ur-

[54] This can be thought of as a higher-order generalization of the principle that 'truth supervenes on being' (Lewis 2001).

[55] Fine (2017b: §6) discusses the connection between truthmaker semantics and ground. Elsewhere, we extend this idea to 'intermediary' grounding connections between the logically complex truths themselves, exploiting their atomist normal forms. The result is a view on which grounding is an object-language 'shadow' of inclusion relations between represented pluralities (in much the way identity is an object language 'shadow' of co-denotation).

language has a term for every entity which it quantifies over).[56] There is a sense, then, in which our approach vindicates a lightweight version of 'grounding atomism'.[57]

So far, we have only described what has been shown about the fictional ur-language. But it might be wondered what relevance this has for the consistency of *Eliminative Atomism* with the claim that we ourselves may be speaking truly when we use logically complex sentences (especially to state and defend our own view). Indeed, there are several notable differences between the ur-language and our own. We've been quantifying plurally (and sometimes over types), which we haven't given a semantics for. And, perhaps more importantly, our language is logically imperfect: its atomic sentences don't correspond to atomic propositions. For example, we've used atomic sentences to describe what linguistic expressions denote. But, in fact, we don't think that there are linguistic expressions or denotation relations in reality, so we don't think there are corresponding atomic propositions. Thus, even if it could be shown that the meanings of our logically complex sentences are settled by these denotation relations, this would not suffice to satisfy the reality constraint. How, then, have we made any progress?

The extension of our approach from the fictional case to our own case relies on a conjecture which we find plausible independently of logical atomism: that all atomic sentences in our language (including those describing denotation) could also be expressed in a logically perfect language with logical resources not dissimilar to the ur-language. If this conjecture is correct, then by making it plausible that ur-language sentences represent pluralities of atomics we thereby make it plausible that our atomic sentences also do so. And, given this, a semantics like the one we've been sketching can be used to derive what complex sentences in our language represent, as a function of the meanings of their logical parts. This would also show that the meaningfulness and truth of

[56] In our fiction, the ur-language only quantifies over the worldly entities in God's creation, and not the representational entities in her mind.

[57] This also vindicates the idea that our semantics can be thought of as a 'metaphysical' semantics, showing how logically complex thought 'fits' into fundamental reality (Sider 2011: §7.4). The comparative flexibility of truth-maker semantics, compared to Sider's preferred biconditional approach, is a crucial advantage in this context (cf. Rubenstein 2024: §5.2). Our approach can also be seen as akin to Cameron's (2010) use of truthmaking to avoid ontological (rather than logical) complexity. It is part of a broader project exploring the advantages of a 'reduction' approach to metaphysical explanation (Rubenstein forthcoming-a, forthcoming-b).

our own logically complex sentences is settled according to the reality constraint, provided that the semantics can be given in a logically perfect meta-language (presumably using logically complex predicates in place of our simple denotation predicates).[58]

There is a caveat, however. In giving the ur-language semantics, we invoked simple denotation relations mapping sentences like Fa and Fx to atomic propositions. But we could have also defined these denotation relations in terms of the denotations of sub-sentential constituents. This doesn't extend to logically imperfect languages. A simple predicate like Grue, for example, cannot denote a corresponding property, because there is no such property. Thus, the meaning of atomic sentences involving Grue cannot be straightforwardly explained in terms of its denotation.

In light of this, our inclination is to take a (somewhat) deflationist attitude toward sub-sentential meaning. We previously mentioned the idea of an 'intra-linguistic' semantics for lambda-terms, which specifies their meanings via the truth-conditions of the sentences they beta-reduce to, when applied to various other terms. A similar idea can be applied to a predicate like Grue: we could give its meaning by describing the pluralities of propositions it represents relative to various possible arguments in the logically imperfect language. (And perhaps we could make sense of the idea that it plays a similar role, in its language, as the role played by some 'corresponding' complex lambda-term in a logically perfect language.)[59]

[58] A complication here is that the proposed approach to quantification involved quantifying over denotation relations themselves. One would hope that this could be implemented with substitutional quantifiers.

[59] If every expression that can occur in argument position itself denotes, then this amounts to the ordinary idea that predicate meanings are functions from entities to propositions. But we don't impose this requirement. Firstly, we want to allow for terms like $\lambda X.X=F$ and $\lambda pq.p \vee q$ which take non-denoting terms as arguments. And secondly, since we like mereological nihilism, we're inclined to think that ordinary names are non-denoting. This raises the question of what kind of meaning these names have. One possibility, consistent with the deflationist approach to sub-sentential meaning, is that their meanings are similarly intra-linguistic.

4. Remaining issues

We've outlined a way of addressing the eliminative atomist's most pressing challenge: the need to give an account of logically complex truths. But of course, many important challenges remain. We finish by discussing a couple of them.

4.1 Moorean arguments

One 'Moorean' concern with *Eliminative Atomism* is that it requires resisting some tempting inferences from obviously true premises. For example:

> Many spiders are poisonous if hairy.
> Therefore, there is a logically complex property that many spiders share.

> George believes that there are mountain lions in the Berkeley hills.
> Therefore, George believes a logically complex proposition.

At first glance, these arguments seem compelling. But, before giving up an elegant and well-motivated metaphysical theory to accommodate their conclusions, we should examine more closely whether their intuitive plausibility remains under the readings of the higher-order quantifiers that are relevant to our thesis.

Regarding the first argument, the inference seems resistible when the quantifier is interpreted in a suitably metaphysical way, as opposed to being loose talk about predicate-meanings, or something similarly representational. Imagine a very simple world, containing just one poisonous hairy spider and one non-poisonous non-hairy spider. In conceiving this scenario, you needn't conceive of the two spiders as instantiating a common property corresponding to the condition 'poisonous if hairy'. It strikes us as perfectly coherent, and even common-sensical, to construe this scenario as one with no such further property.

Common sense does tell in favor of the idea that two hairy spiders share a property. Here we part ways, but only because we would like to combine logical atomism with a reductive physicalist view, on which talk of hairy spiders is equivalent to logically complex talk about underlying micro-physical facts. We suspect that the common-sense verdict is rooted in the naïve assumption that our 'hairy spider' talk is reasonably faithful to reality's structure. Reductive physicalism provides reason to reject this assumption, quite independently of logical atomism.

As for the second argument, there is independent philosophical and scientific motivation for the view that a subject's belief relations to the world are mediated by representational entities, which that-clauses can carry non-truth-conditional information about.[60] This makes it natural to interpret the premise as communicating that George believes* some mental sentence which resembles 'There are mountain lions in the Berkeley hills' in structure and meaning. This entails that George believes* a meaningful logically complex sentence (which is perhaps all the conclusion says in certain contexts). One cannot securely infer from this that George believes a logically complex proposition, without assuming that meaningfulness requires expressing a single proposition. We take the viability of the approach described in §3 to cast doubt on this further assumption.[61]

The dialectic here closely mirrors that regarding mereological nihilism. Some non-nihilists argue from common sense for the view that the particles arranged table-wise compose a table. But we agree with Sider (2013a: §§2-4) and Dorr (2002: §§1.4.3-4) that there are good reasons to resist this kind of argument. For one thing, we think good metaphysical arguments can cast doubt on common-sense views. But, more importantly, the nihilist's departure from common sense is not as big as it might initially appear. The claim that the particles arranged table-wise compose a table, as uttered in the metaphysics room, is naturally read as carrying commitment to more than the claim that there are particles arranged table-wise. It demands a composite object and a composition relation as further additions to reality. It's not obvious to us that this is what common sense

[60] See Nelson 2024 for an overview of the vast literature on this idea.

[61] The conclusion may instead be supported by its explanatory role in psychology. We are optimistic, however, that suitable atomist paraphrases can be given for content-involving explanations, in terms of mental sentences' relations to atomics (conveniently albeit non-perspicuously conveyed via their resemblance to English sentences). Vindicating this optimism is a project for another time.

demands; there may be more deflationist ways of talking about 'composition', on which talk of composites amounts to talk of their simple parts.

Similarly, whilst it wouldn't be devastating for *Eliminative Atomism* to be at odds with common sense, it's not clear to us that it is. There may be 'deflationist' readings of the above arguments' conclusions that are consistent with logical atomism, and are also reasonable construals of the common-sense stance on logical complexity.

4.2 Indispensability

There is a more theoretical way of arguing against the version of logical atomism we have been developing. Philosophers who are friendly to the notion of structure/fundamentality that our view presupposes tend to endorse an abductive methodology for discovering which notions are structural, or what reality is fundamentally like. The idea is that the best explanations of what we know (or, more realistically, our best guesses as to what those explanations might look like) should guide our views of structure. Since the best explanations of our evidence are surely logically complex, this abductive approach seems to support logical realism. Indeed, this is a central motivation for both Sider (2011) and McSweeney (2019b).

One could try to reply to this argument by denying that the best explanations of our evidence are logically complex. Given any candidate explanation of q in terms of p, we can envisage constructing an 'atomist analogue' explanation by replacing p with its atomic grounds pp. We are inclined to think that this atomist analogue will be true, and that it will be a partial ground of the logically complex explanation.[62] But this would not suffice to respond to the abductive argument, since it is implausible that this atomist explanation would be 'better' in the relevant epistemic sense. After all, the new explanation might replace elegant generalizations with long (possibly infinite) lists of their instances. Moreover, it may end up being overly detailed, and therefore non-proportional, since it lacks abstraction devices like disjunction and existential quantification. Thus,

[62] For further discussion of atomist approaches to explanation, see Gómez (ms).

while we hold that all logically complex talk (including logically complex explanations) is made true by atomics, we do not add that the epistemically best explanations are always logically simple.

Our preferred response rests on a distinction between two ways of understanding the abductive methodology. On a 'syntactic' construal, we should take the syntactic structure of our best explanations to reflect reality's structure. Thus, if our best explanations are logically complex, we should take reality itself to be logically complex. We think there are reasons to be suspicious of this syntactic principle, even independently of logical atomism. For, under a natural reading, it invites us to take seriously the idea that conjuncts enter into conjunctions in a certain order, that two objects can occupy the same 'position' in two different relations, and the like.

We prefer a weaker abductive principle, which recommends commitment only to the constituents of our best explanations' 'worldly contents'. Since we deny that logically complex explanations have logically complex contents, we can happily accept this abductive principle and use it to guide us in discovering the simple fundamental entities. Now, because we cannot settle what a theory's semantic constituents are independently of our views of structure, this epistemology of structure is best thought of as yielding a 'coherence' constraint (when coupled with constraints on the semantic theories that our best explanations can plausibly be given).[63] Our point is just that — in light of the approach in §3 — there is a stable package which combines an atomist metaphysics with a natural construal of the abductive methodology.

A serious indispensability challenge remains. One central application of higher-order logic is in mathematics: for example, quantification over logically complex relations is used to define notions like equinumerosity, and the ancestral of a relation. It is unclear whether atomist-friendly analogues of these definitions could be given.[64] This invites a fictionalist approach to these permissive forms of higher-order quantification, which we hope to pursue in future work.[65]

---

[63] Sider's own approach is similarly holistic: he notes (2011: 222) that we should not take our best theory's font or ink color as a guide to reality's structure, because we have independent reason to think that these features of the theory do not carve at the joints.

[64] Thanks to Alex Roberts for articulating this challenge.

[65] Thanks to Shamik Dasgupta, Geoff Lee, Bar Luzon, Alex Roberts, and Ted Sider for feedback on previous drafts. We have also benefitted from discussion with David Builes, Sam Elgin, Ethan Russo, and

References

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

——— (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.

——— (2004). *Truth and Truth-makers*. Cambridge: Cambridge University Press.

Ayer, Alfred Jules (1936). *Language, truth and logic*. London: V. Gollancz.

Bacon, Andrew (2020). Logical Combinatorialism. *Philosophical Review* 129 (4):537-589.

——— (2023). A Theory of Structured Propositions. *Philosophical Review* 132 (2):173-238.

Bennett, Karen (2011). By Our Bootstraps. *Philosophical Perspectives* 25 (1):27-41.

——— (2017). *Making Things Up*. Oxford: Oxford University Press.

Burgess, Alexis (2012). A Puzzle about Identity. *Thought: A Journal of Philosophy* 1 (2):90-99.

Cameron, Ross P. (2010). How to have a radically minimal ontology. *Philosophical Studies* 151(2):249–264.

Carnap, Rudolf (1928). *Der Logische Aufbau der Welt*. Hamburg: Meiner Verlag.

Correia, Fabrice (2010). Grounding and truth-functions. *Logique Et Analyse* 53(211):251–279.

Dasgupta, Shamik (2014). The possibility of physicalism. *Journal of Philosophy, 111,* 557–592.

Demos, Raphael (1917). A discussion of a certain type of negative proposition. *Mind* 26 (102):188-196.

Donaldson, Thomas (2015). Reading the Book of the World. *Philosophical Studies* 172 (4):1051-1077.

Dorr, Cian (2002). *The Simplicity of Everything*. PhD thesis, Princeton University.

——— (2004). Non-symmetric Relations. *Oxford Studies in Metaphysics* 1: 155–92.

——— (2005). What we disagree about when we disagree about ontology. In Mark Eli Kalderon, *Fictionalism in Metaphysics*. New York: Oxford University Press. pp. 234–86.

——— (2016). To Be F Is To Be G. *Philosophical Perspectives* 30(1):39–134.

Dorr, Cian & Hawthorne, John (2013). Naturalness. *Oxford Studies in Metaphysics* 8.

Eddon, Maya (2013). Fundamental Properties of Fundamental Properties. *Oxford Studies in Metaphysics* 8.

Eklund, Matti (2024). *Alien Structure: Language and Reality*. Oxford: Oxford University Press.

participants in the Princeton Metaphysics Workshop, the Higher-Order Metaphysics Workshop at NYU, the Rutgers Metaphysics Reading Group, and the Berkeley Metaphysics and Science Reading Group.

Elgin, Samuel Z. (2023). The Semantic Foundations of Philosophical Analysis. *Review of Symbolic Logic* 16 (2):603-623.

Fine, Kit (2000). Neutral relations. *Philosophical Review* 109 (1):1-33.

——— (2010). Some Puzzles of Ground. *Notre Dame Journal of Formal Logic* 51 (1):97-118.

——— (2012). Guide to Ground. In F. Correia & B. Schnieder (eds.) *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.

——— (2017a). A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic* 46 (6):625-674.

——— (2017b). A Theory of Truthmaker Content II: Subject-matter, Common Content, Remainder and Ground. *Journal of Philosophical Logic* 46 (6):675-702.

Fritz, Peter (2021). Ground and Grain. *Philosophy and Phenomenological Research* 105 (2):299-330.

Fritz, Peter & Bacon, Andrew (ms). The Algebra of Logical Atomism.

Fritz, Peter & Jones, Nicholas K. (eds.) (2024). *Higher-Order Metaphysics*. Oxford University Press.

Gilmore, Cody (2013). Slots in Universals. *Oxford Studies in Metaphysics* 8.

Gómez, Verónica (ms). Nomic Particularism.

Goodman, Jeremy (2023). Grounding Generalizations. *Journal of Philosophical Logic* 52 (3):821-858.

Hirsch, Eli (2011). *Quantifier Variance and Realism: Essays in Metaontology*. Oxford: Oxford University Press.

Jackson, Alexander (2024). The fundamental facts can be logically simple. *Noûs* 58(4):997-1016.

Jago, Mark & Barker, Stephen (2012). Being Positive About Negative Facts. *Philosophy and Phenomenological Research* 85 (1):117-138.

Lewis, David K. (1983). New work for a theory of universals. *Australasian Journal of Philosophy* 61 (4):343-377.

——— (1986). *Philosophical Papers, Volume II*. Oxford: Oxford University Press.

——— (2001). Truth Making and Difference Making. *Noûs* 35: 602–15.

Litland, Jon Erling (2017). Grounding Grounding. *Oxford Studies in Metaphysics* 10.

McSweeney, Michaela Markham (2019a). Following logical realism where it leads. *Philosophical Studies* 176 (1):117–139.

——— (2019b). Logical Realism and the Metaphysics of Logic. *Philosophy Compass* 14 (1):e12563.

——— (2020). Debunking Logical Ground: Distinguishing Metaphysics from Semantics. *Journal of the American Philosophical Association* 6 (2):156-170.

Molnar, G. (2000). Truthmakers for negative truths. *Australasian Journal of Philosophy* 78 (1):72–86.

Nelson, Michael (2024). Propositional Attitude Reports. *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/fall2024/entries/prop-attitude-reports/>.

Plate, Jan (2016). Logically Simple Properties and Relations. *Philosophers' Imprint* 16:1-40.

Prior, Arthur Norman (1971). *Objects of thought*. Edited by P. T. Geach & Anthony Kenny. Oxford: Clarendon Press.

Ramsey, Frank P. (1926). The foundations of mathematics. *Proceedings of the London Mathematical Society* 25:338–384.

——— (1927). Facts and Propositions. *Aristotelian Society Supplementary Volume* 7 (1):153-170.

Raven, Michael J. (2020). Is Logic Out of This World? *Journal of Philosophy* 117 (10): 557-577.

Rosen, Gideon (2010). Metaphysical Dependence: Grounding and Reduction. In Bob Hale & Aviv Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.

Rubenstein, Ezra (2024). Two approaches to metaphysical explanation. *Noûs* 58 (4):1107-1136.

——— (forthcoming-a). Against Grounding Physicalism. *Journal of Philosophy*.

——— (forthcoming-b). Generalism Without Generation. *Journal of Philosophy*.

Russell, Bertrand (1911). Le Réalisme analytique. Translated as 'Analytic Realism,' in *Russell on Metaphysics*, ed. S. Mumford. London: Routledge, 2003. pp. 91–96.

——— (1918). The Philosophy of Logical Atomism. In *Logic and Knowledge*, ed. R.C. Marsh. London: Allen & Unwin, 1956. pp. 177–281.

——— (1924). Logical Atomism. In *Logic and Knowledge*, ed. R.C. Marsh. London: Allen & Unwin, 1956. pp. 323–43.

Russell, Jeffrey Sanford & Hawthorne, John (2018). Possible Patterns. *Oxford Studies in Metaphysics* 11.

Schaffer, Jonathan (2015). What Not to Multiply Without Necessity. *Australasian Journal of Philosophy* 93(4):644–664.

——— (2016). Grounding in the image of causation. *Philosophical Studies* 173(1):49–100.

——— (2017). Laws for Metaphysical Explanation. *Philosophical Issues* 27 (1):302-321.

Sider, Theodore (2011). *Writing the Book of the World*. Oxford: Oxford University Press.

——— (2013). Against Parthood. *Oxford Studies in Metaphysics* 8:237–293.

Skow, Bradford (2016). *Reasons Why*. Oxford: Oxford University Press.

Turner, Jason (2016). Curbing Enthusiasm About Grounding. *Philosophical Perspectives* 30 (1): 366-396.

Van Fraassen, B. (1969). Facts and tautological entailments. *Journal of Philosophy* 66(15): 477–87.

Wang, Jennifer (2013). From Combinatorialism to Primitivism. *Australasian Journal of Philosophy* 91(3): 535–554.

Williamson, Timothy (2003). Everything. *Philosophical Perspectives* 17 (1):415–465.

——— (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Wilson, Jessica (2010). What is Hume's Dictum, and Why Believe It? *Philosophy and Phenomenological Research* 80(3): 595–637.

Wittgenstein, Ludwig (1922). *Tractatus Logico-Philosophicus*. London: Kegan Paul.