

TWO APPROACHES TO METAPHYSICAL EXPLANATION

Ezra Rubenstein

University of California, Berkeley

Explanatory metaphysics aspires to explain the less fundamental in terms of the more fundamental. But we should recognize two importantly different approaches to this task. According to the generation approach, more basic features of reality generate (or give rise to) less basic features. According to the reduction approach, less perspicuous ways of representing reality reduce to (or collapse into) more perspicuous ways of representing reality. The main goals of this paper are to present the core differences between the two approaches (§2), to demonstrate the distinction's significance (§3), to provide some resources for adjudicating between the approaches (§4), and to argue that the project of explanatory metaphysics needs both (§5).

1. Introduction

We come to metaphysics with a wide-ranging and rich conception of the world. Metaphysicians seek to impose some order on this conception (perhaps after some revision). They ask: how can we make sense of the world thus conceived? How do its various aspects 'hang together'? This can be seen as demanding a distinctively metaphysical kind of *explanation*: metaphysicians want to know how various aspects of reality can be metaphysically explained in terms of others.

This demand presupposes that the world has a certain structure: it is not just a disconnected mess of brute details. Two observations seem to justify this assumption. Firstly, reality seems to be certain ways in virtue of being other ways; we capture this general idea by saying that some aspects of reality *ground* others.¹ Secondly, certain aspects of reality seem superficial, derivative, or

¹ *N.B.* I use 'ground' throughout as a generic term for the explanatory connection between more and less fundamental, as opposed to any specific conception of this connection. It can be taken to signify

mundane, whereas others seem deep, basic, or ultimate; we capture this general idea by saying that some aspects of reality are more *fundamental* than others. And these two features appear intimately connected: more fundamental aspects seem to ground less fundamental aspects.

In this way, the world appears to support metaphysical explanation: we can explain its less fundamental aspects in terms of the more fundamental aspects which ground them. And we might hope, or even expect, that the world supports an especially satisfactory form of metaphysical explanation, by providing a foundation of absolutely fundamental aspects which both ground all others and are themselves ungrounded. These, at a highly abstract level, are the roles of grounding and fundamentality in explanatory metaphysics (in the broadest sense of these terms).²

As I see it, contemporary metaphysics exhibits a crucial but neglected divide between two overarching conceptions of what plays these roles, and correspondingly two substantively different approaches to metaphysical explanation. On the ‘generation’ approach, the fundamental generates, or gives rise to, less basic features of reality. This generation is analogous in many ways to causation; in something like the way causation connects earlier to later, generation connects more basic to less basic. On the ‘reduction’ approach, by contrast, explanatory metaphysics seeks to reinterpret our conception of reality in perspicuous, or less metaphysically distorted, terms. There is merely a representational (rather than worldly) difference between fundamental and non-fundamental: non-perspicuous ways of representing reality reduce to, or collapse into, others.

either metaphysical explanation itself (as per ‘unionism’: e.g. Fine 2001, Dasgupta 2017), or a relation backing explanation (as per ‘separatism’: e.g. Audi 2012, Schaffer 2012).

² Although I will speak as if all metaphysical explanation revolves around fundamentality and ground, I recognize that there may be other kinds of metaphysical explanation (especially revolving around *essence* — see Fine 2015, Glazier 2017, Correia & Skiles 2017). But it bears emphasis that I am employing *broad* notions of fundamentality and ground: some kinds of metaphysical explanation may revolve around these notions whilst being disconnected from narrower versions of them. (A similar remark applies to those who would disconnect metaphysical explanation from these notions altogether e.g. Baron & Norton 2019.)

The divide between these approaches is rarely explicitly articulated and discussed.³ But it casts a shadow over ‘first-order’ debates: many disputes within both metaphysics and neighboring disciplines look quite different depending on which approach is taken. In this paper, I aim to articulate and clarify the core differences between generation and reduction (§2), to demonstrate the distinction’s wide-ranging significance (§3), to suggest how we might adjudicate between the two approaches (§4), and to argue that explanatory metaphysics needs both (§5).

2. The Two Approaches

Let me introduce the distinction between generation and reduction by means of some familiar examples of metaphysical explanation.

First, composition. According to a natural view, a composite object’s existence and nature are metaphysically explained by the existences and natures of its parts. For example, a table exists and has a certain mass at least partially because the atoms which make it up exist and have certain masses. On the generation approach to this explanation, the existences and masses of certain atoms generate – or ‘metaphysically cause’ – the table’s existence and mass. The table-portion of reality and the atomic portion are distinct but intimately connected: one gives rise to the other.⁴ On the reduction approach to this explanation, there is only one portion of reality, and two ways of representing it. Our talk of the table’s existence and mass is true, but the state of affairs it represents would in some sense be better represented in terms of the existence and masses of certain atoms. Only the atoms are ‘really there’, but they support our second-rate talk of tables (unlike talk of unicorns, say).⁵

Second, mentality. On a standard understanding of physicalism, the mental is metaphysically explained by the physical. For example, the legend goes that we feel pain in virtue of our C-fibers firing. On the generation approach to this explanation, my C-fibers firing (a physical event)

³ Williams (2012) and deRosset (2017) are notable exceptions.

⁴ E.g. Russell 2003: 92, Fine 2010: §IX.

⁵ E.g. Cameron 2010, Sider 2013a.

generates my pain (a mental event): the events are distinct but intimately connected. (This is a natural interpretation of ‘non-reductive physicalism’.)⁶ On the reduction approach, the pain *just is* the C-fibers firing: our talk of pain is a distinctive way of latching onto what is ‘really’ just some neural activity. (This is a natural interpretation of ‘reductive physicalism’.)⁷

Third, determinables. An object’s determinable properties are naturally explained in terms of its more specific determinate properties: for example, the apple is red because it is scarlet. But again, there are two quite different ways of understanding such explanations. On the generation approach, determinable properties are just as real as their determinates, with the facts involving the former generated or determined by the facts involving the latter.⁸ Alternatively, on the reduction approach, objects only really instantiate maximally specific properties, with determinables being mere ‘shadows’ of linguistic abstraction. On this view, there is a single state of affairs with respect to the apple, and more or less abstract ways of describing it.⁹

These examples suggest that the distinction I am concerned with is very general: many proposed metaphysical explanations can naturally be conceived in two different ways. The rest of this section aims to present these two approaches as clearly as possible. Since my presentation aims to be neutral and abstract, particular examples, general principles, and the choice of primitives are all contentious. Nonetheless, we can ‘triangulate’ onto the distinction by considering a range of reasonably natural cases and connections to surrounding notions. What is important is that a clear conception is conveyed, which readers may then adapt to their own preferred frameworks. I hope that the reader’s grasp of the distinction will strengthen throughout the discussion, as further cases and connections arise.

⁶ E.g. Shoemaker 2007, Schaffer 2021.

⁷ E.g. Loar 1990, Papineau 2002.

⁸ E.g. Wilson 2009, Rosen 2010: §11.

⁹ E.g. Armstrong 1997: 53, Heil 2003: 207.

2.1 Facts vs truths

The two approaches share the idea that certain ‘aspects’ of reality are more or less fundamental and stand in relations of ground to one another. But adherents of the generation approach (‘generators’) take these aspects to be worldly, akin to objects or events, whereas adherents of the reduction approach (‘reducers’) take them to be representational, akin to names or sentences. I will speak in terms of facts on the generation side and truths on the reduction side.¹⁰

Truths, as I use the term, are just true truth-bearers: perhaps sentence-tokens or sentence-meanings/propositions. (Since ‘proposition’ is sometimes used in a worldly sense, I will reserve it as a term which is neutral between facts and truths.) ‘Fact’ is sometimes used as synonymous with my term ‘truth’. However, my use of ‘fact’ follows a tradition which recognizes a more inflationary (or ‘thick’) notion, on the side of reality rather than representation (Russell 1918, Wittgenstein 1922, Armstrong 1997).¹¹ On this conception, facts are entities which constitute the world — often identified with obtaining states of affairs — rather than true ways of representing the world.

Truths stand to facts as sense stands to reference: they are guises by which facts are represented.¹² Just as different names can refer to the same object, and different predicates can denote the same property, so different truths can represent the same fact (Correia 2010: §2, Audi 2012: §IV.1, Rayo 2013: §1.3). For example, the following pairs of sentences plausibly correspond to different truths but the same facts:

‘Hesperus is rising’, ‘Phosphorus is rising’;

‘Snow is white’, ‘It is not not the case that snow is white’;

¹⁰ Although I set this aside, both generator and reducer may also target ‘sub-propositional’ aspects (e.g. Schaffer 2009 on entity-grounding, Fine 2015 on generic grounding).

¹¹ Cf. Correia (2010: 256) on “worldly facts” vs. “conceptual facts” and Jago (2011: 42) on facts vs. propositions.

¹² Following Fine (2013), we might distinguish two senses of ‘representing’: a truth may *describe* a fact, or it may *express* a fact, where a disjunction represents the same fact as its true disjunct in the first sense but not the second. I largely ignore this issue here, but it becomes relevant in §5.2.

‘London is north of Paris’, ‘Paris is south of London’.

In each case, the difference between the sentences seems merely representational, rather than worldly: they are anchored to the same portion of reality, but tied by different ropes.

Of course, there are many ways of carving up both facts and truths, and I cannot explore the issues in any depth here. The crucial point is just that the criteria for individuating them differ, with truths individuated more finely: if p and q correspond to the same truth, they represent the same fact, but perhaps not conversely. In particular, certain paradigmatically ‘opaque’ contexts — which, intuitively, are sensitive to merely representational features — may distinguish truths, but not facts. For example, it might be held that ‘Hesperus is Hesperus’ is analytic, a priori, and believed by the Babylonians, whereas ‘Hesperus is Phosphorus’ is synthetic, a posteriori, and disbelieved by the Babylonians. This indicates that these sentences correspond to distinct truths, but not that they represent distinct facts.

For the sake of regimentation, let me introduce some official primitives.¹³ I assume that for every true sentence p , there is *the truth that* p , which I denote by ‘ $\langle p \rangle$ ’, and which *represents* some fact.¹⁴ I will also use the expression ‘*the fact that* p ’, which I abbreviate by ‘ $[p]$ ’. On one way of using this expression, ‘ $[p]$ ’ denotes whichever fact $\langle p \rangle$ represents. However, following somewhat standard terminology (e.g. Rosen 2010: 115), I adopt a more restricted use, where ‘ $[p]$ ’ only successfully refers if reality provides a matching fact. I explicate this usage below, connecting it to perspicuity.

¹³ In taking these notions as primitive, I simply mean that I will not be defining them; I do not mean that they cannot be defined, or that they are primitive in any metaphysical sense.

¹⁴ This is idealized in several ways. For example, since truths correspond to sentences in contexts, ‘ $\langle p \rangle$ ’ should be understood as a context-dependent expression, and since ordinary sentences are almost always imprecise, their corresponding truths do not determinately represent any particular fact. Also, there may be ‘non-factual’ truths which do not represent any fact e.g., \langle Torture is wrong \rangle , on an expressivist view. (Conversely, truth may be stronger than required for the reducer’s representational aspects; some weaker positive status such as ‘aptness’ may suffice.)

I am speaking as if ground is a relation, with the two approaches conceiving its relata differently. This makes for a vivid picture, but it should be remembered that ground needn't be conceived this way (on either approach). An alternative conception expresses grounding claims using a sentential operator, with no commitment to any entities corresponding to the expressions that flank it (Correia & Schnieder 2012: 10; Fine 2012: 47; Dasgupta 2017: §2). Accordingly, an alternative way of presenting the two approaches eliminates my talk of facts and truths in favor of sentential operators (such as *...is factually equivalent to...*). The difference between my 'reifying' presentation and this non-reifying alternative will not matter for the discussion which follows.

2.2 Generation vs. reduction

Generators take ground to be a matter of worldly aspects *generating* (giving rise to, or building) others. Introduction by paradigm cases is fraught in this context, since it is contested not only whether a given case involves ground at all, but also, assuming that it does, which notion of ground (generation or reduction) it is. Nonetheless, the grounding of sets by their members, and of wholes by their parts, provide intuitive examples. [Socrates exists] is naturally taken to generate [{Socrates} exists], and the fact that the hydrogen atoms are bonded is naturally taken to generate the fact that the H₂ molecule exists. In each case, the facts in question seem clearly distinct: in the first, one involves a person and the other a set; in the second, one involves two atoms and the other one molecule. Nonetheless, the facts seem intimately connected, with one both metaphysically necessitating and wholly relevant to the other. Thus, these cases exhibit the characteristic features of generation: they involve a tight explanatory connection between distinct facts.

The model here is causation — as Jonathan Schaffer (2012: 122) puts it: 'Grounding is something like metaphysical causation. Roughly speaking, just as causation links the world across time, grounding links the world across levels.' Developing this analogy, Schaffer (2016) observes that grounding and causation are similarly connected to laws, necessity, counterfactuals, and explanation.¹⁵ In a similar vein, Karen Bennett (2011a, 2017) characterizes a family of 'building' relations, operating variously on objects, properties, and states of affairs, with composition and constitution as paradigm examples. Building relations are unified by their licensing 'generative'

¹⁵ See also A. Wilson (2018).

talk: ‘builders *generate* the built’ (2017: 58).¹⁶ And Gideon Rosen glosses grounding as ‘metaphysical dependence’ (2010), ‘a much more intimate form of dependence’ than causation (2011: 123), in which facts ‘*give rise to* or *generate*’ others (2011: 130).¹⁷

Generators standardly understand fundamentality in terms of generation: generation orders the world into a hierarchy — ‘the great chain of being’ (Schaffer 2009: 376) — and fundamentality is a matter of position within this hierarchy.¹⁸ If one aspect of reality generates another, it is thereby more fundamental than it, and the absolutely fundamental aspects of reality are those that are ungenerated (Schaffer 2009: 351; Rosen 2011: 124; Bennett 2017: ch.5).¹⁹ I will use ‘basicness’ to denote this conception of fundamentality; thus, I will speak of more basic facts generating less basic facts.

By contrast, reducers take ground to be a matter of truths reducing to (consisting in, or collapsing into) others. With the same caveat about paradigm cases as above, intuitive examples are the reduction of truths about water to truths about H₂O and of truths about heat to truths about molecular motion. It is natural to hold that the truths in question represent the same facts, where this equivalence is directed: the water in the glass ‘is really just’ a collection of H₂O molecules,

¹⁶ Bennett reserves ‘grounding’ for a particular building relation, but notes a more generic sense which may be synonymous with ‘building’ (2011a: n.10; 2017: 12).

¹⁷ Rosen’s credentials as a generator might seem undermined by his ‘grounding-reduction link’ (2010: §10), according to which, if $\langle p \rangle$ ‘reduces’ to $\langle q \rangle$ — in that $\langle p \rangle$ ’s being the case ‘consists in’ $\langle q \rangle$ ’s being the case — then $[p]$ is grounded in $[q]$. However, since Rosen posits a worldly distinction between the facts in question, reduction in his sense is not reduction in my sense: it is not a relation among sentences ‘but rather among the facts ... those sentences purport to describe’ (2011: 122).

¹⁸ One might view Kit Fine as a generator who does not understand fundamentality in terms of generation: in addition to his primitive notion of ground, he posits a primitive, absolute notion of ‘fundamental reality’, denying that it has ‘a relational underpinning’ (2001: 25). However, as discussed below, I see Fine’s notion of fundamentality as paradigmatic of the reducer’s notion of perspicuity, and hence prefer to interpret his rich framework as combining the two approaches. Indeed, he (2001: 15) explicitly distinguishes statements of *ground*, which sometimes connect propositions which ‘hold in reality’, from ‘stricter’ statements of *reduction*, which imply ‘the unreality of what is reduced’.

¹⁹ Analyzing relative fundamentality in terms of generation is non-trivial: see Bennett 2017: §6.5.

and the air getting hotter ‘is really just’ air molecules moving faster. This directedness seems to support explanation: truths about water obtain because truths about H₂O obtain, and likewise truths about heat obtain because truths about molecular motion obtain.

Reducing truths seem to ‘match’ facts which the reduced truths seem to ‘distort’ (and it is this that provides explanatory asymmetry). When discussing the examples of metaphysical explanation above, I spoke of what may be said truly in a ‘second-rate’ language, or what is merely a ‘shadow’ of language, being explained in terms of how things ‘really’ are, or what holds ‘in reality itself’. My preferred regimentation of this way of speaking uses a predicate of truths: ‘<p> is *perspicuous*’ means that <p> is not only true but reflects reality’s structure in a deeper sense.²⁰ I propose to understand reduction in terms of perspicuity as follows. ‘<p> reduces <q>’ means that:

- i) <p> and <q> represent the same fact, and
- ii) <p> is perspicuous and <q> is not.²¹

For example, ‘<the mean molecular energy of the air in the room is x> reduces <the room temperature is y>’ means that: i) these truths represent the same fact, and ii) <the mean molecular energy of the air in the room is x> is perspicuous and <the room temperature is y> is not.

We can get a grip on the key notion of perspicuity by considering the familiar example of ‘the average family has 2.2 children’. Granting the statistic’s accuracy, this sentence is true. But, in another sense, it clearly fails to ‘match’ the fact it represents: there is no family bearing the having

²⁰ Cf. (O’Leary-)Hawthorne & Cortens 1995: §3, Fine 2001: 3, Turner 2010: 8–9, Sider 2013a: 252.

²¹ In the more general many-one case, ‘<p_{1n}> reduces <q>’ means that:

- i) <q> represents the same fact(s) as the conjunctive truth <p₁ &...& p_n>, and
- ii) <p_{1n}> are each perspicuous and <q> is not.

(Note that this may require a many-one relation of representation between a single truth on the one hand and many facts on the other.) This yields a notion of partial reduction, analogous to partial generation: <p> partially reduces <q> iff there are some <p_{1n}> such that <p>, <p_{1n}> reduces <q>.

relation to ‘each’ of 2.2 children. A sentence which appears to better approximate the corresponding fact is: ‘the total number of children divided by the total number of families is 2.2’. (Of course, this sentence may itself be non-perspicuous, e.g. if numbers don’t exist.)²²

As I see it, the notion of perspicuity appears in many guises throughout analytic metaphysics, and there are a number of routes to understanding it.²³ Following Kit Fine (2001: 25), one might view perspicuity as a matter of a truth’s matching the structure of the corresponding fact.²⁴ On this conception, $\langle p \rangle$ is perspicuous just in case for each representational constituent of $\langle p \rangle$, $[p]$ has a corresponding worldly constituent, and for each structuring relation between the constituents of $\langle p \rangle$, the constituents of $[p]$ are correspondingly related.

For example, a simple atomic truth of the form ‘a is P’ is perspicuous if and only if there is some object a and property P, corresponding to the name ‘a’ and predicate ‘P’ respectively, such that a instantiates P. Suppose there are facts involving determinate properties of the form *has x kg mass* but not involving the determinable property *has mass*. Then $\langle \text{Joe has mass} \rangle$ would not be perspicuous since there is no correspondingly structured fact $[\text{Joe has mass}]$, only a fact of the

²² This example may be superficial: the structure of the corresponding sentence-meaning may not match that of the sentence (Kennedy & Stanley 2009). Nonetheless, the mismatch between the sentence itself and the fact represented vividly illustrates non-perspicuity, and it will be helpful in what follows to proceed under the simplifying assumption that sentences’ structures by and large resemble that of their corresponding truths.

²³ Historically, it is central to the logical atomists’ notion of analysis and Quine’s notion of paraphrase. More recently, see van Inwagen (1990) on ‘real’ vs. ‘virtual’ objects, Field (1994) on ‘factually defective discourse’, Dorr (2007) on ‘fundamental’ vs. ‘superficial’ uses of sentences, Cameron (2010) on ‘Ontologese’, Williams (2012) on ‘requirements on reality’, Russell (2013) on ‘objective matters of fact’, deRosset (2017) on conciliatory irrealist views, and Bacon (2020) on fundamentality. I cannot undertake any detailed exegesis here, but — despite important differences between their frameworks — I understand all of these authors as invoking versions of perspicuity.

²⁴ As Fine (2001: 3) puts it, truths which fail to ‘hold in reality’ do ‘not *perspicuously* represent the facts—there will be some divergence between how the facts are “in themselves” and how they are represented as being’. Cf. Bennett (2001: 147) on ‘truths that directly reflect the metaphysical situation’.

form [Joe has x kg mass]. Obviously this requires a structured conception of facts as built up from worldly constituents, such as objects and properties, together with a suitably ‘inflationary’ understanding of these constituents and their relations (Rosen 2010: 114; Audi 2012: 686).²⁵

Alternatively (though not incompatibly), one might understand perspicuity in terms of a sub-propositional notion. Following Ted Sider (2011: 116), one might require that the constituents of a perspicuous truth are each ‘structural’. This notion extends Lewisian naturalness ‘beyond the predicate’, to names, sentential operators, quantifiers, *etc.* Structural notions ‘carve nature at its joints’, in ways that matter for objective similarity, explanation, and confirmation, and which confer a kind of epistemic value that is independent of truth. For example, ‘is grue’ and ‘is taller than’ do not denote perfectly natural properties/relations; hence <this emerald is grue> and <Trump is taller than Obama> are not perspicuous truths. On the other hand, perhaps ‘is negatively charged’ does denote a perfectly natural property; hence <Sparky is negatively charged> may be perspicuous (insofar as the name ‘Sparky’ is itself structural!)²⁶

I have been employing a binary notion of perspicuity. However, reduction is naturally viewed as a gradual process of analyzing or unpacking a non-perspicuous truth, with the underlying fact better approximated as various layers of distortion are removed. One might therefore recognize a

²⁵ Resistance to the idea that facts have sentence-like structure has surfaced recently, especially on the basis of the Russell-Myhill paradox (e.g. Dorr 2016, Goodman 2017). Whilst I cannot adequately address this delicate issue here, we should distinguish the idea that facts are structured from the view that truths with different structures represent different facts. For example, one might hold that <It is not not the case that snow is white> and <Snow is white> are factually equivalent (thus rejecting the latter view) whilst maintaining that the fact represented is structured. See Bacon 2020: §4 for a view along these lines.

²⁶ Along similar lines, one might require that the constituents of a perspicuous truth be ‘metaphysically primitive’, in the sense that they have no real definition. We can use Dorr’s (2016) framework of ‘identification’ to understand the relevant notion. On Dorr’s preferred view (2016: 72), identification obeys a non-circularity constraint: if there is a true identification of the form ‘ $y = z(x)$ ’, then there is no reverse identification of the form ‘ $x = z^*(y)$ ’ (where z and z^* are non-logical). Thus, we can understand a real definition of ‘ y ’ as any true identification of the form ‘ $y = z(x)$ ’ (where z is non-logical).

relative notion of perspicuity, with reduction a matter of transforming some non-perspicuous truth into a more perspicuous (but perhaps not perfectly perspicuous) truth. For example, one might view <the mean molecular energy of the air in the room is x > as more perspicuous than <the room temperature is y >, but still in need of further reduction to yet more perspicuous truths about the masses and velocities of the particular air molecules. Only once we reach perfectly perspicuous truths — those which perfectly reflect reality itself — is reduction complete.²⁷ Although I take the absolute notion of perspicuity to be my official primitive here, there is an important project of characterizing the relative notion (which may or may not be prior to the absolute notion).²⁸

Just as generation can be seen as ‘metaphysical causation’, so reduction can be seen as ‘metaphysical semantics’: it provides something like the ‘metaphysical meaning’ of the reduced truth (Sider 2011: §7.4). Since it requires perspicuity, reduction is more constrained than ordinary semantics — in Sider’s (2011: 112) words, it should ‘show how what we say fits into fundamental reality’. This is no constraint on ordinary meanings, which needn’t be ‘metaphysically privileged’ over the object-language itself. In another way, however, reduction is less constrained: it needn’t integrate with linguistics and psychology. Where ordinary semantics requires cognitive/representational equivalence, reduction only requires factual/worldly equivalence.²⁹

²⁷ From this perspective, one might even deny that there are any perfectly perspicuous truths: we may instead envisage a series of increasingly perspicuous truths ‘tending towards’ reality but never reaching it. (Perhaps reality has an inherently ‘non-semantic’ structure which is necessarily distorted by any truth-bearer.)

²⁸ Relative perspicuity might be understood in terms of relative naturalness, or via Dorr’s (2016) notion of ‘metaphysical priority’.

²⁹ The details here are subtle. Reduction can’t be completely divorced from semantics: for example, water reducing to H_2O is not independent of ‘water’ referring to H_2O (see Williams 2012: §6.3 for discussion). One plausible constraint is that a sentence and its meta-language equivalent in a good semantic theory ought to have the same reduction.

Both generators and reducers hold that ground is (at least typically) irreflexive: grounds and grounded are distinct. But where the generator sees a worldly distinction between facts, the reducer sees a merely representational distinction between truths.³⁰

It is important to note that, although claims about reduction pertain directly to representation, they are also (albeit, indirectly) about the reality being represented: in particular, to say that $\langle p \rangle$ reduces to $\langle q \rangle$ is to say that the fact represented by $\langle p \rangle$ has a structure matching (or at least, better approximated by) that of $\langle q \rangle$. I suspect that some feel an instinctive aversion to the reduction approach, on the basis that metaphysics is concerned with reality itself, independently of how we think or speak.³¹ This reaction is misplaced: reducers should not be seen as lapsed metaphysicians who conflate representation and reality, but rather as approaching reality *via* representation. This approach is no more confused than the idea that in ‘reverse-engineering’ a landscape from an impressionist painting we ought to take into account the stylistic distortions of the artist. Indeed, from a reducer’s perspective, it is generators who conflate representation and reality — it is as if, in comparing the painting with a photograph of the same landscape, they conclude that there must be two landscapes: the photographic one, and the impressionist one it generates!³²

³⁰ Jenkins (2011) challenges the irreflexivity of ‘metaphysical dependence’ on the basis of cases such as pains both depending on and being identical with brain states. However, such cases may be better interpreted as involving reduction.

³¹ Indeed, the opacity of perspicuity/reduction appears to belie Dorr’s (2016: 44) conviction that ‘any operators we might need to appeal to in stating questions that are central to the subject matter of metaphysics should be transparent’.

³² A distinction is sometimes drawn between ‘representationalist’/ ‘conceptual’ grounding and ‘worldly’/ ‘metaphysical’ grounding (e.g. Correia 2010: 258-9; Correia & Schnieder 2012: 21; Correia & Skiles 2017: 656). Those introducing the distinction often disavow the representationalist notion as conflating ‘mere shadows of language with real features of the world’ (Kramer & Roski 2015: 60). As I see it, this distinction is internal to the generation approach — concerning the granularity of facts — rather than corresponding to the broader distinction of interest here. The key difference between the two distinctions is that no distinctive notion of fundamentality is associated with representationalist as opposed to worldly grounding, whereas reduction is tied to perspicuity as opposed to basicness.

2.3 Truth-making

Before turning to the significance of the distinction between reduction and generation, a final point of taxonomy. Having introduced generation as a relation between worldly facts and reduction as a relation between representational truths, it is natural to wonder whether there is some intermediate relation by which truths are connected to underlying facts. The relation of ‘truth-making’ — understood as holding between truths and the facts that make them true — is an obvious candidate (e.g. Mulligan, Simons & Smith 1984, Armstrong 2004, Heil 2016).

The idea that truth-making provides us with a distinctive third approach to metaphysical explanation is certainly worth exploring. However, I am not inclined to separate truth-making from generation and reduction in this way. At least on the face of it, typical truth-making explanations seem to belong to one or other of the two approaches (or, in some cases, to combine them).

Truth-making explanations connect the truth of some truth-bearer to the existence or obtaining of some corresponding entity or fact — for example:

<Snow is white> is made true by [Snow is white].

<a exists> is made true by a.

On one interpretation, these explanations are straightforward applications of the generation approach to the semantic realm: that part of reality which concerns truth. In particular, they may be understood as corresponding to the following generation relations:³³

[<Snow is white> is true] is generated by [Snow is white]
(or by [[Snow is white] exists]).

³³ Schaffer 2008 and Fine 2012: §1.3 might be taken to suggest a generative conception of truth-making. As Fine notes (p.43), one might think that [<p> is true] is generated by [p] together with [<p> represents that p]. (Of course, this is sensitive to how one thinks of truths: cf. Lewis 2001: 604.)

[<a exists> is true] is generated by [a exists].

An alternative interpretation is available to reducers. It might be thought that talk of truth is non-perspicuous, and reduces to those non-semantic aspects of reality which make truth-bearers true. On this picture, the truth-making explanations above may correspond to the following reductions:

<<Snow is white> is true> reduces to <Snow is white>
(or to <[Snow is white] exists>.)

<<a exists> is true> reduces to <a exists>.

Some applications of truth-making seem to target not (merely) the truth of some truth-bearer, but instead (or additionally) the corresponding fact/truth itself. Some of these applications seem to lend themselves more naturally to one approach than to the other. Consider, for example, the way in which logically complex truths are supposed to be made true:

<Snow is white or pigs fly> is made true by [Snow is white].

<There are philosophers> is made true by [Socrates is a philosopher].

These cases do not seem to fit the ‘biconditional’ nature of reduction. Reduction requires factual equivalence between the reduced truth and the reducing truth, whereas logically complex truths do not seem factually equivalent to the logically simple truths corresponding to their truth-makers. For example, it seems strained to hold that <There are philosophers> is factually equivalent to, or ‘represents the same fact as’, <Socrates is a philosopher> (especially since it would by the same token be factually equivalent to <Plato is a philosopher>, making factual equivalence intransitive). Instead, these truth-making claims are naturally understood as corresponding to the following generation relations:

[Snow is white or pigs fly] is generated by [Snow is white].

[There are philosophers] is generated by [Socrates is a philosopher].

By contrast, other applications of truth-making have a decidedly more reductive flavor:

<There is water in the glass> is made true by [There are H₂O molecules in the glass].

<The temperature is y> is made true by [The mean molecular energy is x].

These truth-making explanations are naturally interpreted as corresponding to the following reduction relations:³⁴

<There is water > reduces to <There are H₂O molecules>.

<The temperature is y> reduces to <The mean molecular energy is x>.

Although the truth-making explanations invoke facts as truth-makers, this seems to be a merely superficial difference from the corresponding reductions to perspicuous truths, given the correspondence between these truths and facts. Indeed, for this reason, it is more important that the reduction approach targets truths than that it connects them to other truths. Connecting to truths does have the advantage of flexibility, however: we can allow for intermediate reduction to more perspicuous though not fully perspicuous truths, and we can chain these intermediate reductions together.

Finally, it might be thought that some applications of truth-making cannot be subsumed under either generation or reduction since, on one hand, they lack the biconditional character of reduction, and yet, on the other, they lack the factual character of generation. Take, for example:

<There is water> is made true by [a is an H₂O molecule].

³⁴ Cameron 2008 and Schulte 2011 suggest a reductive conception of truth-making.

It seems that this truth-making explanation cannot be subsumed under reduction, since the targeted truth is not factually equivalent to the truth corresponding to its truth-maker. And yet it may not be subsumed under generation either, since there may be no fact in reality corresponding to the targeted truth (talk of water may be non-perspicuous).

A natural way to understand examples like this is as combining reduction and generation:

<There is water> reduces to <There are H₂O molecules>.

[There are H₂O molecules] is generated by [a is an H₂O molecule].

This relies on there being logically complex facts to be suitable targets of generation. If there is no logical complexity in reality itself, then truth-making (or something akin to it) may indeed be required to provide the connection between logically complex truths and the facts they latch onto.³⁵

3. Significance

I have been speaking as if metaphysicians divide into two camps: generators and reducers. But of course, things aren't so neat. Many metaphysicians do not clearly have either conception in mind (and sometimes seem to veer between them). I suspect that most would regard the two approaches as nuanced variations on the same core idea, and even those who consistently employ one approach rarely do so in self-conscious opposition to the other.

This section argues that the distinction between generation and reduction deserves attention. It enables us to better understand what is at issue in many metaphysical debates and the arguments given by each side (which often implicitly trade on the underlying approach). I focus on some

³⁵ In §5.2 I discuss whether we can make sense of a non-biconditional notion of reduction which provides this connection.

particular case studies, but I hope my discussion will indicate that the distinction's impact on metaphysics is pervasive.

3.1 Ontology

In paradigmatic ontological disputes, e.g. concerning ordinary objects, or abstracta, articulating the disagreement can be difficult. As Bennett (2009) describes, both sides 'minimize the difference', playing down their commitment to/rejection of the controversial entities.

Take mereology. 'Conciliatory nihilists' like Cameron (2010) and Sider (2013a) hold that the table before me exists, but only in a non-fundamental sense: it exists because there are some simples (objects with no smaller parts) arranged table-wise. These simples being arranged table-wise necessitates the existence of the table, but the connection needn't be analytic or a priori: it may be doubted by those who are competent with the relevant concepts. This creates a puzzle: what makes this view 'nihilist'? After all, many 'conciliatory anti-nihilists' agree that the table's existence is non-fundamental, and explained by some simples being arranged table-wise via a necessary a posteriori connection.

Distinguishing two approaches to metaphysical explanation allows us to understand the underlying dispute. Conciliatory nihilists hold that truths about composite objects reduce to truths about simples. On this view, reality itself is merely a swarming of simples; it lacks any compositional structure matching our ordinary conception. For their conciliatory anti-nihilist opponents, by contrast, facts involving composites are generated by facts involving simples. Our talk of composites is not merely true on this view, but reflects reality itself.

Many, I believe, have an inchoate grasp of disagreements like these, and are only skeptical insofar as they cannot clearly articulate them. It is a point in favor of the distinction that it allows us to do so. Moreover, the distinction enables us to better understand the arguments given on each side: the arguments given by 'believers' in the entities in question typically suit a generation approach, whereas those given by 'deniers' typically suit a reduction approach. This is unsurprising: the attractive conciliatory versions of these positions — according to which the entities in question

exist non-fundamentally — are only available once the corresponding approach is recognized. Recognizing both approaches allows us to properly evaluate the arguments.

On the believers' side, take Schaffer's (2009: 360) insistence that the existence of controversial entities is (almost always) obvious: ordinary material objects, numbers, properties, and even fictional characters all exist in the one true non-deflationist sense, as 'full-blown "heavyweight"' entries on the roster of entities'. On Schaffer's view, the interesting question is how these entities are generated by more basic entities. He (2009: 358) motivates this 'permissivism' by various Moorean arguments, such as:

My body has proper parts (e.g., my hands).

Therefore, there are things with proper parts.

This argument does not impress conciliatory nihilists, who deny the perspicuity but not the truth of its conclusion. What, then, underlies Schaffer's dismissive attitude? It is that his framework leaves no room for conciliatory nihilism; since he understands fundamentality as basicness, he interprets the claim that bodies exist 'non-fundamentally' as entailing that they are derivative entities which exist in the one true sense! Thus, the only nihilist opponent he recognizes is revisionary.³⁶ The real issues emerge once we have reduction in sight. For example, are we justified in upholding the perspicuity of Moorean claims? And are we justified in believing that these claims are true, as opposed to 'apt' in some weaker sense?³⁷

On the deniers' side, take Sider's (2013a: §1) argument that nihilism is more ideologically parsimonious since it 'allows us to eliminate 'part' from the ideology of our *fundamental* theories... If one's theory of fundamental matters included an ontology of composite objects, then

³⁶ Cf. his description of Sider's approach to the non-fundamental as 'radically eliminative' (2013: 736).

³⁷ See Sider 2013a: §4 on the metaphysical shallowness of the meta-semantic issue of truth. This becomes important if the English quantifier and the perspicuous quantifier cannot differ in meaning; see Dorr 2014.

that theory would presumably also need a predicate of parthood to connect those composites to their parts...’. This won’t impress conciliatory anti-nihilists: their ‘fundamental theory’ only talks about simples, not the derivative fusions they generate! Sider is evaluating parsimony by the notions which figure in a perspicuous theory, whereas the anti-nihilist I am imagining evaluates parsimony by the constituents involved in the basic facts. Thus, applying the distinction reveals the key issue (to which I return shortly): which sense of fundamentality — basicness or perspicuity — is relevant to parsimony?³⁸

3.2 Meta-ontology

For ‘highbrow’ ontologists, ontology concerns what exists fundamentally:

What we debate in the ontology room is ... what there is *in the most fundamental sense* (Dorr 2005: 24).

... ‘a really (or fundamentally) exists’ is true iff a is an element of our ontology (Cameron 2008: 7).

The ontological question is... “Are there Fs?” where ‘there are’ is understood as having a fundamental sense (Sider 2011: 171).

On this view, when nihilists and anti-nihilists debate the existence of composite objects, nominalists and platonists debate the existence of abstracta, and so on, their disagreement concerns what exists *fundamentally*.

³⁸ It is tempting to argue that the anti-nihilist’s complete fundamental theory must include not only the basic facts but the principles of generation connecting them to derivative facts, and parthood (or the like) is needed to state these connections (Sider 2011: §8.2.1). Sider’s view avoids the analogous issue since he takes truths about reduction to be reducible. But anti-nihilists might defend the parallel view that facts involving generation are themselves generated by basic facts not involving parthood (Bennett 2011b, deRosset 2013; cf. Bennett 2017: 226).

Opposed to this elitism are ‘lowbrow’ ontologists, who separate ontology from fundamentality:

Of course not every entity is a fundamental entity... (Schaffer 2008: 18)

Fundamentality does not lead to greater reality. (Hofweber 2016: 329)

...flatworldism is the result of taking to its extreme the thought that nonfundamentalia are less than fully real. I myself have never seen the appeal of this thought... (Bennett 2017: 216).

On this view, ontology (even ‘serious’ ontology) concerns plain old existence — and non-fundamental things exist too!³⁹

These two views are *meta*-ontological: they characterize ontological disagreements rather than addressing them. But of course meta-ontology guides ontology; insofar as one is drawn to the moderate claim that some controversial entity exists non-fundamentally, one will be a denier/believer if one is a highbrow/lowbrow ontologist.

The disagreement between highbrow and lowbrow ontologists can be perplexing. On one hand, it is liable to seem verbal, with the disputants merely employing the technical term ‘ontology’ differently (cf. Williams 2012: 183). Both sides distinguish three categories: fundamental, non-fundamental, and non-existent (Kriegel 2015). The disagreement is then whether ontology concerns what goes in the first category or what goes in the last. But both are good questions; it doesn’t matter which we call ‘ontology’! Indeed, even the disputants themselves sometimes suggest this reconciliation.⁴⁰

³⁹ Along similar lines, von Solodkoff & Woodward (2013: 568) distinguish ‘inflationism’ about the non-fundamental (~lowbrow ontology) from ‘deflationism’ (~highbrow ontology).

⁴⁰ See Schaffer’s (2008: 18) response to Cameron, and Sider’s (2013c: 761) response to Schaffer.

On the other hand, the disputants' understanding of ontological disputes as substantive belies the idea that the background meta-ontological dispute is merely verbal. For example, Schaffer (2009) and Bennett (2017) are believers in ordinary objects, whilst Cameron (2010) and Sider (2013a) are deniers. There is an obvious reconciliation of these 'ontological' positions: Cameron and Sider address the highbrow question of whether ordinary objects exist fundamentally, whereas Schaffer and Bennett address the lowbrow question of whether they exist *simpliciter*.

Distinguishing generation and reduction clears this up. Highbrow ontologists understand fundamentality in terms of perspicuity —their view is that ontology concerns what exists perspicuously (that is, which objects the perspicuous truths are about).⁴¹ Lowbrow ontologists understand fundamentality in terms of generation — in shunning elitism, they extend ontology to derivative entities (which are no less real).⁴² It is because they have different background conceptions of fundamentality that existing non-fundamentally is 'really' existing for lowbrow but not highbrow ontologists. Since they mean different things by 'fundamental' (*not* 'ontology') the meta-ontological disagreement is verbal after all. This reconciliation accommodates substantive ontological disputes: deniers hold that non-fundamentalia only exist in a non-perspicuous sense, whereas believers hold that they are derivative entities (cf. Sider 2011: 170; 2013a: 253).⁴³

Distinguishing the two approaches also exposes an underlying disagreement on the closely connected question of ontological parsimony. Cameron, Sider, Schaffer and Bennett appear to share the view that only commitment to fundamental entities counts against the ontological parsimony of a theory (leading Cameron (2010: n.2) to speculate that their dispute is 'merely terminological'). But in fact, their views are importantly different: what counts for Cameron (2010:

⁴¹ Not: which truths of the form 'there are Xs' are perspicuous. Table-believers might recognize perspicuous truths about tables whilst denying that 'there are tables' is perspicuous, because either the predicate 'table' or quantification itself may be non-structural.

⁴² Although, see Rosen (2011).

⁴³ Highbrow ontology is an instance of the broader idea that metaphysics concerns what is fundamental, as criticized by Barnes (2014) and Bennett (2017: §8.3). Whichever conception of fundamentality we adopt, Barnes and Bennett are surely correct about the broader point: important metaphysical questions needn't concern either basic facts or perspicuous truths.

250) and Sider (2011: 198) is what exists in a perspicuous sense, whereas Bennett (2017: §8.2.2) and Schaffer (2015) only count the basic entities. This disagreement is a significant driver of ontological disputes: deniers typically reject the believers' claim that commitment to derivative entities is non-costly. (More generally, we can consider what counts against 'ideological parsimony': the constituents of perspicuous truths or of basic facts?)

Two initial points concerning the underlying disagreement are salient. First, insofar as parsimony is tied to explanatory efficiency — the principle being: 'minimize the unexplained' (Schaffer 2009: 361; Bennett 2017: 221) — then derivative structure seems non-costly. Why should it matter that some truth is irreducible, if the corresponding fact can nonetheless be explained via generation? Second, though, an important motivation for denying that non-fundamental structure is costly is the elusive idea that it is 'nothing over and above' the fundamental (Lewis 1991: 81), or 'no addition to being' (Armstrong 1997: 12). For generators, since derivative entities/facts are (numerically) distinct from the basic entities/facts generating them, these slogans are awkwardly misleading (prompting Rosen (2017: n.3) to propose the replacement 'hardly anything over and above').⁴⁴ Reducers, by contrast, can make straightforward sense of the idea. If the table's existing reduces to some particles being arranged table-wise, then: (i) the table is nothing over and above the particles in that there is no worldly distinction between truths about the table and truths about the particles, and (ii) the table is no addition to being in that there are no perspicuous truths about it.

3.3 Realism

An important theme emerges from the discussion so far: in a particular sense, generators take a realist (or inflationist) attitude towards the non-fundamental, whereas reducers take an anti-realist (or deflationist) attitude.

This joint between the two approaches is striking, but spelling it out takes care: 'realism' is a slippery word which is used in many ways. To my mind, the difference in question is best

⁴⁴ Cf. Audi 2012: 709, Bennett 2017: 222.

articulated in terms of perspicuity: when reducers target a non-fundamental proposition, they treat it as a non-perspicuous truth, whereas when generators target a non-fundamental proposition, they treat it as a worldly fact — and hence, view the corresponding truth as perspicuous.

This is vivid in the ontological case. Generation enables a ‘permissivist’ position, according to which the entities in question (e.g. mereological wholes) are ‘real’ but non-basic: truths about them are perspicuous. By contrast, reduction supports a ‘minimalist’ position, according to which the relevant entities only exist in a non-perspicuous sense. In this way, generators take non-fundamentalia more seriously: they accommodate non-fundamental truths about them by positing corresponding entities in reality itself, whereas reducers accommodate the same truths without positing such entities.

But the difference is much more general than this: generators are not just realist about what exists non-fundamentally but, more broadly, about how things are. Consider nominalists who take certain ‘worldly’ predicates, such as *has mass*, more seriously than ‘made-up’ predicates, such as *is grue*. This is not an ontological matter: it does not concern whether the corresponding properties exist, since this is denied in each case. Rather, it is ‘ideological’: they hold that only some predicates reflect ‘real’ (not merely representational) aspects of the world. This difference in attitude may manifest in their approach to non-fundamental propositions expressed using these predicates. For example, they might naturally say that [x has 1kg mass] generates [x has mass], but not that [x is green and observed before t] generates [x is grue], since the world provides no distinctive fact involving x’s grueness. As a ‘made-up’ predicate, *grue* is an unsuitable means of specifying targets for generation; truths about grueness are not ‘self-standing’ but collapse into truths about color and observation times.⁴⁵

There is thus a general (not merely ontological) sense in which generators are more ‘realist’ about the non-fundamental. This realism shouldn’t be confused with conservatism: the view that our ordinary conception of the world is largely correct. Realism (in the sense I have been discussing)

⁴⁵ Since a non-reifying version of the generation approach is available, the dispute is not best cast as an ontological question concerning the fact [x is grue].

concerns the status of the non-fundamental explananda one accepts, whereas conservatism concerns the prior question of which putative explananda should be accepted. Insofar as neither generators nor reducers regard the acceptance of non-fundamental explananda as detracting from the simplicity of their theory, conservatism carries a strong pull for both.⁴⁶

Nonetheless, I suspect that generators' realism may to some extent be attributed to a certain reverence for our ordinary conception which reducers do not share. Contrast 'Weak Mooreanism', the methodological claim that commonsensical Moorean propositions should be presumed true (or at least correct), with 'Strong Mooreanism', the claim that there should be a presumption in favor of the perspicuity of Moorean propositions. To a reducer's eyes, Strong Mooreanism likely seems misguided: perspicuity is appropriately evaluated by wielding Occam's razor in the light of our best scientific theories (Sider 2011: 13) — or by exploiting its connections to surrounding notions — not merely by reflecting on common sense.⁴⁷ However, something like Strong Mooreanism — together with the idea that many Moorean propositions are non-fundamental — may implicitly guide the realist commitments of some generators.

This methodology may be motivated by reflecting on certain 'super-Moorean' propositions:

I exist.

I am conscious.

There is an external world (things outside me).

Time passes (the past was present, the future will be present).

I persist (the person that stands up is the one that sat down).

I control my actions.

Some actions are morally wrong.

⁴⁶ Although, highbrow ontologists may defend radical-sounding ontological theses whilst maintaining the correctness of mundane existence claims.

⁴⁷ G. E. Moore himself likely opposed Strong Mooreanism, given his openness to the 'analysis' of propositions about the external world in terms of sense-data (1925: §IV).

These propositions are commonsensical in the extreme way of being utterly central to our shared ‘sense of reality’, so that denying them is, as Jerry Fodor (1990: 156) puts it, ‘the end of the world’. The idea that such beliefs distort the way things really are — that they are somehow metaphysically ‘off-track’ — can be profoundly disturbing. It is therefore tempting to regard them as deep insights rather than mundane truisms; when we reflect on ourselves, on time, or on morality, we seem to make direct contact with reality itself. This conviction can be seen to underlie a swath of views which are realist in the sense discussed above: non-reductive theories of persons and mentality, A-theories of time, endurantist theories of persistence, libertarian accounts of free will, non-naturalist versions of moral realism, and so on. These views reject the self-alienating idea that super-Moorean beliefs reduce to some austere, unfamiliar truths.

This lends Strong Mooreanism some credibility; if super-Moorean propositions may be presumed perspicuous, regarding other Moorean propositions as irreducible ‘rounds out’ the picture in a natural way. But the analogous position with respect to generation — that Moorean propositions are presumably ungenerated — is not similarly motivated; nothing seems to suggest that our super-Moorean beliefs make direct contact with *basic* reality. We may simply find ourselves metaphysically as well as temporally ‘downstream’. Here too, therefore, distinguishing reduction and generation affords a better understanding of both the disagreement between realist and anti-realist views and their motivations.

4. Adjudication

The choice between reduction and generation matters. How, then, can we tell which of the two approaches is better suited to a given case?

One strategy is to exploit the connections to surrounding notions that I used to introduce the distinction. Most prominently, if $\langle p \rangle$ reduces $\langle q \rangle$, then they latch onto the same fact, whereas if $[p]$ generates $[q]$, then (at least typically) $[p]$ and $[q]$ are distinct. Thus, a theory of fact identity would provide a significant means of adjudication. Likewise, given the connection between perspicuity and Siderian structure described above, a theory of the latter will also be powerful: if $\langle p \rangle$ reduces $\langle q \rangle$, then $\langle q \rangle$ must involve some non-structural notions.

However, fact identity and structure are complex and controversial matters, and it is by no means obvious that they can or should be settled prior to adjudicating between generation and reduction. For example, one might deny that truths latch onto the same facts as their self-conjunctions on the basis that conjunctive facts are generated by their conjuncts, or one might deny that conjunction is a structural notion on the basis that conjunctive truths reduce to their conjuncts. It is likely that several packages of views about fact identity, structure and the scope of generation/reduction are attractive.

This section aims to provide some independent resources for adjudication. The recurring theme is that generation concerns the way the world itself works whereas reduction concerns the way language fits the world. This allows us to discriminate between them on the basis of both the explananda targeted and the nature of the explanatory connections provided.

4.1 Explanatory targets

When we try to explain some proposition using the generation approach, we treat it as a worldly fact, outputted by a worldly process. Hence, taking the generation approach towards indeterminate and/or context-relative propositions involves positing some form of worldly indeterminacy and/or context-relativity. Insofar as reality is plausibly both determinate and absolute, this provides good reason to favor the reduction approach in such cases.

Take, for example, the truth that Tibbles the cat exists. It is uncontroversial that there is no precise cat-like object uniquely worthy of the name ‘Tibbles’; for many particles around Tibbles’s boundary, it is indeterminate whether they are part of Tibbles (think of hairs in the process of being shed). Hence, if <Tibbles exists> latches onto a single fact, it must involve an imprecise cat. But imprecise cats are implausible posits: following Lewis (1999) and McGee & McLaughlin (2000), we should locate the indeterminacy in our language rather than reality itself.⁴⁸

⁴⁸ For defense of worldly indeterminacy, see Barnes & Williams 2011, Wilson 2013. Defending the general idea is one thing, defending its application to particular cases another.

Hence, before we can apply the generation approach, we must first associate this truth with some set of facts: its admissible precisifications. For example, we might associate it with some facts concerning the existence of a range of cat-like micro-aggregates of particles in Tibbles's vicinity. This is an instance of reduction: <Tibbles exists> is being reductively explained in terms of its admissible precisifications.⁴⁹ The association of Tibbles with some cat-like micro-aggregates — as opposed to some cat-like spacetime regions, or some cat-like undulations across the gravitational field, say — is an explanatory hypothesis about what in reality supports <Tibbles exists>.

Similarly, take Lewis's (1983b) counterpart-theoretic account of de re modality. On this view, *x*'s possibly being *P* is metaphysically explained by *x*'s bearing some counterpart relation *R* to some *y* which is *P*. According to Lewis, our de re modal talk is indeterminate between many candidate counterpart relations. For example, there is no determinate answer to the question: how tall could Trump have been? Candidate counterpart relations differ on the height of the tallest Trump-counterpart. For this reason, it would be awkward to view Lewis's account as describing the generation of worldly modal facts, including 'indeterminate facts' about Trump's possible height. It is much more naturally viewed in terms of reduction: as providing a way to understand modal talk in terms of an entirely non-modal reality. On this interpretation, it is not that a modal truth latches onto some indeterminate aspect of reality, but that it is indeterminate which aspect of reality it latches onto.

Another reason to view Lewis's account of modality as exemplifying reduction is that he takes the counterpart relation to be context-relative: in one context, our modal talk is underwritten by counterpart relation *R*₁, and in another, it is underwritten by relation *R*₂. Taking the generation approach in this case would therefore require the strange idea that different modal facts are generated in different contexts — or perhaps, that the modal facts fragment into several realms,

⁴⁹ Depending on the semantics for indeterminate referring expressions like 'Tibbles', it might be that each admissible precisification fully explains <Tibbles exists>, or that all the precisifications collectively do. Either way, it is indeterminate how <Tibbles exists> reduces, since it is indeterminate which precisifications are 'admissible' (i.e. which perspicuous truths represent the same fact as it).

each generated by the facts involving a corresponding counterpart relation. The reduction approach, by contrast, requires only facts involving the various counterpart relations, with the context-relativity built into how our modal talk latches onto them: the world remains both absolute and unfragmented.

For further illustration of context-relativity, consider claims about simultaneity, e.g. the clock strikes five in New York as the clock strikes ten in London. In light of Special Relativity, reality itself has no absolute simultaneity structure, and in this sense such statements do not straightforwardly correspond to any facts. A hot-headed reactionary might hold that this makes simultaneity claims either false or meaningless, but another view counsels cool-headed conciliation: physics just tells us something important about what makes these claims true.⁵⁰ In particular, simultaneity talk is supported by a three-place relation between two events and a contextually supplied reference frame. The context-relativity needn't be built into our linguistic understanding, since our frames of reference (determined by our motion) typically agree on what is (approximately) simultaneous: we move through practically the same 'specious presents'.⁵¹

We could understand this view as fragmenting reality into separate realms of simultaneity facts, one per reference frame, with the facts within each realm generated by the facts about which pairs of events stand in the relevant three-place relation R to the corresponding reference frame (as suggested by Fine 2005: §10).⁵² But it is much more naturally understood in terms of reduction, with truths within a context reducing to truths about pairs of events standing in R to the contextually supplied reference frame. Instead of fragmenting reality itself, this approach simply

⁵⁰ Even supposing that all such statements fall short of truth, a distinction surely remains between those which are apt or 'approximately true' and those which are straightforwardly false. The explanatory demand might then be reconceived as concerning this weaker status.

⁵¹ More generally, appeal to context is much less costly in theories of reduction than in ordinary semantics, since it needn't be accessible to competent speakers.

⁵² For similarly 'fragmentalist' proposals, see Hare 2009, Lipman 2016, and Spencer 2016.

understands simultaneity talk as latching onto reality in a context-relative way, with reduction revealing how context contributes to determining which fact is represented.⁵³

4.2 Explanatory goodness

When evaluating the plausibility of an explanation, we are guided by certain epistemic principles. Two constraints seem especially important. First, explanations should be systematic: they should invoke reasonably concise general principles that subsume many cases at once, rather than proceeding on an ad hoc, case-by-case basis. For example, there is a function which maps, for each object x , the truth about x 's color to the truth about x 's location. But since this function does not obey any simple rule, it is not plausibly taken to describe an explanatory connection. Second, explanations shouldn't be arbitrary: the general principle invoked ought to be privileged in some way over relevant alternatives. For example, the mereological principle that composition always and only takes place on Mars is implausible because it unjustifiably assigns Mars a special status.⁵⁴ (Note that this principle seems acceptably systematic, so systematicity and non-arbitrariness come apart.)⁵⁵

These constraints apply differently in the case of generation and reduction, however. Generative connections describe how worldly entities are built up from others, whereas reductive connections describe how our non-perspicuous conception 'projects' onto reality. Consequently, what it is for

⁵³ I am envisaging the reducer making a claim of the following form: in a context C , $\langle x$ is simultaneous with $y \rangle$ reduces to $\langle x$ and y stand in R to $F_C \rangle$, where F_C is the reference frame supplied by C , and $\langle x$ is simultaneous with $y \rangle$ is the truth expressed by sentences of the form 'x is simultaneous with y' in the given context (see n.14). The explanatory advantage of the reduction approach here crucially involves the ability to state a meta-linguistic principle which generalizes across contexts; it cannot be achieved simply by making the various corresponding reduction claims (within various contexts).

⁵⁴ For arbitrariness arguments, see Lewis 1986: 212-3; van Inwagen 1990: 126-7; Sider 2001: §4.9. For discussion see Korman 2010, Fairchild & Hawthorne 2018, Builes 2021.

⁵⁵ It is an open question how they are connected. From one perspective, systematicity is a matter of being subsumed under a simple theory, and non-arbitrariness is a matter of being subsumed under a theory which is simpler than the relevant alternatives.

these connections to be systematic and non-arbitrary looks quite different in the two cases. This provides another means of adjudicating between the approaches: an explanation may be plausibly systematic and non-arbitrary on one approach but not on the other.

With respect to systematicity: in the generative case, this constraint reflects the assumption that the world works systematically, with respect to both dynamical evolution through time and metaphysical generation across levels (Kment 2014: 5). We should expect that, as in the causal case, all generation ultimately derives from some concise list of ‘root principles’ (perhaps: principles of logic and/or connecting determinates to determinables, and set formation and/or mereology). Hence, generation should be ‘nomicallly’ systematic in the way that causation is — backed by some sparse ‘lawbook’ of simple principles.⁵⁶

By contrast, in the reductive case, the systematicity constraint reflects the assumption that language works systematically, with respect to both ordinary meaning and metaphysical analysis (Sider 2011: 118). We should expect that, as with ordinary semantics, reduction may be subsumed under a reasonably systematic meta-semantic account, which explains why a notion reduces in the way it does in terms of its meta-semantic profile (e.g. its history of use, its causal/nomic connections, etc.). Hence, reduction should be ‘meta-semantically’ systematic in the way that semantics is — backed by some ‘dictionary’ of metaphysical analyses (together with some principles of meaning composition).

To illustrate the difference between nomic and meta-semantic systematicity, consider the task of explaining conscious experiences in physical terms. Suppose this task is accomplished by pairing phenomenal properties with functional properties: *x* is in pain because *x* is in a state which plays nomic role R1, *x* sees red because *x* is in a state which plays nomic role R2, and so on. If this connection describes generation, it seems problematically case-by-case: the list cannot be subsumed under any root principle of functional-to-phenomenal generation, since the outputted

⁵⁶ Witness the widespread practice of formulating general principles describing how facts of certain kinds are grounded. Schaffer (2017) argues that metaphysical explanation requires ‘laws of metaphysics’, and Wilsch (2016) offers an account of metaphysical explanation in terms of these laws. These authors are best interpreted as focusing on generative explanation.

phenomenal properties cannot be concisely defined in terms of the inputted functional properties.⁵⁷ However, if this connection describes reduction instead, it may be suitably systematic: the list might be subsumed under a meta-semantic account which explains, in terms of the meta-semantic profiles of our phenomenal concepts, why they reduce to the corresponding functional properties. Phenomenal truths would then reduce compositionally via the metaphysical meanings of the predicates they involve.

With respect to arbitrariness: in the generative case, it is costly to posit a principle describing how the world works which fails to stand out over relevant alternatives. For example, it is costly to posit a mereological law which specifies how tightly some parts must be stuck together in order to compose, just as it is costly to posit a dynamical law which specifies some threshold mass-value above which particles no longer move according to 'F = ma'. Generative connections should have a kind of 'objective' privilege, conferred by reality itself.

By contrast, in the reductive case, there should be some account of why it is that we represent one aspect of reality rather than some relevantly similar aspect. Such an account may appeal to somewhat parochial features of our own 'window on the world' (where this would be an inappropriate source of privilege in the context of generation). For example, any reduction of our talk of composite objects which specifies how tightly some parts must be stuck together is implausible insofar as nothing in that talk seems to distinguish any particular degree of tightness. Reductive connections should have a kind of 'perspectival' privilege, which partly reflects our own vantage point.

By analogy, compare two constructions of the natural numbers. They can be built up from powers of ten: for example, 476 is $(4 \times 10^2) + (7 \times 10^1) + (6 \times 10^0)$. This construction is perspectivally but not objectively privileged: it is highlighted by our decimal system of representation, but nothing in the numbers themselves distinguishes it. Alternatively, they can be

⁵⁷ As Levine (1983: 357) remarks: 'there seems to be nothing about C-fiber firing which makes it naturally "fit" the phenomenal properties of pain, any more than it would fit some other set of phenomenal properties'.

built up from their prime factors: for example, 476 is $2^2 \times 7 \times 17$. This construction is objectively but not perspectively privileged: it reflects the deep nature of the numbers themselves, but appears strangely random from our own perspective.

To illustrate the difference between objective and perspectival privilege, consider the best system account of laws (Lewis 1983a, Loewer 1996). According to this account, laws are the theorems of the deductive system which best summarizes the ‘Humean mosaic’: the complete collection of ‘local matters of particular fact’. A system is good to the extent that it balances various virtues, especially simplicity (roughly, length in fundamental terms) and informativeness (roughly, amount of worlds excluded). Thus, its being a law that p is metaphysically explained by p ’s being a theorem of the best system (which is explained in turn by the Humean mosaic itself).

Arbitrariness provides strong reason to view this explanation as reductive.⁵⁸ As Lewis (1994: 479) admits, how the virtues are defined and traded off may partly be a psychological matter: the account is to be fleshed out in accordance with the role that laws play in our own cognitive lives and in the practice of science. What counts as the ‘best’ system may differ for aliens with a different psychology who practice science differently. We may therefore distinguish the ‘human’ best system from the ‘alien’ best system, with neither providing an ‘objectively correct’ notion of law. Thus, it seems problematically arbitrary to posit that p ’s lawhood is generated by its belonging to the human best system, since this explanatory connection lacks objective privilege: why would nomic facts be attuned to our conception rather than the alien conception?⁵⁹ But there is no analogous issue with positing that p ’s lawhood reduces to its belonging to the human best system, since this reduction is perspectively privileged: we would have an explanation — in terms of our

⁵⁸ Indeed, Lewis himself seems to have reduction in mind: he describes Humean Supervenience as the view that ‘all there is in the world is a vast mosaic of local matters of particular fact’ (1986: ix).

⁵⁹ Lewis famously hoped that the human best system is best ‘under any standards’ (see Massimi (2017) and Gómez Sánchez (forthcoming) for critical discussion). Even if this turns out to be true, it is not plausibly necessary, and so doesn’t provide the kind of objective privilege which a generative approach to the account would require.

psychology and scientific practices — for why it is that we latch onto this notion of lawhood rather than the alien notion.⁶⁰

5. Redundancy

So far, I have presented the distinction between generation and reduction, illustrated the significance of its impact on metaphysics, and suggested some ways of adjudicating between the two approaches. This final section argues that both approaches have a role to play in explanatory metaphysics; we can't rely on one of them alone.

5.1 Against pure generation

According to pure generators, all metaphysical explanation is (or is backed by) generation amongst facts.⁶¹ I know of no explicit defenses of this extreme position, and it might be that no generator would, upon reflection, endorse it. Nonetheless, the rhetoric of many generators, together with their readiness to widely apply their framework, suggests that it may not be uncommon as an implicit presupposition.⁶² Sociology aside, discussing the view is theoretically important for my purposes: it helps to bring out what makes reduction distinctively important.

Now, I have already argued above (§4.1) that certain imprecise truths, such as <Tibbles exists>, and certain context-relative truths, such as those concerning simultaneity, are better treated in terms of reduction than generation, at least insofar as there are no corresponding imprecise or context-relative facts. In each case, generators must find some precise, absolute facts associated

⁶⁰ In order to avoid arbitrariness, generators might posit plenitudinous generation: for example, in the case of the best system account, they might posit many principles of the form 'if p is a theorem of the best _{i} system, then p is a law _{i} '. This strategy still requires supplementation by (merely perspectively privileged) reduction: once we have generated the various laws _{i} , it remains to specify which of them our own notion represents.

⁶¹ At least, all explanation revolving around fundamentality and ground (see n.2).

⁶² A pure generation picture is (perhaps) suggested by Bennett 2017. I suspect that many generators would subscribe to Schaffer's (2017: 303) characterization of metaphysical explanation as 'an explanation backed by grounding relations' (where 'grounding' refers to generation).

with the truth in question. But this prior step is reductive: the non-perspicuous imprecise/context-relative truths are metaphysically explainable in terms of perspicuous truths which reveal the associated facts. These kinds of cases may have already convinced you that pure generation is untenable. Nonetheless, in this section, I wish to add a further kind of example, before considering what I take to be the pure generator's most natural response.

Consider a case of 'cross-level conflict': non-fundamentally, this table occupies a continuous, table-shaped region; fundamentally, this region is largely empty space, inhabited only by a sparse network of particles. How can it be both a non-basic fact that the region is occupied and a basic fact that it is not? Their obtaining at different 'levels' seems irrelevant — as Martin Lipman (2018: 597) puts it: 'stacking the facts does not help remove the conflict.'⁶³ Thus, in order to explain the region's being occupied, generators must first associate it with an appropriate fact, by clarifying the sense in which the region is 'occupied' despite being largely empty. But again, this prior association of the explanandum with a fact appears reductive: it reveals, in more perspicuous terms, what supports the truth in question.⁶⁴

Pure generators might insist that, in each of the cases I have highlighted — involving indeterminacy, context-relativity, and cross-level conflict — the association with a fact is not part of the explanation itself, but rather provides a suitably clear gloss on the explanandum being targeted. The idea is that the semantics in question does not amount to reduction: it is merely 'descriptive' rather than genuinely explanatory.

This response doesn't do justice to the phenomena in question. Take Tibbles. The association of 'Tibbles' with its precisifications seems importantly different from merely being taught the name, or having the name disambiguated, neither of which seems to yield the same kind of metaphysical understanding. If you were to say, upon learning that 'Tibbles' refers to some cat in the vicinity,

⁶³ Lipman (2018) argues that cross-level conflict makes trouble for 'fundamentality-based metaphysics' in general, but does not focus on reduction as a distinctive approach.

⁶⁴ Again, there's a radical response (akin to making reality imprecise and/or fragmented): pure generators might embrace an inconsistent reality.

‘Ah, now I see – Tibbles exists in virtue of *that cat* existing!’), you would not have advanced the cause of explanatory metaphysics. Nor should the understanding yielded by precisifying ‘Tibbles’ be conflated with that attained when an ordinary Frege puzzle is resolved. Having realized that Tibbles is the very same cat already known as ‘Hamish’, you would not satisfy metaphysicians by informing them that for Tibbles to exist is just for Hamish to exist. The notion of perspicuity allows us to capture the important difference between these ways of understanding <Tibbles exists>: it is only when it is glossed in more perspicuous terms that metaphysical explanation is achieved.

Similarly, pure generators might see the context-relative ‘explanation’ of truths about simultaneity as akin to the merely descriptive association of a context-dependent meaning with straightforwardly indexical truths, such as <I exist>. It does not seem especially natural to think that <I exist> may be understood in any metaphysical sense in terms of <x exists>, where x is the speaker/thinker in the relevant context. It seems more natural to say that these are just the same truth (in the context), and so no explanatory progress has been made.⁶⁵ By contrast, however, it is natural to think that the truth that two events are simultaneous may be metaphysically understood in terms of the truth that they are simultaneous relative to a contextually determined reference frame, and not at all natural to think that these are just the same truth (as opposed to representing the same fact in the relevant context). It is a substantive discovery about reality (not merely about language) that our simultaneity talk is misleading in this way.

Finally, on the proposed response, pure generators would conflate the resolution of cross-level conflict with ordinary polysemy: in one sense of ‘door’, she walked through the door, but in another (related) sense, she didn’t. In this case, clarifying that what she walked through was the airy gap within the doorframe pedantically spells out the speaker’s original intention. The task of explicating the sense in which the table-shaped region is occupied is importantly different. When we discover that this occupation is a matter of the dynamical laws underwriting the exclusion of

⁶⁵ Although, even here it might be argued that it is a genuine metaphysical discovery (albeit one typically made early in life) that the I-concept picks out something which may also (and indeed, more perspicuously) be regarded third-personally.

other (macroscopic) objects, we haven't merely described what was intended all along: we have explained what in reality supports the truth in question.⁶⁶

Some generators might still prefer to use 'explanation' in a narrower 'worldly' sense, according to which the reduction step may be a necessary prelude to explanation but is not itself The Real Thing. I don't wish to quibble about terminology; the substantive point is that characterizing worldly generation relations is not the only way to satisfy the aims of 'explanatory' metaphysics. On this narrower usage of 'explanation', neither the association of truths about pain with truths about C-fibers nor the association of truths about laws with truths about the Humean mosaic would constitute proposals for metaphysical explanation. But these hypotheses are surely intended as contributions to explanatory metaphysics (indeed, as crucial parts of the defense of overarching metaphysical positions such as physicalism and Humeanism), and if true, they yield the kind of understanding that is its distinctive aim.

This discussion has revealed that the reduction approach has resources which the generation approach lacks. Reduction can advance our understanding by focusing on features of representation itself, uncovering the ways in which it makes contact (or fails to make contact) with reality.⁶⁷ Reducers can appeal to semantic indeterminacy, avoiding an imprecise reality, to the representational context, avoiding a context-relative reality, and to fundamental and non-fundamental uses of an expression, avoiding an inconsistent reality. These resources have wide potential application: our ordinary conception of the world abounds with imprecision, and may turn out, upon analysis, to be thoroughly context-relative and to conflict radically with fundamental reality.⁶⁸

⁶⁶ Cf. Dorr 2007: 33 on fundamental and non-fundamental uses of an expression.

⁶⁷ Cf. Sider (2011: §7.9) on the advantages of semantic ascent in the context of grounding.

⁶⁸ Let me briefly outline a further argument against pure generation. I argued above (§4.2) that generation ought to be backed by a concise lawbook of general principles. But it is hard to see how such principles could back the generation of truths involving names and/or atomic predicates, unless supplemented by some auxiliary reductions. For example, by what suitable general principle could the existence of Obama be outputted? Perhaps there is a suitable principle governing the generation of mereological fusions; but this principle is only applicable once Obama has been 'located' within the

5.2 Against pure reduction

According to pure reducers, all metaphysical explanation is (or is backed by) reduction amongst truths.⁶⁹ The most direct way to challenge this view is to identify certain irreducible (i.e. perspicuous) truths which plausibly represent generated facts.⁷⁰ This section argues that, even allowing for the success of a breathtakingly ambitious reductive project — in which all truths are reduced to logically simple truths about ontologically simple entities — such cases are hard to avoid. In conceding so much to reducers, this argument departs significantly from the existing arguments for the importance of generation in the literature (such as those in Schaffer 2009, Rosen 2010, and Bennett 2017), which either do not consider the reduction approach or effectively assume that it is much more limited in scope.

Schaffer (2004) argues that some ‘macro-properties’, such as being an H₂O molecule, ‘serve as the ontological basis for linguistic truths’. (In his words: ‘Molecules aren’t merely manners of speaking.’) On this view, the truths about such properties are perspicuous, and yet correspond to higher-level facts which are generated from microphysical facts. Pure reducers must reject this picture. This requires faith in an ambitious enterprise: all ‘higher-level’ truths, stated both in ordinary vocabulary as well as that of the special sciences, will need to be shown to be reducible

mereological hierarchy. This location step amounts to a reduction of <Obama exists> in perspicuous mereological terms. (See Rosen 2010, Glazier 2016 and Wilsch 2016 for discussion of metaphysical laws requiring ‘auxiliary identities’.)

⁶⁹ As with pure generation, I am not here targeting a position that has (at least to my knowledge) been explicitly defended. My discussion aims to foreground the distinctive importance of generation.

⁷⁰ Jones (forthcoming) has recently argued against something resembling pure reduction: the thesis that all non-fundamental phenomena are representational (such as the truth of certain sentences/propositions). I do not regard pure reducers as committed to this thesis, however, since non-perspicuous truths needn’t be representational phenomena in the relevant sense.

to fundamental physics. Embarking on such an enterprise is certainly brave, but may not be foolish (for words of encouragement, see Sider 2011: §7.11.1).⁷¹

However, we needn't defend anything as anti-reductionist as Schaffer's picture to refute the pure reducer. Instead, we can grant that all truths are reducible to fundamental physics, and argue that some 'fundamental' physical facts are plausibly generated.

One remaining obstacle to pure reduction is logical complexity. If reduction requires factual equivalence, it is hard to see how all disjunctive truths, or all quantificational truths, could reduce to logically simple truths. But if there are logically complex facts — such as [Fa or Fb] and [There are Fs] — they would cry out for explanation in terms of logically simple facts — such as [Fa] — and this explanation is naturally taken up by the generation approach. Hence, pure reducers need an alternative conception of reduction which does not require factual equivalence.

Taking inspiration from Fine (2013), we might envisage a version of the reduction approach — the 'D-project' — which asks what in reality truths 'describe', as opposed to 'express' (see n.12).⁷² One way to think about the distinction between expression and description is in terms of a truth's possible truth-makers: a truth expresses the existential claim that one of its possible truth-makers obtains, and it describes each of the possible truth-makers which do obtain. The idea is then that logically complex truths may be taken to 'represent the same fact as' logically simple truths in the sense that they both describe the same fact. For example, <Fa or Fb> and <There are Fs> both describe [Fa], and hence can both be reduced to <Fa>, which perspicuously describes this same fact. This differs from the ordinary biconditional conception of reduction — the 'E-project' —

⁷¹ For discussion of the tension between reduction and multiple realizability, see Schaffer 2013 and Sider 2013c: §3. For reductionist progress, see Gómez Sánchez 2023.

⁷² Rather than proposing an expanded conception of reduction, Fine himself might be interpreted as insisting on the importance of generation in addition to reduction.

which would prohibit such reductions on the grounds that $\langle Fa \text{ or } Fb \rangle$ and $\langle \text{There are } Fs \rangle$ do not express $[Fa]$.⁷³

It is far from clear that any such conception of reduction is legitimate, however. One question is whether the D-project is really a version of the reduction approach, or is better thought of as some third approach (akin to truth-making, as described in §2.3). This threatens to be terminological. On one hand, the D-project bears important similarities to the E-project: they each target non-perspicuous truths, and connect them to perspicuous truths which ‘represent’ the same fact. On the other, they treat the key notion of representation (or ‘latching onto’) very differently. The substantive (albeit elusive) question, to my mind, is whether the D-project adequately addresses the demands of explanatory metaphysics. Suppose we had a complete account of what all truths describe: would this be a complete account of reality itself? Or would some aspects of what these truths are ‘getting at’ be left unexplained: those aspects of reality which they express but do not describe? This seems to me a very deep question; I wish I knew how to go about answering it.⁷⁴

Set this challenge aside, however: suppose that reality consists in particular details, non-perspicuously described by logically complex patterns.⁷⁵ A further obstacle remains: ontological complexity. The perspicuous truths seemingly involve composite entities of some sort, such as mereological fusions. For example, to completely describe spacetime, we plausibly need to specify the pattern of topological properties and/or distance relations. As Sider (2013a: §11) points out, the topological notion of ‘openness’ is a predicate of regions, and distances in General Relativity are path-dependent (the distance between two points is the length of the shortest path between them). So our perspicuous description plausibly includes sentences of the form ‘region R is open’ and ‘the distance along path P is x’. But the corresponding facts are not plausibly basic: facts

⁷³ More generally, pure reducers might look to forms of semantics ‘beyond the biconditional’, such as truthmaker semantics, which gives meanings in terms of ‘verifiers’ and ‘falsifiers’ (see e.g. Fine 2017, Elgin forthcoming).

⁷⁴ See Sider 2013b for discussion.

⁷⁵ This may require an atomic totality fact of the form ‘ $T(a, b, \dots)$ ’ to reduce generalizations (Russell 1918: 503; Fine 2012: 62). Pure reducers might regard (some) negative truths as perspicuous, with the corresponding negative facts ungenerated.

involving composite entities like regions and paths are plausibly generated by facts involving the sub-regions and sub-paths (and perhaps ultimately the points) which make them up.⁷⁶

Pure reducers face a dilemma here. They can either (i) deny these plausible generation claims and insist that regions and paths are fundamental in the same sense that points are, or (ii) provide alternative primitive notions which apply to points directly (rather than regions or paths), and which all truths about open regions and path-dependent distances reduce to.

Proponents of (ii) could, for example, take openness and distance to be captured by plural predicates applied to continuous infinities of points. For example, truths of the form ‘region R is open’ might reduce to infinitary truths of the form ‘points p, q, ... are open’. However, since it requires perspicuous plural predicates, this exchanges ontological complexity for a significant jump in ideological complexity, and — following Sider (2011: §9.15) — we may doubt that this sort of tradeoff is worthwhile.⁷⁷

Besides, the explanatory connections we were worried about survive the transition from ontological to ideological complexity. The openness of certain regions is determined by their being the union of certain other open regions, and the distances along paths are determined by the distances along various sub-paths, regardless of whether these facts involve monadic properties of composite entities or plural properties of their constituents. These pervasive patterns indicate explanatory connections, with the more ‘local’ truths (pertaining to smaller regions/paths) plausibly explaining the less local truths which they determine (pertaining to larger regions/paths). But these explanatory connections are better treated in terms of generation than reduction, since ‘locality’ seems irrelevant to perspicuity: it seems implausible to claim that $\langle p, \dots \text{ are open} \rangle$ and $\langle q, \dots \text{ are open} \rangle$ are perspicuous, whereas $\langle p, \dots, q, \dots \text{ are open} \rangle$ is not!⁷⁸

⁷⁶ Similarly, the richness of set theory suggests that mathematical reality cannot be completely described without using sets (or the like). And yet, facts involving sets would appear to be generated by facts involving their members.

⁷⁷ Sider is concerned with plural quantification, but analogous points apply to plural predication.

⁷⁸ The point here is not that there are necessary connections between the atomic truths which must be explained by invoking generation. Pure reducers may instead seek to explain necessary connections

Sider (2011: 146) points out that, at least in some cases, this sort of generation would not be well-founded: in a continuous space there are no smallest open regions and no shortest paths. But it is unclear that this is reason to deny the generation in question, especially in this context, where we are supposing that it merely supplements the (well-founded) explanatory structure of reduction. Methodologically speaking, we must decide whether general principles like well-foundedness should be prioritized over intuitions about particular cases. But besides, denying any generation amongst the atomic facts would preserve well-foundedness at the cost of violating another (to my mind, at least as plausible) general principle: that the fundamental facts ought to be non-redundant (i.e., free of necessary connections). If all openness facts are ungenerated, for example, then the basic facts involve pervasive redundancy.

Recognizing reduction and generation offers an attractive diagnosis of the conflict between well-foundedness and non-redundancy. An intuitive motivation for well-foundedness is that all truths have some ultimate basis in reality: this seems to fit reduction better than generation.⁷⁹ Some collection of geometric facts ought to serve as foundation for all the geometric truths, but positing some bottomless generative hierarchy within this foundation seems no more problematic than, in the causal case, a universe without initial conditions (cf. Fine 2001: 27). An intuitive motivation for non-redundancy is that necessary connections between contingent truths ought to be explicable, and that ground (either through their grounding one another, or sharing some common ground) provides a plausible explanation. This tells against necessary connections between basic geometric facts, but not between perspicuous geometric truths, since the latter connections may be explained

via some reduction of modality. The point is rather that there are general patterns amongst the atomic truths, of the kind which indicate explanatory connections (cf. Kim 1993: 167). This is consistent with the necessary connections explaining the generational connections, or something else explaining both — for example, following Sider 2011: §12.5 and Wilsch 2016 respectively, we might account for both necessity and generation in terms of ‘metaphysical laws’.

⁷⁹ Strictly speaking, it fits reduction as supplemented with the representation relation (and/or truth-making), which bridges the gap between facts and truths. As mentioned above (n.27), there may be no perfectly perspicuous truths. But we may still demand that all truths be supported by some foundation of facts, even if these facts have no perfectly perspicuous representation.

in terms of the generation relations between the corresponding facts. Hence, an appealing overall picture involves a highly redundant foundation of perspicuous truths, corresponding to a bottomless hierarchy of facts with no unexplained necessary connections.

In sum: even granting that all truths reduce to purely physical, logically atomic truths about ontologically simple entities, it still seems plausible that some of the corresponding facts are generated. Bennett (2017: 237) likens generative structure to ‘a growing plant, unfolding upwards and outwards’. My discussion suggests that this plant is a resilient one, surviving even the harshest desert landscape.

In characteristic style, reducers might countenance truths about generation whilst denying that reality itself features any generative structure. Indeed, this kind of position is defended by Sider (2011: 145): ‘we reductionists can take the same attitude towards metaphysical causation [i.e. generation] as towards everyday causation: it reduces in some way to facts that don’t involve metaphysical causation’. For example, Sider suggests, we might reduce generation to metaphysical laws (and reduce these laws in the style of the best system account). Thus, reducers might concede that facts involving logical, ontological, or ideological complexity are ‘generated’, but deny that this way of talking is perspicuous.

I suspect that generators would not regard this ‘Humean’ position as a genuine reconciliation: they typically view truths about generation as perspicuous (although they may not put it that way). Nonetheless, it does amount to a retreat from pure reduction, in favor of a hybrid approach to metaphysical explanation. First, Humeans treat reduction and generation alike, since they cannot plausibly hold that truths about reduction (which feature non-perspicuous truths as constituents) are themselves perspicuous. On this picture, all explanatory structure — generative and reductive — disappears from reality. Second, the crucial question as I see it is: what makes metaphysical explanations explanatory? Humeans recognize two sources of explanatoriness: generation amongst facts and reduction amongst truths. For Humeans — unlike staunch pure reducers — not all metaphysical explanation is tied to perspicuity.

6. Conclusion

I have laid out the distinction between two broad approaches to metaphysical explanation, and have demonstrated some of the ways in which it matters. I hope to have made it clear that whenever one is engaged in explanatory metaphysics, the question of approach should be front and center.

I have also suggested some ways of adjudicating between the two approaches, and argued that neither can sustain metaphysical explanation alone. This motivates a ‘hybrid’ approach which employs both. The hybrid approach recognizes two dimensions of relative fundamentality: perspicuity amongst truths, and basicness amongst facts. But it does not make explanatory metaphysics a disconnected enterprise. Rather, generation and reduction complement each other: reduction ‘lays bare’ the facts which generation targets, whilst generation provides the ‘raw material’ for reduction. For example, an explanation of water’s existence in terms of hydrogen’s existence might involve two components: water’s existence reduces to H₂O’s existence, and hydrogen’s existence generates H₂O’s existence.⁸⁰

The hybrid approach is a broad church, with room for disagreement about where to place the line between generation and reduction (and, correspondingly, their relative importance to the overall project). At one extreme is the view that our conception is more or less perspicuous: reduction is largely ‘tidying around the edges’, with most interesting action involving the formulation of ‘heavyweight’ principles of generation. At the other extreme is the view that our conception is radically non-perspicuous: the main task is to recover this conception from an unfamiliar reality, with all generation principles being ‘lightweight’ in character (purely logico-mathematical, perhaps).

Recognizing both approaches allows us to raise new questions and to sharpen some old questions. For example, many contemporary metaphysicians have considered whether fundamentality is fundamental, and how ground should be grounded. But generators are asking: ‘are facts involving

⁸⁰ In future work, I plan to develop a version of physicalism which employs both reduction and generation, and to extend this hybrid approach to causal and mathematical explanation.

generation and basicness basic, and if not, how are they generated?’, whereas reducers are asking the very different question: ‘are truths about reduction and perspicuity perspicuous, and if not, how are they reduced?’ We can also ask whether there are basic facts involving perspicuity, and whether truths about generation are perspicuous.

The answers to these questions are likely to mark significant divisions within the hybrid approach, and no doubt further important divisions remain to be made. Nonetheless, what unifies the hybrid approach — as I see it — is the recognition of two approaches to metaphysical explanation which are both different in important ways and, in different ways, important.⁸¹

REFERENCES

- Armstrong, David (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Armstrong, D. M. (2004). *Truth and truthmakers*. Cambridge: Cambridge University Press.
- Audi, Paul (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy* 109(12):685–711.
- Bacon, Andrew (2020). Logical Combinatorialism. *Philosophical Review* 129 (4):537-589.
- Barnes, Elizabeth (2014). Going Beyond the Fundamental: Feminism in Contemporary Metaphysics. *Proceedings of the Aristotelian Society* 114 (3pt3):335-351.
- Barnes, Elizabeth & J. Robert G. Williams (2011). A Theory of Metaphysical Indeterminacy. In Karen Bennett & Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics volume 6*. Oxford University Press. pp. 103-148.
- Baron, Sam & James Norton (2019). Metaphysical Explanation: The Kitcher Picture. *Erkenntnis*:1–21.
- Bennett, Jonathan (2001). *Learning from Six Philosophers: Descartes, Spinoza, Leibniz, Locke, Berkeley, Hume (Vol. 1)*. New York: Clarendon Press.

⁸¹ Thanks to Karen Bennett, David Builes, Alex Skiles, the Spring 2020 Rutgers Metaphysics Group, and two referees for *The Philosophical Review* and for *Noûs*. Special thanks to Verónica Gómez, Jonathan Schaffer, and Ted Sider for many discussions (and insights into the minds of generators and reducers) which have greatly benefitted this paper.

- Bennett, Karen (2004). Spatio-temporal coincidence and the grounding problem. *Philosophical Studies* 118(3):339–371.
- (2009). Composition, colocation, and metaontology. In David J. Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- (2011a). Construction area (no hard hat required). *Philosophical Studies* 154(1):79–104.
- (2011b). By Our Bootstraps. *Philosophical Perspectives* 25 (1):27–41.
- (2017). *Making Things Up*. Oxford: Oxford University Press.
- Builes, David (2021). Ontology and Arbitrariness. *Australasian Journal of Philosophy*. doi: [10.1080/00048402.2021.1931377](https://doi.org/10.1080/00048402.2021.1931377)
- Cameron, Ross P. (2008). Truthmakers and ontological commitment: or how to deal with complex objects and mathematical ontology without getting into trouble. *Philosophical Studies* 140 (1):1–18.
- (2010). How to have a radically minimal ontology. *Philosophical Studies* 151(2):249–264.
- Correia, Fabrice (2010). Grounding and truth-functions. *Logique Et Analyse* 53(211):251–279.
- Correia, Fabrice & Benjamin Schnieder (eds.) (2012). *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge: Cambridge University Press.
- Correia, Fabrice & Alexander Skiles (2017). Grounding, Essence, and Identity. *Philosophy and Phenomenological Research* (3):642–670.
- Dasgupta, Shamik (2017). Constitutive Explanation. *Philosophical Issues* 27(1):74–97.
- deRosset, Louis (2013). Grounding Explanations. *Philosophers' Imprint* 13.
- (2017). Grounding the Unreal. *Philosophy and Phenomenological Research* 95 (3):535–563.
- Dorr, Cian (2005). ‘What We Disagree About When We Disagree About Ontology’. In Mark Kalderon (ed.), *Fictionalist Approaches to Metaphysics*, Oxford: Oxford University Press.
- (2007). ‘There Are No Abstract Objects’. In Theodore Sider, John Hawthorne and Dean Zimmerman (eds.), *Contemporary Debates in Metaphysics*, Malden: Wiley-Blackwell.
- (2014). Quantifier Variance and the Collapse Theorems. *The Monist* 97 (4):503–570.
- (2016). To Be F Is To Be G. *Philosophical Perspectives* 30(1):39–134.

- Elgin, Samuel Z. (forthcoming). The semantic foundations of philosophical analysis. *Review of Symbolic Logic*:1-25.
- Fairchild, Maegan & John Hawthorne (2018). Against Conservatism in Metaphysics. *Royal Institute of Philosophy Supplement* 82:45-75.
- Fine, Kit (2001). The question of realism. *Philosophers' Imprint* 1:1–30.
- (2005). Tense and reality. In *Modality and Tense*. Oxford: Oxford University Press.
- (2010). Towards a Theory of Part. *Journal of Philosophy* 107 (11):559-589.
- (2012). Guide to Ground. In Correia & Schnieder (2012).
- (2013). Fundamental Truth and Fundamental Terms. *Philosophy and Phenomenological Research* 87(3):725–732.
- (2015). Unified Foundations for Essence and Ground. *Journal of the American Philosophical Association* 1(2):296–311.
- (2017). A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic* 46 (6):625-674.
- Field, Hartry (1994). Disquotational truth and factually defective discourse. *Philosophical Review* 103 (3):405-452.
- Fodor, Jerry (1990). *A Theory of Content and Other Essays*. Cambridge: MIT Press.
- Glazier, Martin (2016). Laws and the Completeness of the Fundamental. In Mark Jago (ed.), *Reality Making*. Oxford: Oxford University Press. pp. 11-37.
- (2017). Essentialist Explanation. *Philosophical Studies* 174(11):2871–2889.
- Gómez Sánchez, Verónica (2023). Naturalness by law. *Noûs* 57 (1):100-127.
- (forthcoming). From Nomic Humeanism to Normative Relativism. *Philosophical Perspectives*.
- Goodman, Jeremy. (2017). Reality Is Not Structured. *Analysis* 77, no. 1: 43–53.
- Hare, Caspar (2009). *On Myself, and Other, Less Important Subjects*. Princeton: Princeton University Press.
- Heil, John (2003). Levels of Reality. *Ratio* 16(3): 205–21. doi:10.1111/1467-9329.00218
- . Truthmaking and fundamentality. *Synthese* 198 (Suppl 3):849-860.
- Jenkins, Carrie (2011). Is Metaphysical Dependence Irreflexive? *The Monist* 94(2):267–276.

- Jones, Nicholas K. (forthcoming). Against Representational Levels. *Philosophical Perspectives*.
- Kennedy, Christopher & Stanley, Jason (2009). On ‘Average’. *Mind* 118(471):583–646.
- Kim, Jaegwon (1993). Postscripts on Supervenience. *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- Kment, Boris (2014). *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.
- Korman, Daniel Z. (2010). Strange Kinds, Familiar Kinds, and the Charge of Arbitrariness. *Oxford Studies in Metaphysics*:119-144.
- Krämer, Stephan & Roski, Stefan (2015). A Note on the Logic of Worldly Ground. *Thought: A Journal of Philosophy* 4(1):59–68.
- Kriegel, Uriah (2015). Existence, Fundamentality, and the Scope of Ontology. *Argumenta* 1 (1):97-109.
- Levine, Joseph (1983). Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64: 354–61.
- Lewis, David (1983a). New work for a theory of universals. *Australasian Journal of Philosophy* 61(4):343–377.
- (1983b). Counterpart Theory and Quantified Modal Logic. Reprinted with additional “Postscripts” in *Philosophical Papers, vol. 1*. Oxford: Oxford University Press.
- (1986). *Philosophical Papers, vol. 2*. Oxford: Oxford University Press.
- (1991). *Parts of Classes*. Oxford: Blackwell.
- (1994). Humean Supervenience Debugged. *Mind* 103:473–490.
- (1999). Many but Almost One. In *Papers in Metaphysics and Epistemology: Volume 2*. Cambridge: Cambridge University Press.
- (2001). Truth Making and Difference Making. *Noûs* 35: 602–15.
- Lipman, Martin A. (2016). Perspectival Variance and Worldly Fragmentation. *Australasian Journal of Philosophy* 94 (1):42-57.
- (2018). Against Fundamentality-Based Metaphysics. *Noûs* 52(3):587–610.
- Loewer, Barry (1996). Humean Supervenience. *Philosophical Topics* 24(1):101-127.
- Massimi, Michela (2017). Laws of Nature, Natural Properties, and the Robustly Best System. *The Monist* 100 (3):406-421.

- McGee, Vann & McLaughlin, Brian P. (2000). The Lessons of the Many. *Philosophical Topics* 28:129–152.
- Moore, George E. (1925). A Defence of Common Sense. In J. H. Muirhead (ed.) *Contemporary British Philosophy* (2nd series). London: Allen and Unwin.
- Mulligan, Kevin, Peter Simons & Barry Smith (1984). Truth-Makers. *Philosophy and Phenomenological Research*. 44: 287–321.
- (O’Leary-)Hawthorne, John & Andrew Cortens (1995). Towards ontological nihilism. *Philosophical Studies* 79(2):143–65.
- Rayo, Agustín (2013). *The Construction of Logical Space*. Oxford: Oxford University Press.
- Rosen, Gideon (2010). Metaphysical Dependence: Grounding and Reduction. In Bob Hale & Aviv Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.
- (2011). The Reality of Mathematical Objects. In John Polkinghorne (ed.), *Meaning in Mathematics*. Oxford: Oxford University Press.
- (2017). Ground by Law. *Philosophical Issues* 27(1):279–301.
- Russell, Bertrand (1918). The Philosophy of Logical Atomism. In R. C. Marsh (ed.), *Bertrand Russell: Logic and Knowledge, Essays 1901–1950*, London: Unwin Hyman, 1956.
- (2003). Analytic Realism. In S. Mumford (ed.), *Russell on Metaphysics*. London: Routledge, 2003.
- Russell, Jeffrey S. (2013). Possible Worlds and the Objective World. *Philosophy and Phenomenological Research* 87 (1).
- Schaffer, Jonathan. (2004). Two conceptions of sparse properties. *Pacific Philosophical Quarterly* 85:92–102.
- (2008). Truthmaker commitments. *Philosophical Studies* 141 (1):7-19.
- (2009). On what grounds what. In David Manley, David J. Chalmers & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- (2012). Grounding, transitivity, and contrastivity. In Correia & Schnieder (2012).
- (2013). Metaphysical Semantics Meets Multiple Realizability. *Analysis* 73(4):736–751.

- (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.
- (2021). Ground Functionalism. *Oxford Studies in the Philosophy of Mind*.
- Schulte, Peter (2011). Truthmakers: A Tale of Two Explanatory Projects. *Synthese* 181: 413–31.
- Shoemaker, Sydney (2007). *Physical Realization*. Oxford: Oxford University Press.
- Sider, Theodore (2001). *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- (2013a). Against Parthood. *Oxford Studies in Metaphysics* 8:237–293.
- (2013b). Replies to Dorr, Fine, and Hirsch. *Philosophy and Phenomenological Research* 87(3):733–754.
- (2013c). Symposium on Writing the Book of the World. *Analysis* 73(4):751–770.
- Spencer, Jack (2016). Relativity and Degrees of Relationality. *Philosophy and Phenomenological Research* 92 (2):432–459.
- Turner, Jason (2010). Ontological Pluralism. *Journal of Philosophy* 107(1):5–34.
- Van Inwagen, Peter (1990). *Material Beings*. Ithaca: Cornell University Press.
- Williams, J. Robert G. (2012). Requirements on Reality. In Correia & Schnieder (2012).
- Wilsch, Tobias (2016). The Deductive-Nomological Account of Metaphysical Explanation. *Australasian Journal of Philosophy* 94 (1):1–23.
- Wilson, Alastair (2018). Metaphysical Causation. *Noûs* 52 (4):723–751.
- Wilson, Jessica (2009). Determination, Realization, and Mental Causation. *Philosophical Studies* 145(1): 149–169. doi:10.1007/s11098-009-9384-8
- (2013). A determinable-based account of metaphysical indeterminacy. *Inquiry: An Interdisciplinary Journal of Philosophy* 56 (4):359–385
- Wittgenstein, Ludwig (1922). *Tractatus Logico-Philosophicus*. Trans. C. K. Ogden. London: Routledge and Kegan Paul Ltd., 1981.