

# RELEVANCE AND THE AIM OF FUNDAMENTAL METAPHYSICS

MARTIN GLAZIER

ABSTRACT. Fundamental metaphysics is the branch of metaphysics concerned with fundamental reality. I propose that fundamental metaphysics should aim, at least in part, to summarize the fundamental truths and to determine, for each non-fundamental truth, what is relevant to its explanation in the fundamental truths. I develop an account of the notion of grounding-explanatory relevance to which this proposal appeals.

Keywords: fundamentality, metametaphysics, ground, explanation, relevance

## 1. FUNDAMENTAL METAPHYSICS

What is the aim of fundamental metaphysics? Philosophers since Thales have pursued this project; if we complete it, what will we have made? Put another way, what is, or should be, the product or ‘deliverable’ of fundamental metaphysics? This paper proposes a partial answer to this question.

My question concerns fundamental metaphysics rather than metaphysics simpliciter. Metaphysics encompasses a range of domains, many of which are plausibly nonfundamental. A metaphysician, for instance, might investigate the nature of artifacts, or gender, or race, and these phenomena, important though they are, are presumably no part of fundamental reality.<sup>1</sup> Fundamental metaphysics is the branch of metaphysics concerned with fundamental reality.

But how should the notion of fundamental reality be understood? This question has been much discussed in recent years and I will not recapitulate that discussion here.<sup>2</sup> I will simply assume that we understand the notion well enough to proceed. I further assume that we have available the notion of a fundamental fact, a fact which partly constitutes fundamental reality—and thus that we have available the notion of a fundamental truth, a proposition which states a fundamental fact.

Within fundamental metaphysics we may distinguish two subbranches. Intrinsic fundamental metaphysics is concerned with fundamental reality as it is in itself, without regard to the rest of reality. Extrinsic fundamental metaphysics, by contrast, is concerned with the relation between fundamental reality and the rest of reality.

We began by asking after the aim of fundamental metaphysics. But we now see that what seemed to be one question is really two. What is the aim of intrinsic fundamental metaphysics? And what is the aim of extrinsic fundamental metaphysics? The answer to the more general question with which we began will consist in large part of answers to these two more specific questions. I will consider each in turn.

---

<sup>1</sup>On the nonfundamentality of gender, for instance, see Barnes (2014).

<sup>2</sup>For discussion of fundamentality see Fine (2001), Schaffer (2010), Jenkins (2011), Sider (2011), Dasgupta (2014b), Raven (2016) and Bennett (2017), among others.

## 2. INTRINSIC FUNDAMENTAL METAPHYSICS

What is or should be the aim, the product, the deliverable, of intrinsic fundamental metaphysics?

A natural first thought is that it should be a list of the fundamental truths—that, in this sense, intrinsic fundamental metaphysics should aim to write the book of the world.<sup>3</sup> But there are two reasons to demur from this proposal.

The first reason appeals to the idea that some fundamental truths are not worth knowing. Consider Rawls's (1971: 379) case of the person who spends her days counting the blades of grass in various lawns.<sup>4</sup> Or consider Dennett's case of the theoretical study of *chmess*, a game which is 'just like chess except that the king can move two squares in any direction, not one' (2006: 39–40). We have the strong sense that these activities are largely worthless as inquiries. But why are they worthless? Because, at least in part, the truths they aim to discover are truths that are not worth knowing.

The sense in which these truths are not worth knowing is a relatively objective one, not one that is sensitive to an individual inquirer's parochial circumstances.<sup>5</sup> There is, to be sure, a different sense that *is* sensitive in this way. Dennett imagines an academic discipline devoted to the study of *chmess*, and for its hapless practitioners, the exigencies of academia will render the truths of *chmess* very much worth knowing. But it is only given these professional circumstances that such truths are worth knowing. In a more objective sense, they are not.

Some fundamental truths may fare scarcely better. Consider two arbitrarily chosen electrons  $e_1$  and  $e_2$  in different galaxies. It may well be a fundamental truth that they are exactly  $10^{55}$  Planck lengths apart. But if so, it is a fundamental truth not worth knowing. The point, of course, does not depend on this specific example (or on its microreductionist flavor). On any reasonable view, fundamental reality will be complex and detailed enough for there to be plenty of truths which, though fundamental, are mere trivialities.

If intrinsic fundamental metaphysics aims to list the fundamental truths, then some of what it aims to know will not be worth knowing. And so it will not be a wholly worthwhile branch of inquiry. Of course this does not entail that it is without any value at all. But we metaphysicians should aim higher. We should aim at something other than writing the book of the world.

This conclusion may also be supported in a different way by appeal to the thought that the fundamental truths may not include a description of the *character* of fundamental reality. Consider, for example, David Lewis's (1986: ix) pithy formulation of his doctrine of Humean supervenience: 'All there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another.' This can be read as a description of the character of fundamental reality, a statement of what, in general, the fundamental is like. But the description itself does not state a local matter of fact, for it concerns the world as a whole. Nor does it state a matter of particular fact, for it is highly general. So the description, even if it is correct, does not itself state a fundamental truth.

<sup>3</sup>I borrow this phrase from the title of Sider (2011), but I do not mean to suggest that he endorses this claim.

<sup>4</sup>The case is anticipated by Stace (1944: 238).

<sup>5</sup>It may be sensitive to less parochial aspects of our collective situation, such as the fact that we are human beings living on Earth.

Yet surely formulations like Lewis's are at least part of what we should want from intrinsic fundamental metaphysics. The project should, if all goes well, reveal to us the character of fundamental reality. Since a description of that character may not itself state a fundamental truth, the aim of intrinsic fundamental metaphysics cannot be to list the fundamental truths.

Indeed, these considerations suggest an alternative conception of what that aim should be. The intrinsic fundamental metaphysician should aim, at least in part, not to list the fundamental truths but to summarize them. For example, if Lewis's statement of Humean supervenience is correct, then that statement itself, or something like it, will surely be part of any summary of the fundamental truths. The precise distance between the electrons  $e_1$  and  $e_2$ , by contrast, will not be.

But although not every fundamental truth will form part of the intrinsic fundamental metaphysician's 'output', that does not mean that the fundamental truths are wholly disconnected from her project. On the contrary, they are essentially bound up with it, since her aim is to summarize those truths.

But what makes something a summary of the fundamental truths? Although a detailed account of this notion is beyond the scope of this paper, I would like to sketch one proposal, which I advance only tentatively. (I do not claim that this proposal is the only possible one.) It repurposes a different idea associated with Lewis, the 'Humean' or 'best system' account of laws of nature (Lewis 1973: 73–4). The laws, on this account, are the contingent regularities or generalizations which appear as theorems in the best system (or systems, if there are ties). The best system is the true deductive system that best balances simple axiomatization with informativeness about the world.<sup>6</sup>

What is so special about the theorems of the best system? Why should the laws of nature be taken to be among them? These theorems are often characterized as *summaries* of the totality of facts about the world.<sup>7</sup> Although I would not wish to endorse the best system account of laws of nature, I agree that the theorems of the best system constitute a kind of summary of the world.

So if the intrinsic fundamental metaphysician's aim should be to summarize the fundamental truths, perhaps she can accomplish this by seeking the best system. Not, as the Humean does, the best system of the whole of reality, but rather the best system of the fundamental: the true deductive system that best balances simple axiomatization with informativeness *about fundamental reality*. The truths contained in the best fundamental system (or in any of the best fundamental systems, if there are ties) will constitute a summary of the fundamental truths and may therefore be taken to be, at least in part, the deliverable of intrinsic fundamental metaphysics.<sup>8</sup>

### 3. EXTRINSIC FUNDAMENTAL METAPHYSICS

I turn now to extrinsic fundamental metaphysics, which is concerned with the relation between fundamental reality and the rest of reality. What should its aim be?

<sup>6</sup>The literature on the best system is vast. In addition to Lewis's seminal (1973) presentation, important discussions include Armstrong (1983: 66–73), Lewis (1983), Carroll (1990: 197–206), Lewis (1994), Loewer (1996), and Loewer (2012). See Bhogal (2020) for an overview of the debate.

<sup>7</sup>See, for instance, Loewer (1996: 192), Miller (2015: 1314), and Loewer and Hüttemann (2022: 5).

<sup>8</sup>As Lewis (1983: 367) observed, unless the language of the best system is restricted in some way, the system will turn out to contain all truths whatsoever. To avoid this consequence, Lewis required the language of the best system to contain only predicates for what he called 'perfectly natural' properties and relations. We may impose a similar requirement: the language of the best fundamental system is to contain only predicates occurring in one or more fundamental truths.

Again a natural first thought presents itself. Extrinsic fundamental metaphysics, it might be proposed, should aim to identify, for each nonfundamental truth, which fundamental truths, if any, explain it. With many others, I assume that the relevant form of explanation here is grounding explanation.<sup>9</sup> And so the aim, on this proposal, will be to list the truths of the form ‘ $f_1, f_2, \dots$  ground  $g$ ’, where  $f_1, f_2, \dots$  are fundamental truths and  $g$  is a nonfundamental truth. In this sense, extrinsic fundamental metaphysics will aim to say what grounds what.<sup>10</sup>

But there are again two reasons to demur from this proposal. In the first place, some truths of the form ‘ $f_1, f_2, \dots$  ground  $g$ ’ will involve a  $g$  that is not worth knowing, and if some truth is not worth knowing, then which fundamental truths explain it is not worth knowing either. To make the point vivid, suppose that  $g$  is the truth about how many blades of grass there are in Central Park’s Great Lawn. Imagine that Rawls’s grass counter, having completed her enumeration of these blades of grass and thereby discovered  $g$ , sets to work in identifying the bazillions of subatomic particles constituting each blade and the precise distance relations in which they stand. Truths about these matters may well form part of the fundamental ground of  $g$ . But the grass counter’s new fundamental inquiry seems scarcely more worthwhile than her previous nonfundamental one.

In the second place, even if we restrict our attention to those nonfundamental truths that *are* worth knowing, it may not be worth knowing which fundamental truths explain them. Consider, for example, the truth that there are stars; presumably this is a nonfundamental truth worth knowing. But this truth may well be explained in part by certain fundamental truths about the bazillions of subatomic particles  $a_1, \dots, a_n$  that constitute the stars. And it does not seem worth knowing much of anything about particles  $a_1, \dots, a_n$  in particular.

We will therefore consider an alternative proposal concerning the aim of extrinsic fundamental metaphysics. For each nonfundamental truth worth knowing, the extrinsic fundamental metaphysician should aim, at least in part, to identify, not which fundamental truths explain it, but what is *relevant* to its explanation in the fundamental truths, if it has one.<sup>11</sup>

Below I offer an account of this particular form of explanatory relevance. But we can already get a sense of the concept and its application to the aim of metaphysics by returning to the truth that there are stars. This nonfundamental truth is partly explained, let us suppose, by certain fundamental truths about the bazillion subatomic particles  $a_1, \dots, a_n$  to the effect that their precise relation is thus-and-so. This explanatory connection, in all its gory particularity, does not seem much worth knowing. But why not? The natural answer is that it contains a huge amount of irrelevant detail. It does not matter to why there are stars that  $a_1, \dots, a_n$  in particular are precisely thus-and-so related. All that matters is that *some* huge collection of

<sup>9</sup>Fine (2001) is a classic discussion of fundamentality and ground. For simplicity I will make no distinction between ground and grounding explanation and will speak of them interchangeably; see Raven (2015: 326) and Glazier (2020: 128–9) for discussion of some of the issues here. I will work with a conception of ground on which it is full rather than partial (see Fine 2012: 50) and on which it relates propositions rather than facts (see Fine 2012: 46–7).

<sup>10</sup>I am again borrowing a turn of phrase, this time from the title of Schaffer (2009), without wishing to suggest that he would endorse this proposal.

<sup>11</sup>Perhaps she should aim at other things too, such as an account of the ‘metaphysical laws’ governing the relation between the fundamental and the nonfundamental; see Wilsch (2015a; 2015b), Glazier (2016) and Schaffer (2017) for discussion of such laws.

particles is *roughly* thus-and-so related. In what follows, I spell out this sense of ‘all that matters’—that is, of grounding-explanatory relevance.

As the example of the preceding paragraph shows, what explains, or partly explains, a proposition may not be in this sense relevant to why that proposition obtains. There is, of course, a sense in which whatever explains a proposition *must* be relevant to it. In that relatively undemanding sense, one proposition is relevant to another just in case the former plays at least some role in making it the case that the latter obtains (Dasgupta 2014a: 4). This undemanding sense of relevance is not the target of this paper.

#### 4. GROUNDING-EXPLANATORY RELEVANCE

The notion of explanatory relevance we require is not specific to the grounding of a nonfundamental truth in fundamental truths. It is a notion that applies to all cases of ground. We will not need to introduce any considerations specific to fundamentality.

Indeed, the notion of grounding-explanatory relevance is important quite apart from the question of the aim of fundamental metaphysics. For to know what is relevant to the grounding explanation of *A* is to know what matters, in a certain sense, to why *A*. And so even those philosophers whose focus lies elsewhere may find the following account of interest.

**4.1. The notion in question.** We begin with an everyday example of grounding-explanatory relevance. Suppose I receive the sum of \$100 in unemployment compensation. Responsible citizen that I am, I want to know how this affects my tax liability. According to US tax law, an individual’s tax liability depends in part on her gross income. So before I can figure the tax I owe, I need to determine whether the \$100 sum is part of my gross income. Poring over the law, I discover it provides that ‘gross income includes unemployment compensation’<sup>12</sup> and so the \$100 sum is indeed part of my gross income.

What grounds this last proposition? Note that the law lays down certain conditions on what is to count as gross income and the \$100 sum counts as gross income according to those conditions. It is not implausible to hold that the sum’s being part of my gross income is grounded in the law’s being what it is, together with the proposition that the sum is part of my income, together with the proposition that the sum is unemployment compensation.<sup>13</sup> And indeed it is not implausible to hold in general that if the law lays down certain conditions on what is to count as *F*, and if *x* counts as *F* according to those conditions, then those propositions together ground the proposition that *x* is *F*.

Examining the law further, however, I discover it also provides that ‘gross income means all income from whatever source derived’.<sup>14</sup> And so the proposition that the \$100 sum is part of my gross income may be taken to have a second ground in addition to the one specified in the preceding paragraph. It is grounded in the law’s being what it is together with the proposition that the sum is part of my income.

On this way of thinking, then, the proposition that the sum is part of my gross income has multiple overlapping grounds. Reflection on these grounds suggests that it does not actually matter to why the sum is part of my gross income that it is

<sup>12</sup>26 U.S.C. §85(a).

<sup>13</sup>Rosen (2017: 285–7) discusses some related legal applications of ground.

<sup>14</sup>26 U.S.C. §61(a).

unemployment compensation. All that matters is that it is income from some source or other. In this sense, the proposition that the sum is unemployment compensation is not explanatorily relevant.<sup>15</sup>

**4.2. Indispensable ground-parts.** But how should this notion of grounding-explanatory relevance be understood? The case discussed in §4.1 suggests a natural answer. The proposition that the sum is unemployment compensation does not matter to why it is part of my gross income. Why not? Because that proposition is an idle wheel, wholly dispensable to the explanation. To be sure, when put together with the law and the proposition that the sum is part of my income, it forms a ground for the proposition that the sum is part of my gross income. But the law and the proposition that the sum is part of my income already form a ground all on their own. The sum's being unemployment compensation is irrelevant.

The general thought is that relevance is a matter of being an indispensable part of some ground.<sup>16</sup> Let  $A$  and  $B$  be propositions. We say that  $B$  is an indispensable part of a ground of  $A$ , or an *indispensable ground-part* of  $A$  for short, if there is some set  $\Gamma$  of propositions such that (i)  $\Gamma$  grounds  $A$ , (ii)  $\Gamma$  contains  $B$ , and (iii)  $\Gamma \setminus \{B\}$  does not contain a ground of  $A$ . We then take  $B$  to be explanatorily relevant to  $A$  just in case  $B$  is an indispensable ground-part of  $A$ .

But this natural account faces a difficulty.<sup>17</sup> Let us say that something is *signal-colored* just in case it is red, yellow or green. And suppose the illuminated light on this traffic signal is red. Since disjunctive propositions are grounded in their true disjuncts, the proposition that the light is red grounds the proposition that it is signal-colored. But of course the light will be some shade of red in particular—crimson, say. Since an object's having some determinable property is grounded in its having a corresponding determinate property, the proposition that the light is crimson grounds the proposition that it is red. And so by transitivity, the proposition that the light is crimson grounds the proposition that it is signal-colored.

The light's being crimson is irrelevant to why it is signal-colored. The proposition that the light is crimson simply includes too much detail. It does not matter that the light is crimson in particular; all that matters is that it is some shade of red.

But the above account of relevance entails that the proposition that the light is crimson *is* relevant. For if we start with this proposition, and remove from it the proposition that the light is crimson (i.e. itself), we no longer have any propositions at all, and so we no longer have a *set* of propositions that includes a ground of the proposition that the light is signal-colored. The light's being crimson is an indispensable ground-part of its being signal-colored.

This difficulty stems from a general feature of the account: it entails that (full) grounds are always relevant. For if  $A$  is grounded in  $B$ , then  $\{B\} \setminus \{B\} = \emptyset$  does not

<sup>15</sup>I do not say that the text offers the only reasonable account, or even the account most likely to be correct, of the grounds of the sum's being part of my gross income. My aim in this subsection is only to help communicate my target notion of relevance, and so as long as the account in the text is coherent and somewhat plausible, that is sufficient for my purposes. Some similar cases, which might also serve the same aim, are given by Krämer and Roski (2017: 1196–8).

<sup>16</sup>Krämer and Roski (2017), building on Strevens (2008), offer an account in this vein of the closely related notion of a difference-making partial ground. They also employ the 'idle wheel' metaphor (1195).

<sup>17</sup>I take the following case from Krämer and Roski (2017).

contain a ground of  $A$  and so  $B$  is an indispensable ground-part of  $A$ .<sup>18</sup> To avoid this difficulty we need an account that allows a ground of a proposition to be irrelevant to its explanation.

4.3. **Abstraction.** One might attempt to develop such an account in counterfactual terms. I will consider an account inspired by Yablo's (1997) classic discussion of proportionality, but the points I make will have clear application to many other counterfactual accounts.

To understand the Yablo-inspired account, return to the signal-colored light. Here are three true propositions about it: it is crimson, it is red, and it is signal-colored. There is a sense in which these propositions are listed in order of their 'degree of detail': the first is a more detailed version of the second, which is in turn a more detailed version of the third. (We will not need to make this notion precise.)

Now had the light been red without being crimson, it would still have been signal-colored. In this sense, the light's being red is said to *screen off* its being crimson from its being signal-colored. And the same would appear to hold of *any* more detailed version of the proposition that the light is red, such as that the light is crimsonish (i.e. crimson or nearly so). Any such proposition will be screened off by the proposition that the light is red.

But it might be thought that there is no less detailed version of the latter which screens *it* off. To see why, consider some less detailed version; to fix ideas, let it be the proposition that the light is reddish (i.e. red or nearly so). Had that proposition obtained without the light's being red, the light would *not* have been signal-colored. And so it does not screen off the proposition that the light is red.

The proposition that the light is red, then, may appear to have the following special feature. It screens off every more detailed version of itself, but no less detailed version of itself screens *it* off. The Yablo-inspired thought is that explanatory relevance is a matter of possessing this special feature.

Generalizing, the account is this. We say that  $B_1$  screens off  $B_2$  from  $A$  just in case, had  $B_1$  been true without  $B_2$ 's being true,  $A$  would still have been true. We further say that  $B$  is *required* by  $A$  just in case no less detailed version of  $B$  screens off  $B$  from  $A$ ; and we say that  $B$  is *enough* for  $A$  just in case  $B$  screens off every more detailed version of  $B$  from  $A$ . The account takes  $B$  to be relevant to the explanation of  $A$  just in case it is both required by and enough for  $A$  (cf. Yablo 1997: 266–7).

This account is not without its appeal. But it faces a difficulty over necessity. If a proposition is necessary, then the claim that it is screened off by some other proposition is a counterfactual whose antecedent is impossible. The standard view of counterfactuals with impossible antecedents is that they are vacuously true. And so a necessary proposition will be screened off by any proposition whatsoever, including any of its less detailed versions. A necessary proposition, then, can never be required by another proposition. And so the account entails that it can never be relevant to any explanation.

But consider the proposition that  $5^2 = 25$ . It is necessary. The account therefore entails that it cannot be relevant to any explanation. And yet it is. It is relevant, for instance, to the explanation of why 25 is a square number.

<sup>18</sup>If some propositions are zero-grounded (Fine 2012: 47–8) then this statement must be qualified. I ignore this complication here.

I do not wish to claim that this objection cannot possibly be met. For example, one might endorse a nonstandard view of counterfactuals with impossible antecedents on which they avoid vacuity.<sup>19</sup> But while I would not wish to rule out the possibility that this or some other counterfactual account could be made to work, I would submit that this difficulty gives us good reason to look elsewhere first.

Even if the Yablo-inspired account is ultimately to be left aside, it does capture a certain insight. The reason the proposition that the light is crimson is irrelevant to why it is signal-colored is that it contains too much detail. What is relevant is just what remains after one ‘abstracts away’ this detail. By abstracting the right propositions in the right way, we arrive at what is truly relevant to the explanation.<sup>20</sup> But which propositions are the right ones? And what is the right way?

Let us begin with the first question. Some propositions are hopeless as a starting point for abstraction. If a proposition is not an indispensable ground-part of *A*, then no amount of abstraction should be expected to deliver relevance to the explanation of *A*. Take the proposition that the \$100 sum I received was unemployment compensation. It is not an indispensable ground-part of the proposition that the sum is part of my gross income, and so no matter how much detail is removed the result will be irrelevant to the explanation of this proposition.

But if we begin with the indispensable ground-parts of a proposition, then by abstracting in the right way we will arrive at what is relevant. The proposition that the light is crimson, for example, is an indispensable ground-part of the proposition that it is signal-colored. If we abstract from the former proposition in the right way, we will jettison the irrelevant information about the light’s particular shade of red and retain only the relevant information that it is red.

Each indispensable ground-part, then, is an appropriate starting point for abstraction (though see §4.4 for a qualification). But further starting points must also be added. To see why, notice that the crimson light is something of a special case. For the proposition that the light is crimson forms not just a part, but the whole, of a ground of the proposition that the light is signal-colored. It is a part of a ground, to be sure, but an improper one.

In general, however, this will not be the case. An indispensable ground-part of a proposition may be only a proper part of a ground. In such cases we may need to abstract multiple indispensable ground-parts taken together, rather than one in isolation, in order to arrive at what is explanatorily relevant.

Consider, for instance, the proposition that this bucket of paint is sufficient to coat this wall. It is grounded, let us suppose, in the proposition that the wall is twelve feet wide, together with the proposition that it is eight feet high, together with some propositions about the paint. Both propositions about the wall’s dimensions are clearly indispensable ground-parts: remove either and you no longer have a ground. But if we abstract either in isolation we will miss what is truly relevant: that the total *area* of the wall is small enough to be coated by the bucket. If you start only with a proposition about the wall’s width, then no matter how much detail you throw away you will not reach a proposition about the wall’s area, and similarly for

<sup>19</sup>It will not help to simply take counterfactuals with impossible antecedents to be vacuously false rather than vacuously true. For then a necessary proposition will not screen off any of its more detailed versions and so will never be enough for another proposition. This again delivers the conclusion that a necessary proposition can never be relevant to any explanation.

<sup>20</sup>Strevens (2008) has developed an account in this vein of causal-explanatory relevance. My discussion owes much to his work.

the wall's height. It is only by abstracting width and height together that we obtain the relevant proposition about area.

The right starting points for abstraction, then, include not only each indispensable ground-part on its own, but also each plurality of indispensable ground-parts which are part of a common ground.

So much for the question of which propositions are the right ones to abstract. What about the question of the right way to abstract these propositions? Put another way, how exactly should abstraction be understood?

One might think to understand it in terms of modal entailment: one proposition is an abstraction of another just in case the second modally entails the first. But this will not work. By abstracting the indispensable ground-parts of  $A$ , we are supposed to obtain what is relevant to the explanation of  $A$ . But an indispensable ground-part of  $A$  will modally entail many propositions that are irrelevant to  $A$ . Consider again the proposition that the light is crimson; it is an indispensable ground-part of the proposition that the light is signal-colored. That the light is crimson modally entails that the light is crimson or zucchini is tasty, but the tastiness of zucchini is irrelevant to the light's being signal-colored, and so this disjunctive proposition is irrelevant as well. Or again, that the light is crimson modally entails that  $2 + 2 = 4$ , yet this too is irrelevant.

These difficulties stem from modal entailment's well-known disregard for any constraint of relevance. To avoid difficulty we require a notion of entailment that is sensitive to such constraints. Although there may be suitable alternatives, I will work with the notion of exact entailment (Fine and Jago 2019). One proposition exactly entails another just in case every state which is a truthmaker for the first is a truthmaker for the second. Since a state must be wholly relevant to any proposition for which it is a truthmaker, exact entailment is sensitive to relevance.

This form of entailment can deliver the reduction in detail which is abstraction's *raison d'être*. A proposition's truthmakers may be thought of as the ways for that proposition to be true. If one proposition exactly entails another, then every way for the first to be true is a way for the second to be true, and so there is a sense in which the second can be no more detailed than the first.

But we cannot simply take abstraction to be exact entailment. To arrive at what is relevant to the explanation of a given 'target' proposition, the process of abstraction must not only remove detail from an indispensable ground-part, it must remove it in the right way. If we are not careful we can break the explanatory link with the target.

One way this can happen is when the result of abstraction fails to guarantee the target. For example, the light's being red exactly entails its being reddish. But the light's being reddish is not relevant to why it is signal-colored; indeed it does not even guarantee that it *is* signal-colored. To block such cases we need to ensure that, if we start with some ground  $\Gamma$  of the target, and we abstract one or more indispensable parts of that ground, then the result of this abstraction, together with the other, 'remainder' propositions in  $\Gamma$  (if there are any), guarantees that the target will obtain.

Such guaranteeing, however, cannot be understood in terms of modal entailment. To see why, consider again the proposition that 25 is a square number. One of its indispensable ground-parts is the proposition that  $5^2 = 25$ . This proposition exactly entails the disjunctive proposition that  $5^2 = 25$  or zucchini is tasty, which in turn modally entails that 25 is a square number. But the disjunctive proposition is not relevant to why 25 is a square number.

Guaranteeing should instead be understood in terms of exact entailment. We should require that the abstraction of one or more indispensable parts of a ground of the target, together with the remainder propositions (if there are any), exactly entails the target.<sup>21</sup> In the special case in which the abstracted proposition grounds the target all on its own, this reduces to the requirement that the abstraction exactly entails the target. The reddish case is thereby blocked, for the proposition that the light is reddish has as one of its truthmakers a state in which the light is orange rather than red, and this state is not a truthmaker for the proposition that the light is signal-colored.

But though we have now solved one problem, others remain. A second way to break the explanatory link between indispensable ground-parts and target can arise when abstraction results in a proposition that is too ‘close’ to the target. Suppose we take as our target the proposition that the light is red. It has as an indispensable ground-part the proposition that the light is crimson, and so by abstracting this proposition we should be able to obtain what is relevant. Now the proposition that the light is crimson exactly entails the proposition that the light is red. And the latter proposition exactly entails the target (i.e., itself) and so guarantees it. Still, we do not want to say that what matters to why the light is red is that it is red.

To block such cases we should require an abstraction to be a certain distance from the target in the following sense. The abstraction of some indispensable ground-parts from a ground  $\Gamma$  of the target must, together with the remainder propositions (if there are any), exactly entail something that the target does not. In the special case in which the abstracted proposition on its own grounds the target, this reduces to the requirement that the abstraction exactly entails something that the target does not and hence that the target does not exactly entail the abstraction. This prevents the proposition that the light is red from being relevant to why the light is red.

**4.4. Cohesion.** Abstraction can go awry in a third way by removing too little detail. The proposition that the light is crimson is an indispensable ground-part of the proposition that the light is signal-colored. And the former proposition exactly entails the proposition that it is crimsonish. And *that* proposition exactly entails and so guarantees that it is signal-colored. *And* the proposition that the light is signal-colored does not exactly entail the proposition that it is crimsonish. So all the requirements we have imposed thus far are satisfied by the proposition that the light is crimsonish. All the same, that proposition is not relevant to the explanation of why the light is signal-colored. For it does not matter that the light is crimsonish in particular; all that matters is that it is red. To say that it is crimsonish is to include too much detail.

What this case shows is that too little abstraction can leave us with something irrelevant. We need to abstract ‘past’ the proposition that the light is crimsonish to arrive at the proposition that it is red. To get at what is relevant, we should abstract the indispensable ground-parts as much as possible.

As much as possible, but no more. For one might continue to abstract the proposition that this light is red to obtain the proposition that it is red or green. But surely what matters to why the light is signal-colored is that it is red, not that it is red or green. So where do we draw the line? How much abstraction is too much?

<sup>21</sup>Fine and Jago (2019) distinguish ‘collective’ and ‘distributive’ notions of exact entailment, which coincide for single-premise entailments but come apart for multiple-premise entailments. I must leave for another time the question of which notion, if either, is to be preferred.

The problem with the proposition that the light is red or green is that it lacks what Strevens (2008: 101–9) calls *cohesion*. Rather than pinpointing a single explanatory factor, it teeters between two of them. Such a proposition can never be relevant.

We must therefore qualify §4.3's claim that each indispensable ground-part is an appropriate starting point for abstraction. For if an indispensable ground-part lacks cohesion, then neither it nor any proposition obtained from it by abstraction will be explanatorily relevant. Only cohesive indispensable ground-parts are appropriate starting points.

It should not be assumed that disjunctive propositions inevitably lack cohesion. The UN Convention on the Law of the Sea provides that states bordering a sea 'connected to another sea or the ocean by a narrow outlet' must coordinate the exploitation of the sea's resources (Art. 123). The Baltic Sea is connected to the North Sea by the Danish straits and so the states bordering the Baltic are legally obligated to coordinate its exploitation. It is not implausible to suppose that this obligation is grounded in the Baltic's connection to the North Sea together with propositions about the law. But it does not matter to the explanation of this obligation that the connection is to the North Sea in particular. Nor does it matter that it is to a sea. For the category of sea is continuous with the category of ocean and either is sufficient to trigger the convention. What matters is simply that the connection is to a sea or ocean.

Cohesion, then, should not be taken to be a matter of whether a proposition is disjunctive or not. Instead, we should take a proposition to be cohesive just in case its truthmakers form a connected set.

But what is it for a set of truthmakers to be connected? Here we may follow the topologists, who have defined various notions of connectedness, such as path-connectedness, simple connectedness, contractibility, and so on. A set of states will be path-connected, for example, just in case one can continuously move between any two states in the set without leaving the set. (I leave for another time a comparison of the different notions of cohesion that correspond to different notions of connectedness.) To be sure, the definitions of the various notions of connectedness all presuppose an underlying topology on the space of all states. But if we make the plausible assumption that states stand in relations of objective similarity to one another (cf. Sider (2011: 88–90) on similarity relations between facts), it can be argued that the space of states does indeed admit of a suitable topology.<sup>22</sup>

So where do we draw the line between too little abstraction and too much? The answer is that we should abstract as much as possible while still retaining cohesion. Thus the proposition that this light is crimson may be abstracted to the proposition that it is red, since presumably the truthmakers of 'the light is red' form a connected set. But the proposition that this light is crimson cannot be abstracted to the proposition that it is red or green, since presumably the truthmakers of 'the light is red or green' do not form a connected set.

<sup>22</sup>Here is how. Presumably we should admit not just one state-similarity relation but a family of (binary) relations of state-similarity in this or that respect or to this or that degree. Each such relation is clearly reflexive and symmetric, and it is plausible that given one such relation  $S$  there will be some other such relation  $T$  which is strict enough that any two-link chain of  $T$ -similarity relations corresponds to an  $S$ -similarity relation (i.e. if  $a$  is  $T$ -similar to  $b$  and  $b$  is  $T$ -similar to  $c$  then  $a$  is  $S$ -similar to  $c$ ). Given these assumptions, by Theorem 6.3 of Kelley (1955) the family of state-similarity relations will form a subbase for some uniformity for the space of states, which induces a uniform topology on that space.

The directive to abstract as much as possible while retaining cohesion must be properly understood. For there are cases in which the directive cannot be obeyed, and if it *cannot* be obeyed, then it *should* not be obeyed either. In such cases we should avoid any abstraction at all.

Consider, for example, the Burj Khalifa skyscraper.<sup>23</sup> The proposition that it is more than 800 meters tall is grounded in the proposition that it is 829.8 meters tall (call this  $B$ ), and so the latter is an indispensable ground-part of the former. The proposition that the building is 829.8 meters tall may be abstracted to the proposition that it is more than 829 meters tall (call this  $B_{829}$ ). This, however, is no stopping point. For we may abstract further to reach the proposition that the building is more than 828 meters tall ( $B_{828}$ ), and then the proposition that it is more than 827 meters tall ( $B_{827}$ ), and so on. We must, of course, stop before we reach the proposition that it is more than 800 meters tall, for this is just the target of our explanation (and any proposition ‘beyond’ this would no longer guarantee the target). But wherever we stop we could have gone further. We cannot abstract as much as possible.

In this sort of case, the right response is not to abstract at all. For we do not have the sense that, in moving from  $B_{829}$  to  $B_{828}$  to  $B_{827}$ , we are getting closer to what really matters to why the building is more than 800 meters tall. There is no apparent gain in relevance over our starting point of  $B$ , and so we should never have quit that point in the first place.

The upshot of this section and the preceding one is the following. Given a plurality  $\Delta$  of one or more propositions that are members of a common ground  $\Gamma$  for a target  $A$ , the abstraction  $B$  of  $\Delta$ , if one exists, is the weakest (in the sense of exact entailment) cohesive proposition such that (i)  $\Delta$  exactly entails  $B$ , (ii)  $B$  together with the remainder propositions in  $\Gamma$  exactly entails  $A$ , and (iii)  $B$  together with the remainder propositions exactly entails some proposition not entailed by  $A$ .

**4.5. The final step.** I earlier said that we obtain what is relevant to the explanation of  $A$  by abstracting each cohesive indispensable ground-part (or plurality of cohesive indispensable ground-parts) of  $A$ . But even when this has been done we require one final step.

To see why, return to the wall case. We assume that the proposition that the bucket of paint is sufficient to coat the wall is grounded in the proposition that the wall is twelve feet wide, together with the proposition that it is eight feet high, together with some propositions about the paint (which I will ignore). Both propositions about the wall’s dimensions are indispensable ground-parts. Let us suppose that the bucket of paint is sufficient to coat any surface whose area is under 120 square feet. Then if we abstract only the proposition about the wall’s width, we will obtain the proposition that the wall is under fifteen feet wide. If we abstract only the proposition about its height, we will obtain the proposition that it is under ten feet high. If we abstract both together, we will obtain the proposition that the wall’s area is under 120 square feet. But it is only this last proposition that is truly relevant to why the paint is sufficient to coat this wall.

How can we ensure the account delivers this result? Notice that the proposition that the wall is under fifteen feet wide (call this  $A$ ), together with its being eight feet high, exactly entails the proposition that its area is under 120 square feet (call this  $B$ ). This last proposition is also exactly entailed by the proposition that the wall is under ten feet high ( $C$ ) together with its being twelve feet wide. These entailments

<sup>23</sup>I am indebted to Louis deRosset here.

reflect the fact that the propositions  $A$  and  $C$  contain irrelevant detail that  $B$  does not.

This suggests that we should delete from our list of relevant propositions any proposition that yields in this way some other proposition on the list. That is, let  $A_1, \dots, A_n$  and  $B_1, \dots, B_m$  be distinct pluralities of indispensable ground-parts all of which belong to the same ground. Let  $A$  and  $B$  be their abstractions and  $\Gamma$  and  $\Delta$  their remainders. If  $\{A\} \cup \Gamma$  exactly entails each member of  $\{B\} \cup \Delta$ , but the reverse does not hold, then we delete  $A$  from our list. In this way, we ensure that our final list of relevant facts is as abstract as possible.

Here, then, is our procedure for determining what is relevant to the explanation of  $A$ :

- (1) Find the cohesive indispensable ground-parts of  $A$ . That is, for each plurality of propositions which together ground  $A$ , determine which of its cohesive members are such that the remaining members do not contain a ground of  $A$ .
- (2) For each plurality of cohesive indispensable ground-parts that are members of a common ground, write down its abstraction, if it has one; otherwise write down the plurality itself. Each abstraction or plurality written down is said to be an *item*.
- (3) Erase any item which, together with its remainders, exactly entails some other item and each of *its* remainders, provided the reverse entailment does not hold, and provided both items were obtained from the same ground of  $A$ .

The result is the list of relevant propositions.

This account captures the natural and intuitive thought, which has been a theme throughout the discussion, that relevance is what remains when superfluous detail has been eliminated from all explanatory factors. For the account may be understood in the following way. To obtain what is relevant, we start with the collection of all propositions. We junk those that lack cohesion, for these can never succeed in pinpointing a single explanatory factor. We junk those that are not indispensable ground-parts of the target, for these are wholly superfluous. We take each remaining explanatory factor and junk as much detail as possible from it. And then we junk any of the resulting explanatory factors which strictly contains the detail of another, for the extra detail of the first is superfluous. Since this procedure plausibly succeeds in eliminating all superfluous detail, we may have some substantial measure of confidence that no counterexample to the account will emerge.

I would finally like to note a further nice feature of the account. It does not require a proposition relevant to the explanation of  $A$  to be a ground of  $A$  or even part of a ground of  $A$ . And this is welcome, for there may be cases in which such a requirement would be inappropriate.

To see why, return to the case of the light. The proposition that it is crimson grounds the proposition that it is red, which in turn grounds the existential proposition that *something* is red. But what is relevant to the explanation of this last proposition? Not the proposition that the light is crimson, for it does not matter that it is that light in particular that is crimson. Plausibly all that matters is that *something* is crimson. But it is sometimes assumed that existential propositions can be grounded only ‘by way of’ their instances. In the present case, this amounts to the assumption that any ground of the proposition that something is red must either be a proposition of the form ‘ $x$  is red’ or must ground a proposition of that

form.<sup>24</sup> Given this assumption, the proposition that something is red cannot be grounded in the proposition that something is crimson, and nor will the proposition that something is crimson be a part of some other ground of the proposition that something is red. Yet it still seems that the proposition that something is crimson is highly relevant to why something is red.

## 5. CONCLUSION

Let us now return to the proposal of §3 and see how §4's account of grounding-explanatory relevance fits with it. The proposal was that the extrinsic fundamental metaphysician should aim, at least in part, to identify, for each nonfundamental truth worth knowing, what is relevant to its explanation in the fundamental truths. To do this, for each nonfundamental truth *A* worth knowing, she should follow §4.5's procedure for determining what is relevant to its explanation—but with one slight modification. In the first step, she should ignore those indispensable ground-parts of *A* that are not fundamental. For her aim is not to determine everything relevant to the explanation of *A* but only what is relevant to its explanation in the fundamental truths.

Our partial account of the aim of fundamental metaphysics may be generalized. A closely related account may be given of the aim of any branch of metaphysics concerned with a domain that reduces to another domain in the sense that all truths of the first are grounded in truths of the second. For example, suppose one takes the moral domain to reduce to the descriptive domain in the sense that all moral truths are grounded in descriptive truths. The metaphysics of morality may then be divided into intrinsic and extrinsic subbranches. The aim of the intrinsic subbranch will be, at least in part, to produce a good summary of the moral truths, perhaps through discovering the Lewisian best system of the moral domain. And the aim of the extrinsic subbranch will be, at least in part, to identify, for each moral truth worth knowing, what is relevant to its explanation in the descriptive truths. Or again, suppose one takes the mental domain to reduce to the physical domain in the sense that all mental truths are grounded in physical truths. Then the aim of the intrinsic metaphysics of mind will be, at least in part, to produce a good summary of the mental truths, while that of the extrinsic metaphysics of mind will be, at least in part, to identify, for each mental truth worth knowing, what is relevant to its explanation in the physical truths.<sup>25</sup>

## REFERENCES

- Armstrong, D. M. 1983. *What is a Law of Nature?* Cambridge, England: Cambridge University Press.
- Barnes, E. 2014. Going beyond the fundamental: Feminism in contemporary metaphysics. *Proceedings of the Aristotelian Society* 114(3): 335–351.
- Bennett, K. 2017. *Making Things Up*. Oxford: Oxford University Press.
- Bhagal, H. 2020. Humeanism about laws of nature. *Philosophy Compass* 15(8).
- Carroll, J. 1990. The Humean tradition. *Philosophical Review* 99(2): 185–219.

<sup>24</sup>Fine (2012: 63–6) states the assumption in a weaker form, but the differences between Fine's formulation and ours do not matter here.

<sup>25</sup>My thanks to Singa Behrens, Louis deRosset, Stephan Krämer, Stephan Leuenberger, Jon Litland, Stefan Roski, Benjamin Schnieder, Alex Skiles, Michael Strevens, and Rivka Weinberg, as well as to audiences in Hamburg, Padua, and Vienna. I am grateful for the support of the German Research Foundation (grant KR 4516/2-1).

- Dasgupta, S. 2014a. On the plurality of grounds. *Philosophers' Imprint* 14(20): 1–28.
- Dasgupta, S. 2014b. The possibility of physicalism. *Journal of Philosophy* 111(9): 557–592.
- Dennett, D. C. 2006. Higher-order truths about chess. *Topoi* 25: 39–41.
- Fine, K. 2001. The question of realism. *Philosophers' Imprint* 1(1): 1–30.
- Fine, K. 2012. Guide to ground. In *Metaphysical Grounding: Understanding the Structure of Reality*, eds. F. Correia and B. Schnieder, 37–80. Cambridge, England: Cambridge University Press.
- Fine, K. and M. Jago. 2019. Logic for exact entailment. *Review of Symbolic Logic* 12(3): 536–556.
- Glazier, M. 2016. Laws and the completeness of the fundamental. In *Reality Making*, ed. M. Jago, 11–37. Oxford: Oxford University Press.
- Glazier, M. 2020. Explanation. In *The Routledge Handbook of Metaphysical Grounding*, ed. M. J. Raven, 121–132. London: Routledge.
- Jenkins, C. 2011. Explanation and fundamentality. In *Ontological Dependence, Supervenience, and Response-Dependence*, eds. B. Schnieder, A. Steinberg, and M. Hoeltje. Munich: Philosophia Verlag.
- Kelley, J. L. 1955. *General Topology*. New York: Springer.
- Krämer, S. and S. Roski. 2017. Difference-making grounds. *Philosophical Studies* 174(5): 1191–1215.
- Lewis, D. 1973. *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. 1983. New work for a theory of universals. *Australasian Journal of Philosophy* 61(4): 343–377.
- Lewis, D. 1986. *Philosophical Papers*, Volume 2. Oxford: Oxford University Press.
- Lewis, D. 1994. Humean supervenience debugged. *Mind* 103(412): 473–490.
- Loew, C. and A. Hüttemann. 2022. Are we free to make the laws? *Synthese* 4: 1–16.
- Loewer, B. 1996. Humean supervenience. *Philosophical Topics* 24: 101–127.
- Loewer, B. 2012. Two accounts of laws and time. *Philosophical Studies* 160(1): 115–137.
- Miller, E. 2015. Humean scientific explanation. *Philosophical Studies* 172: 1311–1332.
- Raven, M. J. 2015. Ground. *Philosophy Compass* 10(5): 322–333.
- Raven, M. J. 2016. Fundamentality without foundations. *Philosophy and Phenomenological Research* 93(3): 607–626.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, Mass.: Harvard University Press.
- Rosen, G. 2017. Ground by law. *Philosophical Issues*: 279–301.
- Schaffer, J. 2009. On what grounds what. In *Metametaphysics: New Essays on the Foundations of Ontology*, eds. D. J. Chalmers, D. Manley, and R. Wasserman, 347–383. Oxford: Oxford University Press.
- Schaffer, J. 2010. Monism: The priority of the whole. *Philosophical Review* 119(1): 31–76.
- Schaffer, J. 2017. Laws for metaphysical explanation. *Philosophical Issues* 27: 302–321.
- Sider, T. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Stace, W. T. 1944. Interestingness. *Philosophy* 19: 233–41.
- Strevens, M. 2008. *Depth: An Account of Scientific Explanation*. Cambridge, Mass.: Harvard University Press.

- Wilsch, T. 2015a. The deductive-nomological account of metaphysical explanation. *Australasian Journal of Philosophy* 94(1): 1–23.
- Wilsch, T. 2015b. The nomological account of ground. *Philosophical Studies* 172: 3293–3312.
- Yablo, S. 1997. Wide causation. *Philosophical Perspectives* 11: 251–281.