The Monist, 2023, **106**, 239–254
https://doi.org/10.1093/monist/onad012
Article

An Epistemic Approach to Ground

Ned Hall*

Recent enthusiasm for grounding often begins by observing that inquiry in metaphysics (and other areas) features a distinctive species of noncausal explanation. Having labeled this species "grounding explanation," it's a short step to the conclusion that we need a philosophical theory of grounding itself: an allegedly fundamental relation of metaphysical dependency between facts, such that a "grounding explanation" of some fact succeeds by providing information about what "grounds" that fact. This short step is hasty. For another live option is to accept that grounding explanation is a legitimate form of explanation, but to give it a thoroughly epistemic treatment, one that does not see it as involving any sort of special metaphysical relationship or structure at all. This paper sketches such a treatment, drawing inspiration from reflections on explanatory structure in mathematics.

0. INTRODUCTION

If we want to become grounding partisans, there is a prominent and perfectly Plausible Path we can follow. The Path starts with examples of simple explanations—examples that are by now canonical. It adds philosophical commentary designed to show that they are *grounding* explanations. It then zooms out to take stock of metaphysical inquiry as a whole, arguing that the characteristic questions that metaphysics takes up aim at the *very same kind* of explanation: they are, at bottom, questions about *what grounds what*. (See Schaffer [2009] for a bracing defense of this "what grounds what" conception of metaphysical inquiry.) Its final short step concludes that grounding *itself*—its logic, its properties, the extent of its coverage, etc.—should be a central topic *in metaphysics*. The recent literature has apparently obliged: as of this writing, the philpapers.org page on "grounding" lists 477 works, just two of which were published before the year 2000, and only 17 of which were published before 2010.[1]

The Plausible Path has had some persuasive force. (See for example deRosset [2020] for an especially crisp presentation and defense of it. For that matter, just ask a grounding enthusiast whence their enthusiasm, and you'll likely get some version of it.) But its final step looks hasty. Maybe metaphysics is up to its ears in grounding claims. Maybe these claims are *true*—some of them, anyway; and if not, not because the notion of "ground" is confused or incoherent. Maybe they all draw on the *same* notion of "ground"; i.e., this term isn't just a placeholder for some more specific explanatory metaphysical relationship that can vary from case to case. (Compare Wilson [2014].) Still, for all that it might be that the nature of grounding is purely a topic in

*Harvard University, US

epistemology; that when one fact P is grounded in some other facts $Q_1, \ldots, Q_n$, *that* fact itself holds purely in virtue of a certain kind of *epistemic* relationship that $Q_1, \ldots, Q_n$ bear to P. (And likewise *that* fact, and so on; in case you're wondering, I'm fully aware that the last sentence is itself a grounding claim. See §6 below.) The aim of this essay is to sketch such an epistemic approach to ground, by drawing inspiration from some reflections on *mathematical* inquiry.

Is the epistemic approach the *right* approach? I don't know. (Nor does anyone else, unless they possess powers of philosophical revelation denied the rest of us.) But it passes the right test: an important corrective if true; and if false, worth the effort of exposing as such. Consider this essay, then, an invitation to develop the epistemic approach in much more detail than I will be able to. Our sketch will begin by tracing the Plausible Path in a little more detail.

1. THE PLAUSIBLE PATH

1.1 The canonical examples

Billy and Suzy are having lunch. So: They are either having lunch or watching a movie together. Let us grant for the sake of discussion that corresponding to each of these true statements is a *fact*:

(Fact 1): The fact that Billy and Suzy are having lunch.
(Fact 2): The fact that Billy and Suzy are either having lunch or watching a movie together.

Now we get the first of our canonical examples (Compare Rosen 2010, p. 117.). For all of the following seem clearly true, and indeed seem to say the same thing: Fact 2 obtains *in virtue of* Fact 1. Fact 1 *explains* Fact 2. Fact 2 obtains *because* Fact 1 obtains. And, importantly, the reverse claims seem clearly false.

A natural first thought is that these claims are nothing more than stylized ways of reporting on an obvious logical asymmetry: Fact 1 entails Fact 2, but not vice versa. Could that be all that is going on? Plausibly not. For suppose that Suzy and Billy's friend Ahmed is watching a movie. Then we have two more facts (the second of our canonical examples):

(Fact 3): The fact that Ahmed is watching a movie.
(Fact 4): The fact that Billy and Suzy are having lunch and Ahmed is watching a movie.

These additional claims also seem clearly true, and indeed seem to say the same thing: Fact 4 obtains *in virtue of* Fact 1 and Fact 3. Fact 1 and Fact 3 together *explain* Fact 4. Fact 4 obtains *because* Fact 1 and Fact 3 obtain. Again, the reverse claims all seem clearly false. But the logical asymmetry goes the wrong way: Fact 4 entails Facts 1 and 3, not the other way around.

Let us grant all these claims, and take them at face value. Focus on those that use "because." The sense of this term looks distinctively different from the sense it has in standard causal-explanatory claims: To see the contrast, compare (for example) "Fact 2 obtains because Suzy invited Billy to lunch earlier in the day." But the "because" in "Fact 2 obtains because Fact 1 obtains" still looks *explanatory*, as does the "because" in "Fact 4 obtains because Fact 1 and Fact 3 obtain." Let us take the further step of agreeing that it *is* explanatory, that these claims present a distinctive kind of *noncausal explanation*.

We have taken the initial steps along the Plausible Path, and reached a point where we have identified a particular kind of explanation—a kind, note, that can at least sometimes (as in our canonical examples) be discovered a priori. How tempting it is, then, to take the next step, and see these "grounding" explanations cropping up all over philosophy.[2] To take just a few examples, perhaps the utilitarian is giving a grounding explanation of moral rightness when she says that morally right actions are those that maximize happiness: more carefully put, her claim should be that morally right actions are morally right *because* they maximize happiness, with this "because" having the same sense that it does in our canonical examples. Similarly, perhaps compatibilism should be understood as claiming that free actions are free *because* they have such-and-such psychological causes. Perhaps something has a mind *because* part of it (e.g., its brain) has a certain kind of causal-functional organization. And so on. What emerges is a distinctive and undoubtedly tempting image of what we are *doing* when we engage in philosophical inquiry (much of it, anyway): We are developing *explanatory theories* of philosophically central subjects, where the common kind of explanation being aimed for is *grounding* explanation.

We're almost there. To take the final step along the Path, it will help to consider an influential view within the literature on *scientific* explanation, and a tempting extension of it to grounding explanation.

1.2 A background view about "explanation" and explanation

David Lewis, in his influential paper "Causal Explanation" (1986), advances an attractive view about how a philosophical account of explanation ought to divide up its labor. This division of labor rests on a distinction between two senses of "explanation," a distinction Lewis borrows from Sylvain Bromberger (1965):

> To quote Sylvain Bromberger, "an explanation may be something about which it makes sense to ask: How long did it take? Was it interrupted at any point? Who gave it? When? Where? What were the exact words used? For whose benefit was it given?" But it is not clear whether just any act of explaining counts as an explanation. Some acts of explaining are unsatisfactory; for instance the explanatory information provided might be incorrect, or there might not be enough of it, or it might be stale news. If so, do we say that the performance was no explanation at all? Or that it was an unsatisfactory explanation? The answer, I think, is that we will gladly say either—thereby making life hard for those who want to settle, once and for all, the necessary and sufficient conditions for something to count as an explanation. Fortunately that is a project we needn't undertake.
>
> Bromberger goes on to say that an explanation "may be something about which none of [the previous] questions makes sense, but about which it makes sense to ask: Does anyone know it? Who thought of it first? Is it very complicated?" An explanation in this second sense of the word is not an act of explaining. It is a chunk of explanatory information—information that may once, or often, or never, have been conveyed in an act of explaining. (It might even be information that never could be conveyed, for it might have no finite expression in any language we could ever use.) (Lewis 1986, 218)

For short, we should distinguish explanatory *acts* from explanatory *information*, understanding the former to be attempts to convey some of the latter.[3] And with this distinction in hand, we can separate two philosophical tasks. One task is to articulate the standards that govern explanatory acts, so that we can say with some precision what distinguishes performing them *well* from performing them *poorly*. A distinct task is to give an account of what makes some body of information count as "explanatory" with respect to some given topic. On Lewis's view, these tasks don't have much to do with each other. In fact, the first task doesn't really have anything to do with *explanation*. After all, an explanatory *act* is (at least typically) just an attempt to grant someone who has requested it some explanatory *information*. While there is a *general* task of articulating standards that govern attempts to provide information, there is, according to Lewis, no interestingly specific version of this task focused on *explanatory* information:

> An act of explaining may be more or less satisfactory, in several different ways. It will be instructive to list them. . . . This list covers much that philosophers have said about the merits and

> demerits of explanations, or about what does and what doesn't deserve the name. And yet I have not been talking specifically about explanation at all! What I have been saying applies just as well to acts of providing information about *any* large and complicated structure. It might as well have been the rail and tram network of Melbourne rather than the causal history of some explanandum event. *The information provided, and the act of providing it, can be satisfactory or not in precisely the same ways.* There is no special subject: pragmatics of explanation. (1986, 226–28; second italics added)

Philosophical work on explanation proper should therefore focus solely on the *second* task, that of giving an account of what makes some body of information count as "explanatory" with respect to some given topic. Thus Lewis's main thesis in his paper: When the topic is a particular *event*, the explanatory information is exactly the information about that event's causal history.

Notice that the distinction between explanatory acts and explanatory information helps insulate this account from certain overly hasty criticisms. For example, suppose you ask why a certain window broke, and I say that Suzy woke up with a migraine. You're confused: you don't know who Suzy is, let alone that she is the one who broke the window, let alone that she did so because committing acts of vandalism is her way of dealing with her migraines. So you can fairly complain that what I have offered is a terrible explanation, perhaps no explanation at all. But examples like this don't automatically refute Lewis's main thesis, since he can reply that my explanatory act fails *not* because it conveys no explanatory information, but for some other reason. And here it's worth remembering that there are *many* ways that you can botch an attempt to grant a request for information.[4]

I think the distinction has had another effect on the literature, which has been to shunt questions about human psychology and reasoning off-stage as not relevant to the second of the two tasks. Here is what I mean. Suppose you haven't yet cottoned on to the distinction between explanatory acts and explanatory information. At this point in your philosophical investigations, it strikes you as just obvious that a good philosophical account of explanation will need to pay *some* attention to human psychology, human cognition, and certain related issues in epistemology. After all, an explanation of some fact or phenomenon counts as successful only if it is capable of enhancing the *understanding* of those who receive it. So a good theory of explanation needs to rest in part on a good theory of what "understanding" is and what it takes to enhance it. And that, in turn, will require delving into how human cognition works, and on the epistemic standards that govern *good* cognition.

But with the act/information distinction in hand, you now have the option of *sequestering* epistemology and human psychology: yes, they matter; but only to an account of what makes an attempt to convey information (of any sort) successful, *not* to an account of the nature of explanatory information. Lewis clearly takes this option. Alluding to his list of ways that an explanatory act can fail, he writes

> Another proposed *desideratum* is that a good explanation ought to produce understanding. If understanding involves seeing the causal history of the explanandum as simple, familiar, or whatnot, I have already registered my objection. [Namely, that there is no reason to expect the causal history to oblige.] But understanding why an event took place might, I think, just mean possession of explanatory information about it—the more of that you possess, the better you understand. If so, of course a good explanation produces understanding. It produces possession of that which it provides. But this *desideratum*, so construed, is empty. It adds nothing to our understanding of explanation. (1986, 28)

Lewis rushes here. To begin, his proposal about what "understanding why an event took place" consists in is not especially plausible: detail for me the first billion years of the causal history of the window's breaking, and I'll be no closer to *understanding* why it broke. But for our purposes, the more important point is that sequestering psychology and epistemology in the way he does *is not obligatory*. The act/information distinction makes such sequestering possible, but does not require it. We can respect the distinction, all the while insisting that a proper account *of explanatory information* needs to build upon some relevant considerations in epistemology and psychology.

As we'll see below, that is an approach worth pursuing when it comes to "grounding" explanations, and perhaps even more generally. For now, let's assemble the pieces and consider how the Plausible Path gets smoothed over if we adapt Lewis's approach to the case of grounding.

1.3 An attractive foil

Remember where we'd got to: we'd identified a kind of noncausal explanation that appears to crop up all over the place in philosophy (and perhaps in the sciences). Suppose that, following Lewis's example, we take any attempt to provide such an explanation to be an attempt to convey some distinctively noncausal kind of explanatory information. What is the best theory of the nature of this "metaphysically" (let's call it) explanatory information? Here it's tempting to follow Lewis one step further. In his case, explanatory information pertaining to some event is information about a vast, complicated structure picked out by a central metaphysical relationship: *causation*. We have the event; and we have the structure consisting of all of its causes, together with all of *their* causes, and all of *their* causes, and so on.[5] Explanatory information, for Lewis, is information *about that structure*.

Now we can offer a straightforwardly parallel treatment of grounding explanations and the noncausal information they attempt to convey. Explanatory information pertaining to some fact is information about a vast, complicated structure picked out by a central metaphysical relationship: *grounding*. We have the fact; and we have the structure consisting of all of its grounds, together with all of *their* grounds, and all of *their* grounds, and so on.[6] Metaphysically explanatory information is information *about that structure*.

We've taken the final step along the Path. We began with some beguiling examples; read those examples a certain way, as instances of an important kind of explanation; conjectured that the pursuit of this kind of explanation is central to vast swathes of philosophy; adapted Lewis's popular take on what we are doing in offering explanations; and arrived at the view that *grounding* is a centrally important metaphysical relationship—perhaps even *the* central metaphysical relationship.

Time to retrace our steps to see where we might have taken a different direction. As hinted at above, we'll focus on that critical decision to leave epistemic and psychological considerations out of our account of explanatory information.

2. TWO CASE STUDIES

We're going to work our way towards an outline of an epistemology-first approach to metaphysically (yes: *metaphysically*) explanatory information. It will help to start with a very simple case study, followed by a much richer and more complex one. The first illustrates how epistemological considerations—directly tied to human psychology—can play a role in determining, for certain kinds of questions, what counts as the right kind of information to convey in an answer (and not just what counts as the right *way* to convey that information). The second (i) illustrates how structures that are vitally important to inquiry, and that might *look* metaphysical, can in fact have purely epistemic foundations; (ii) provides an important precedent for developing one kind of epistemic approach to ground. It will, along the way, highlight the key open questions that such an approach would need to answer.

2.1 Numbers

Consider this question:

(*) What is 55 times 27?

If you like, pause for a minute to figure out the answer.

Now let's step back and ask an odd question *about* this question: What makes something *count* as "the answer" to (*)? Put in the terms introduced above, what qualifies as the sort of information that a proper response to this question should convey? To answer *those* questions, let's consider some possible responses to (*).

There are responses that say something false. E.g.: "17."

There are responses that say something true, but off topic: "Billy had cereal for breakfast."

There are responses that say something true and on topic, but insufficiently informative: "An integer less than 1500."

None of these responses succeeds at giving us the answer. Why not? Well, not *just* for the reasons noted. That is, it is *not* enough for a response to say something true, on topic, and maximally informative (in the sense of uniquely identifying the number which is the product of 55 and 27): Consider "the product of 15 and 99."[7] No, what's wanted is the *decimal* representation of the product: "The number 1485."

But why? You can get yourself into a frame of mind where that question seems genuinely puzzling. After all, what's so special about the *decimal* representation? There are plenty of other representations: the binary (10111001101), the ternary (2001000), the hexadecimal (5cd).... For that matter, what's so special about *these* kinds of representations? They are, mathematically speaking, nothing more than compact definite descriptions: "the product of 1 and 1000, plus the product of 4 and 100, plus the product of 8 and 10, plus 5." But so is, e.g., "the product of 15 and 99."

By this point you're probably getting impatient, since there is such an obvious explanation for why "1485" is special. But hold that thought for just a moment, long enough to appreciate that there is at least a *temptation*, even if one easily resisted, to "go metaphysical" here, and speculate that the representation "1485" is special because of the way it directly reveals the *metaphysical essence* of the number in question. You can provoke this temptation by appropriate use of italics: "Look, we don't know what 55 times 27 *is*—we don't know its *identity*—until we're told, or figure out, that it is 1485." Perhaps it is part of the number's *essence* that it has this decimal representation, but *not* part of its essence that it uniquely satisfies "product of 15 and 99."

Or perhaps we should return to sanity, and point out the obvious: the decimal representation is special because of *us*, because of the way that almost all of us have been trained from early childhood to use various algorithms that involve manipulating decimal representations. Thanks to that training, the decimal representation is *distinctively useful*: by means of it, we can engage in all sorts of reasoning about numbers that would otherwise be very difficult if not (for most of us) impossible.

Of course there is plenty more that could be said about this case. For example, it's worth noting that the special status of the decimal representation is context-sensitive: if you're talking to a bunch of people thoroughly fluent in binary, the answer "10111001101" might be the right one to give. More controversially, you might want to invoke the act/information distinction, and argue that the question (*) itself simply functions as a request for identifying information—*any* identifying information, as long as it accurately singles out the product of 55 and 27. (So even "the product of 55 and 27" would do.) But (you say) any *act* of providing such information is subject to various pragmatic desiderata, which somehow make it the case that (in typical contexts) the answer "1485" is, of all the many correct answers, the best one to supply. Myself, I think that's stretching the act/information distinction too far; but at any rate it will be enough for our purposes to observe that even if this is the right analysis of the case, it can't be *fully generic* pragmatic desiderata, of the kind Lewis considers, that single out the decimal representation. No, facts about the distinctive epistemic utility of the decimal representation play a special role, one way or another.

2.2 Axioms, theorems, and mathematical concepts

For our second (and much more valuable) case study, consider a branch of mathematics such as, say, the theory of arithmetic. As a body of knowledge, we could take this branch to consist in a large bunch of arithmetical claims now known to be true.[8] But in fact mathematicians impose *much* more structure than that, in at least three different ways. First, they designate certain arithmetical truths as *axioms*. Second, they introduce a host of arithmetical concepts and distinctions in order to organize inquiry; consider for example the distinction between composite and prime numbers, or the meaning of "mod" in (e.g.) "15 mod 4 = 3." Third, they routinely distinguish certain known arithmetical truths as more "central," "important," or "fundamental" than others. Such distinctions appear routinely in *proofs*: consider the difference between a "lemma," "theorem," and "corollary." But among statements that earn "theorem" status (as opposed to mere lemmas or corollaries), some are understood to be much more *central*, *important*, or *fundamental* than others; as an example, consider the Fundamental Theorem of Arithmetic.

Here's a thought experiment designed to help make vivid how much this extra structure of axioms, concepts, and key theorems matters, *epistemically*. Imagine we encounter some Alien Arithmeticians (AAs) and figure out how to communicate with them. We discover that, like us, the AAs have a great interest in figuring out truths about the numbers. We discover that our most basic vocabularies have a lot in common; in particular, they deploy all the same basic logical notions that we do, and make central use of the concepts *natural number* and *successor*. We discover, moreover, that their conception of *proof* is fundamentally the same as ours: a demonstration deploying rules of inference licensed by first-order logic with identity. And proofs are just as important to them as they are to us as a basis for claiming knowledge of some arithmetical statement.

But that's where the similarities end. And here I'm going to have the thought experiment branch in two. In Version 1 of the thought experiment, the AAs strike us as utterly alien in their choice of *axioms*. We can recognize that the arithmetical claims *they* treat as axioms are *true*; and they likewise can recognize the same of us. But each of our communities finds the other's choice of axioms utterly baffling. In Version 2 of the thought experiment, we discover, happily, that the AAs have hit upon the Peano Axioms as their preferred axioms for arithmetic (they call them by a name we can't pronounce). But they do not make use of—and see no use for—any nonlogical arithmetical vocabulary beyond what is used in their statement of their axioms. So while, for example, they can understand perfectly well which numbers are the ones we call "prime," they have no idea why we introduce a distinct term for them. What's more, every proof they produce of an arithmetical result proceeds directly from their axioms. (So, yes, their proofs are *long*.) They can't really comprehend what we're doing in distinguishing "lemmas" from "theorems" from "corollaries"; nor does our habit of singling out some results as more "mathematically significant" than others make any sense to them. Nor—crucially—can they see any point to producing *different* proofs of the same result. It's as if they view the world of known arithmetical results as wholly undifferentiated—just one damn result after another, with no discernable structure beyond the fact that all of their proofs proceed from the same starting points.

It might just be possible for us to adopt the mathematical practice of the AAs in Version 1 of the thought experiment. (Though we would likely just cheat: if we can, derive our Peano axioms from their axioms, and then proceed as usual.) But I strongly suspect that it would, in both senses of the term, be unthinkable for us to adopt their practice in Version 2.

Why? In large part, I think, because we would lose the extensive conceptual superstructure by means of which we can gain not just mathematical knowledge but mathematical *insight* and *understanding*. Indeed, the concepts of mathematical insight and understanding wouldn't really have any purchase, any more.

Example. We all know that $\sqrt{2}$ is irrational. Why? Well, the usual proof goes as follows: Suppose $\sqrt{2}$ is rational. Then for some integers m and n, $\sqrt{2} = (m/n)$. We can assume that m and n are not *both* even. Squaring both sides and rearranging gives us $2n^2 = m^2$. So m^2 is even. But then m must be even, since an odd times an odd is an odd. So $m = 2k$ for some integer k. Substituting and simplifying, it follows that $n^2 = 2k^2$. But then n^2 is even, whence n is too—a contradiction.

That bog standard proof is fine, but here's a better one—*not* because, somehow, it secures the result more firmly, but because it conveys the deeper reason why the result obtains (which has nothing to do with properties of even numbers in particular): Take any integers $k > 1$ and a where $a \neq n^k$, for any integer n. Then we can show that the kth root of a is irrational. For suppose otherwise. Then $a^{1/k} = m/n$, for some integers m and n, and where m and n share no common factors. So $a = (m^k/n^k)$. By the Fundamental Theorem, $m = m_1 \cdot m_2 \cdot \ldots \cdot m_i$ and $n = n_1 \cdot n_2 \cdot \ldots \cdot n_j$, where all the terms on the right-hand sides of these equations are prime numbers. No prime number factors both m and n. So no prime number factors both m^k and n^k. But then (m^k/n^k) cannot be an integer, and so cannot equal a. We recover the result that $\sqrt{2}$ is irrational as just one not-so-special case. More to the point, we can now appreciate that $\sqrt{2}$ is irrational for *fundamentally the same reason* that, say, the 17th root of 592 is irrational. In this way, the second proof helps put on display the *value* of singling out the Fundamental Theorem for special attention. (Value to *us*, anyway. The Version 2 aliens presumably think differently.)

I think we should generalize: the rich structure of axioms, key concepts, and important theorems we impose on our arithmetical practice earns its keep at least in part because of the way in which it facilitates explanation and understanding. That generalization raises a good philosophical question: Why, exactly, does this structure facilitate explanation and understanding? I will advance a negative answer, along with the barest beginnings of a positive answer. (For one classic and important account of explanation that bears on this question, see Kitcher [1989].)

2.3 An epistemic approach to explanatory structure

The negative answer is just that whatever is going on, it has nothing to do with *metaphysics*. What makes the Peano axioms a good choice is not that they somehow most directly capture the metaphysical nature of *number*. What makes the concept "prime number" valuable is not that it marks an "objective joint" in the numerical universe, or corresponds to (adapting Lewis's popular phrase) a "perfectly natural arithmetical property." Important theorems are not important because they mark some metaphysically distinguished location in a hierarchy of arithmetical truths. Something else—something more thoroughly *epistemic*—is going on. For short: explanatory structure in mathematics counts *as such* not because it captures or corresponds to some human-cognition-independent kind of special information about its objects, but because—for reasons having wholly to do with the nature of human cognition and reasoning—it enhances the mathematical understanding of those who grasp it.

As we'll see below, that thesis by itself is enough to get us a fair way towards an epistemic approach to ground. Still, it's also a thesis positively begging for elaboration. What, exactly, are these "reasons having wholly to do with the nature of human cognition and reasoning"? I don't know. But there are clues, clues that point to a suggestion that is at least *slightly* more specific: perhaps what makes a certain structure explanatory, at least in the mathematical case, is that grasp of it dramatically enhances one's ability to engage in relevant sorts of reasoning. An example will illustrate.

Consider the following game. Two people play. You start with a pile of coins. On each player's turn, they remove either one or two coins from the pile. The player to remove the last coin wins.

Equipped with the right list of instructions, you can play this game expertly: "If the pile contains 16 coins, take one"; "if the pile contains 17 coins, take two"; "if the pile contains 18 coins, take either one or two"; etc. *Why* does this list work? Here is an explanation:

Suppose it is your opponent's turn, and she faces a pile of three coins. Then you are guaranteed to win: if she takes one, then you take two; if she takes two, then you take one. So a 3-coin pile is a 'losing situation', precisely because 3 is one more than the maximum number of coins that can be taken. It follows that a 6-coin pile is a losing situation, since the player whose turn it is will—if their opponent plays correctly—face a 3-coin pile on their next turn. And so on: any pile of coins that is a multiple of 3 is a losing situation. Correspondingly, if it is your turn and the pile of coins is *not* a multiple of 3, then (with proper play) you are guaranteed to win if you take enough coins that the remainder is a multiple of 3.

Some observations about the explanation in the last paragraph. To begin, it really *is* an explanation. But it does not succeed as an explanation by providing causal information, or indeed any information of a metaphysically special kind. Rather, it appears to succeed because of the way it puts its recipient in a *better position to conduct inquiry* about our simple game than someone who only has the list of instructions. And this in two ways. First, it enables vastly more efficient reasoning about the game itself. Second, it makes it easier to spot generalizations to *other* games. To drive home this latter point, try the following exercises:

- Figure out how to play the game, if you can take 1, 2, or 3 coins.
- Figure out how to play the game, if you can take any number of coins up to n.
- Figure out how to play the game, if the player who takes the last coin *loses*.
- Figure out how to play the game, if there are *two* piles, and on your turn you remove coins from just one of them.

Now I am going to go out on a limb, with a fair amount of hand-waving, and suggest that the lessons from our toy example generalize in ways that point to an inquiry-centric account of the nature and value of explanatory information and explanatory frameworks (at least in the mathematical case—and if the approach works there, it of course makes sense to investigate whether it extends more broadly). Here is the idea. Just as our ability to reason effectively about takeaway games gets enhanced quite a lot by the explanation provided above, so too our ability to inquire effectively in *any* domain will hinge on our possession of an appropriately *organizing framework*, and our possession of information that allows us to make best use of that framework. So Lewis had matters precisely Euthyphro-backwards (at least, in the mathematical case): what understanding consists in is *not* possession of explanatory information; rather, what makes something *count* as explanatory information is that its possession enhances understanding, and more specifically one's ability to conduct relevant inquiry.[9]

3. RESPECTFUL DEFLATIONISM

It will be helpful to zoom out to take stock of where we are, and to set the discussion in a broader philosophical context. If the negative answer/suggestion of a positive answer we have just considered are correct, then the right philosophical approach to the concept of *arithmetical explanation* (and to mathematical explanation more generally) will be an instance of what I've elsewhere called "respectful deflationism" (see Hall [2023]). Here is the idea. Given some putatively philosophically significant concept X, we can distinguish three stances we could adopt:

Modest eliminativism:

X does not have an important role to play in any serious philosophical theorizing.

Robust realism:

X has an important role to play in at least some serious philosophical theorizing, because it marks out or closely corresponds to a distinctive kind of metaphysical structure.

Respectful deflationism:

X has an important role to play in at least some serious philosophical theorizing, but not because it marks out any distinctive kind of metaphysical structure.

Examples: Our everyday concept of "object"—the concept that distinguishes between "genuine" objects and mere arbitrary aggregates—is, arguably, a good candidate for modest eliminativism. More controversially, some philosophers (e.g., Russell 1913; Norton 2003) have argued for modest eliminativism about "cause." Next, a number of philosophical accounts of laws of nature—e.g., Maudlin's (2007) view that laws are metaphysically fundamental features of reality that govern how earlier complete states of the universe generate later states—count as examples of robust realism about laws. By contrast, Lewis's "Humean" account of laws as generalizations belonging to that set of truths about our world that best optimizes simplicity and informativeness counts as a species of respectful deflationism about laws: respectful because the resulting analysis of laws is put to so much philosophical work by Lewis (and other Humeans); deflationist because laws, so understood, do not (contra Maudlin et. al.) count as such because of the way they mark out or correspond to any kind of metaphysical structure. As another example, "axiom of arithmetic" looks ripe for respectful deflationist treatment. (Again, see Hall [2023].) Notice, finally, that while the official definition of "respectful deflationism" leaves it fairly open just *why* the given concept X has an important role to play in at least some serious philosophical theorizing, our examples suggest that this will often, if not invariably, be because of the *epistemic benefits* of employing this concept. At any rate, that's the guiding idea I'll pursue in what follows.

What about "ground"? Current orthodoxy seems to hold that we should give this and cognate terms a robustly realist treatment (with a few modest eliminativist holdouts). But respectful deflationism needs investigating as an alternative. After all, our discussion of explanatory structure in mathematics provides a clear precedent. And, reflecting on that precedent, we can now appreciate how hasty it was to follow Lewis in dismissing "understanding" as a possible source of insight into the nature of explanation. In the case of mathematics, understanding why some fact obtains does *not* appear to "just mean possession of explanatory information about it." Perhaps the same is true of metaphysics. Perhaps "grounding" isn't a *metaphysical* relation at all, any more than "illuminating proof of" is.

4. GROUND-CLEARING: THE CANONICAL EXAMPLES REVISITED

But to get a "respectfully deflationist" approach to grounding into proper view, we first need to clear up what is—at least, from this perspective—a serious mistake in the literature, which is the undue importance it gives to the canonical examples. If you follow the masses in thinking of grounding as a fundamental metaphysical relation, you'll happily treat the examples as illustrative. But if you favor a broadly epistemic approach to ground, these examples have little to teach. And that's because there's an obvious, simple, highly nongeneralizable story to tell about *why* we find them "explanatory."

4.1 How (not) to teach sentential logic

We'll begin the story with a little detour through pedagogy, focused on how and how *not* to teach the semantics of sentential logic. Suppose your task is to explain to your students what an "interpretation" is. Suppose you proceed like so:

"Consider all those functions from sentences of our formal language to truth-values. We will call such a function an 'interpretation' exactly if it meets the following conditions: It assigns opposite truth-values to each pair of a sentence and its negation. If it assigns 'true' to a conjunction (A. B), then it assigns 'true' to both A and B; otherwise, it assigns 'false' to at least one of A and B. . . . [And so on, through the rest of the connectives.]"

If your students are mathematically sophisticated, this introduction might work just fine. But if they're at all like *my* newbie logic students, it will be a disaster. What you should do instead is say something like this:

"Think of an interpretation as, in the first instance, assigning truth-values to the *atomic sentences*: p, q, r, and so forth. These truth-values then *percolate up* into truth-values for more complex sentences that are *built* from the atomic sentences. The *way* that truth-values percolate up is, in turn, determined by the specific way the more complex sentences are built. Thus, the *negation* of an atomic sentence gets the opposite truth-value to that atomic sentence; the *conjunction* of two atomic sentences gets assigned 'true' if *both* of its constituents were assigned 'true', and gets assigned 'false' otherwise. . . . [And so on, through the rest of the connectives.] Then the truth-values for these slightly more complex sentences determine, in turn, truth-values for the next level of sentences, and so on."

What you will have instilled in your students, if you proceed in something like this second way, is a kind of *inheritance model* of truth-value assignment, a model which will encourage them to think of truth-values for complex sentences as *depending on* truth-values of their parts in a hierarchical fashion. And that's useful! You *should* teach them this way, precisely because doing so will make it much easier for them understand the semantics and reason about it. But—and this is the punchline—the obvious utility of thinking in this way discloses *exactly nothing* of metaphysical interest. Just because it's *helpful* to think of, say, the truth-value of a conjunction as being "determined" by the truth-values of its conjuncts (and not vice versa) doesn't mean that there is some sort of mind-independent, objectively asymmetric relation of "determination" at work.

4.2 The examples deflated

And that means, in turn, that we should be super skeptical that anything at all deep is going on in the canonical examples. In fact, I think that what's going on is quite shallow: The examples naturally call to mind the hierarchical inheritance model, and so—finding it ever so easy to think in terms of this model—we acquiesce to the standard characterization of the examples. But if we're on our guard, that acquiescence should only go so far: "Sure, we can usefully *think* of the fact that Billy and Suzy are either having lunch or watching a movie together as 'depending on' the fact that Billy and Suzy are having lunch. After all, that way of thinking is both familiar and very helpful in other contexts, namely where we're dealing with much more complicated truth-functional compounds. But so what? It hardly follows that in so thinking, we are recognizing the presence of some distinctively metaphysical relation connecting these two facts."

Unfortunately, once the canonical examples get deflated in this way, they no longer serve as particularly useful models for thinking about how metaphysical inquiry writ large might be structured. Yes, we can agree (e.g.) that "disjunctive facts are grounded in their true disjuncts." But all that's going on is that we're tacitly invoking a heuristic model that is *specifically* useful when we're trying to understand the semantics of sentential logic. How is that supposed to help illuminate what we're doing when, say, we hypothesize that the voluntariness of an action is grounded in its psychological causes? Compare our two case studies from §2, above: It's as if we examined the desiderata on good answers to multiplication questions ("What is 55 times 27?"), worked out that they derive from the distinctive utility *for us* of decimal notation—and tried to use that observation as the basis for a fully general account of mathematical explanation. I conclude that—absent, at least, some boldly Tractarian metaphysics according to which the totality of facts is structured in a way perfectly limned by sentential connections—the canonical examples just aren't.

5. A MODEST PROPOSAL

Fortunately, we do not *need* the canonical examples in order to motivate an interest in grounding, or in metaphysical explanation more generally. We need merely step back for a moment, and reflect on the extent to which serious philosophical inquiry has as one of its aims the provision of integrated accounts of a wide range of topics of central philosophical importance. An account of what it is for an action to be *free* might appeal to the mental states of the agent, and the way they cause that action. An account of the *mental states* of an agent might appeal to the functional role of certain of their brain states. An account of *functional role* might appeal to causation; so causation gets into the story of free action via two different routes. An account of *causation* might appeal to counterfactual dependence. An account of *counterfactual dependence* might appeal to laws of nature. And so on. We're all perfectly familiar with the way in which philosophical investigations can hang together in an organized, systematic fashion.

So it is a perfectly respectable philosophical project to try to figure out what this "hanging together" amounts to. The contemporary grounding enthusiast offers *one* answer, elegantly expressed by Schaffer (2009, 351):

> . . . the neo-Aristotelian will begin from a hierarchical view of reality ordered by priority in nature. The primary entities form the sparse structure of being, while the grounding relations generate an abundant superstructure of posterior entities. The primary is (as it were) all God would need to create. The posterior is grounded in, dependent on, and derivative from it. The task of metaphysics is to limn this structure.

But another answer is available, an answer we can almost directly read off from our discussion of the second "case study" in §2.2 above. What we are doing in offering integrated, hierarchically structured philosophical explanations—explanations that we are perfectly entitled to use the language of "grounding" to convey—is simply imposing an explanatory framework on a set of truths in order to facilitate our understanding of them. But the framework itself can serve that purpose even though it corresponds to *no* distinguished structure of metaphysical dependency relations.

Pushing the (possible!) analogy between mathematical and metaphysical inquiry helps bring this option into sharper focus. In the mathematical case, we begin with a set of truths: say, all the arithmetical truths. Some are known, some not. We have a generic interest in expanding the latter category. But not *just* that: we also have, it seems, a basic epistemic interest in achieving mathematical *understanding* (witness the fact that the Alien Arithmeticians really are, well, *alien*). To that end, we impose a structure of axioms, key concepts, and important theorems, and by appeal to that structure do such things as distinguishing certain proofs as "more illuminating" than others. But this really is an *imposition*: as far as the arithmetical truths themselves are concerned, there may be no more "metaphysical structure" than what is given by bare relations of logical entailment.

The "modest proposal" of this section simply invites us to view *metaphysical* inquiry through the same lens. Of course there are differences. The subject matter of metaphysics is strictly broader; in fact, there may be no limits on it at all, at least in principle. In the mathematical case, one necessary condition on "axioms" is that they strike us as exceedingly *obvious*; good luck imposing that condition, in the metaphysical case. And, most importantly, it's an open question what the analog of "logical entailment" should be in the metaphysical case. Here I will simply float one straightforward option: as far as the metaphysical truths themselves are concerned, there may be no more "metaphysical structure" than what is given by bare relations of metaphysical necessitation. Some truths (about free action, mental states, causation, counterfactuals, what have you) supervene with metaphysical necessity on others, and that's it.

In the contemporary climate, that last statement will likely be met with howls of outrage. Haven't we known since at least the 90s that "supervenience with metaphysical necessity" isn't enough, that inquiry in metaphysics must attend to relations much finer-grained than *that*? Yes. We *have* known this. (And still do!) But so what? By the same token, mathematicians have known for ages that, in order to achieve the epistemic aims they set for themselves, they must attend to an explanatory structure much richer than what could ever be articulated *just* by focusing on logical entailment relations. But for all that, it is entirely plausible that this structure (of axioms, key concepts, important theorems, etc.) is *imposed*. And it's just very hard to see how treating it as such threatens in the least to deprive it of its credentials *as* explanatory.

So too—perhaps—in the case of metaphysics. We have, say, a set of *metaphysically possible worlds*, and a set of *propositions*.[10] As far as intrinsic metaphysical structure is concerned, propositions stand in relations of necessitation to each other: p necessitates q iff every world in which p is true is a world in which q is true. There simply is no more metaphysical structure than *that*. But the epistemic aims we metaphysicians set for ourselves require us to attend to an explanatory structure much richer than what could ever be articulated *just* by focusing on these relations of necessitation. So we *impose* such a structure, treating certain facts as "fundamental," and organizing our understanding of the necessitation relations we care about by saying that some facts "ground" other facts. And we thereby enhance our understanding, potentially quite dramatically.

Granted, the contemporary grounding partisans think that this is not *imposition*, but *discovery*. Maybe they're right. But it's very hard to see what we'd lose if they weren't.

6. UPSHOTS AND OPEN QUESTIONS

The central idea behind the sketch of an epistemic approach to ground is this: claims to the effect that one fact is "grounded in" some other facts are *not* explanatory because they provide information about relations of metaphysical dependency; rather, they are explanatory because of the way that the juxtaposition of explanandum with explanantia enhances understanding—where "enhances understanding," in turn, should be understood as a matter of enhancing the recipient's ability to conduct inquiry. And our lodestar in developing this sketch into a proper theory is, once again, the mathematical case.

The most obvious and urgent open questions, then, are whether we *can* develop this sketch into a proper theory—and if so, what that theory looks like, what its basic posits are, whether it can be extended to cover scientific and everyday explanation, and so on. But another question may have been nagging at you: In setting forth this approach, even in its highly sketchy form, I've helped myself to the kinds of locutions that approach itself is meant to weigh in on. Let's make this as explicit as possible: A grounding claim counts as explanatory not *in virtue of* the fact that it provides information about a special kind of metaphysical structure (constituted, say, by metaphysically fundamental grounding relations that obtain between facts), but *in virtue of* the kind of epistemic benefit it confers on its recipients (where this benefit itself can be characterized without needing the very suspect metaphysical posits that this approach is meant to let us avoid). Our first open question is about what, exactly, this epistemic benefit is. Our second is about whether the approach can, in good conscience, be extended to the "in virtue of" locution used in its very statement.

Yes, it can—at least, as far as I can see. On the epistemic approach, what we're *doing* when we provide good explanations in metaphysics is organizing our understanding of the facts about what metaphysically necessitates what in a way that greatly enhances our ability to conduct inquiry into what metaphysically necessitates what. So, similarly, what we're *doing* when we provide an epistemic explanation *of explanation* is organizing our understanding of the facts about what explains what in a way that greatly enhances our ability to conduct inquiry into what explains what. In short, if the epistemic approach can be made to work at all, it can be made to work *on itself*.

A final open question. What might *motivate* us to pursue the epistemic approach? Well, philosophical curiosity, obviously. But the approach also promises some upshots that are, I think, quite philosophically interesting. Here are three.

First, on the epistemic approach as sketched here, there is nothing particularly puzzling about the question, "What grounds grounding facts?" Take for example the claim that the fact that an action is free is grounded in facts about its psychological causes. Okay, in virtue of what is *that* claim true? On the epistemic approach, the answer will involve facts about inquiry, what it takes to enhance it, and perhaps relevant details about human cognition in particular. In short, "What grounds grounding facts?" becomes a question squarely in epistemology-cum-cognitive psychology.

Second, the oddly popular thesis that every fact is either metaphysically fundamental or is grounded in other facts simply goes by the wayside—or at least, looks like a case of pure speculation, with no particularly strong reason to believe it a priori. That seems to me clearly a feature, and not a bug. Compare the mathematical case, where it's child's play to ask for explanations where none are forthcoming. We can explain why $\sqrt{2}$ is irrational: we can convey information that deepens one's understanding of this fact. Can we, in the same sense, explain why the eighth digit in its decimal expansion is 5? I doubt it. But that doesn't make this fact "mathematically fundamental." A good question for fans of the popular thesis is why metaphysical explanation should work any differently. On the epistemic approach, at any rate, there's no reason to think it will. What's more, recognizing sensible limits on what can be "metaphysically explained" will help guard us against tempting pseudoproblems—such as (in my view) the "problem" of saying what grounds contingent negative existentials (see, e.g., Muñoz [2020], which characterizes this problem as a "notorious paradox").

Finally, the epistemic approach may, depending on its details, allow room for a kind of metaphysical pluralism that will (but shouldn't) strike many as quite radical. Suppose, again, that as far as objective metaphysical structure is concerned, there are propositions that stand in necessitation relations, but that's it. And suppose that human inquiry into that metaphysical structure inevitably requires more, so that we find ourselves imposing additional explanatory structure, just as we do in the case of mathematics. Well, why not structure*s*, plural? If one core epistemic aim is to improve our understanding of what necessitates what, maybe there are multiple ways to do so. For example, maybe we get one kind of understanding if we model facts about the existence of wholes as "depending on" facts about the existence of their parts. Maybe we get a different kind of understanding if we model in the reverse, taking parts to "depend on" the wholes they compose (and ultimately, on The One). Such a result would be fascinating, and perhaps even surprising. What it would *not* be is an occasion to fight about who's right.

NOTES

1. To be sure, not all of these works are specifically *about* grounding.
2. And perhaps elsewhere: for example, perhaps a scientific explanation of the solubility of table salt in terms of its molecular constitution is the same kind of noncausal explanation. We'll set this possibility to one side.
3. It's worth noting that a similar distinction shows up elsewhere. For example, an "answer" can be something that takes a certain amount of time, that is interrupted, that is given by someone at a certain place and time, etc.; but it can also be something which no one yet knows, which Suzy was the first to discover, which is not expressible in English, etc.
4. That said, there is a serious methodological issue here that Lewis's discussion does not adequately come to terms with. To bring the issue into focus, imagine the following bonkers philosophical account of explanatory information: The explanatory information pertaining to a given event is exactly the information about every event that preceded it (whether part of the causal history or not). You attempt to refute this account by pointing out that providing information about some random event that preceded the window's breaking will, typically, strike us as not explaining that breaking *at all*. I reply that these acts of explaining succeed at providing some information of the right type, but fail for other reasons. That's clearly a cheat, but *why*, exactly? Here it's well to remember that the causal history of the breaking stretches *very* far into the past, and thus likely includes such events as (say) the munching of a certain plant by a certain dinosaur, millions of years ago. We would never count an attempt to explain the window's breaking by citing that munching event as remotely successful. Yet according to Lewis, that cannot be because such attempts convey no explanatory information; it must be for some other reason. Isn't that, too, a cheat?
5. If causation is transitive, then we can say more simply: we have the structure consisting of all of the event's causes.
6. If grounding is transitive, then we can again say more simply: we have the structure consisting of all of the fact's grounds. Note in addition that we might want to amend Lewis's thesis about *events*: for in addition to an event's causes, there are the facts that ground that event's occurrence. We might want to count information about such grounds as explanatory information pertaining to the event itself. See Skow (2016). Yet another option—not so much for amendment, but for unification—is to argue that grounding just *is* causation, or at least a species of it. See Wilson (2018).
7. For that matter: "The product of 55 and 27."
8. So we are ignoring other significant, philosophically interesting aspects of mathematical practice: for example, the distinction between *important* and *trivial* open arithmetical questions, or distinctions between different techniques of proof.
9. And yes, I'm *still* aware that I'm using what are in effect "grounding" locutions in drawing this contrast. We'll get back to this.
10. We may or may not decide to *identify* the propositions with sets of possible worlds. All that matters here is that for any world w and proposition p, p is either true (and not false) at w, or false (and not true).

REFERENCES

Bromberger, Sylvain 1965. "An Approach to Explanation," in *Analytical Philosophy: Second Series*, ed. R.J. Butler, Oxford: Blackwell.

deRosset, Louis 2020. "Anti-Skeptical Rejoinders," in M. Raven, ed., *The Routledge Handbook of Metaphysical Grounding*, New York: Routledge, ch. 12.

Hall, N. 2023. "Respectful Deflationism," in H. Andersen and S. Mitchell, eds., *The Pragmatist Challenge*, Oxford: Oxford University Press.

Kitcher, Philip 1989. "Explanatory Unification and the Causal Structure of the World," in P. Kitcher and W. Salmon, eds., *Scientific Explanation*, Minneapolis: University of Minnesota Press, 410–505.

Lewis, David 1986. "Causal Explanation," in *Philosophical Papers Vol. II*, Oxford: Oxford University Press, 214–40.

Maudlin, Tim 2007. "A Modest Proposal Concerning Laws, Counterfactuals, and Explanation," in *The Metaphysics within Physics*, Oxford: Oxford University Press, 5–49.

Muñoz, Daniel 2020. "Grounding Nonexistence," *Inquiry: An Interdisciplinary Journal of Philosophy* 63(2): 209–29.

Norton, J. 2003. "Causation as Folk Science," *Philosophers' Imprint* 3: www.philosophersimprint.org/003004/

Raven, Michael J., ed. 2020. *The Routledge Handbook of Metaphysical Grounding*, New York: Routledge.
Rosen, Gideon 2010. "Metaphysical Dependence: Grounding and Reduction," in B. Hale and A. Hoffmann, eds., *Modality: Metaphysics, Logic, and Epistemology*, Oxford: Oxford University Press, 109–35.
Russell, Bertrand 1913. "On the Notion of Cause," *Proceedings of the Aristotelian Society* 13: 1–26.
Schaffer, Jonathan 2009. "On What Grounds What," in D. Manley, D.J. Chalmers and R. Wasserman, eds., *Metametaphysics: New Essays on the Foundations of Ontology*, Oxford: Oxford University Press, 347–83.
Skow, Bradford 2016. *Reasons Why*, Oxford: Oxford University Press UK.
Wilson, Alastair 2018. "Metaphysical Causation," *Noûs* 52(4): 723–51.
Wilson, Jessica M. 2014. "No Work for a Theory of Grounding," *Inquiry: An Interdisciplinary Journal of Philosophy* 57(5–6): 535–79.