A priori internal explanation

For *Grounding and Consciousness*, (ed.) G. Rabin, OUP
Draft: August 1, 2025

Abstract: Some claim that experiential truths lack transparent constitutive explanations in non-mental terms. But what is a *transparent* constitutive explanation? On the standard view, a constitutive explanation is transparent just in case its explanans a priori entails its explanandum. Some challenge the extensional adequacy of this view. Rather than try to resolve this matter, I propose an alternative view, one available to proponents and opponents of the standard view alike. My ecumenical proposal is that a constitutive explanation is transparent just in case it's *a priori internal* in the sense that its explanans and explanandum together a priori entail that the former constitutively explains the latter. An upshot is that proponents and opponents of the standard view can in effect understand the claim that there is an explanatory gap between the experiential and physical in the same way. Relatedly, if you think there is an explanatory gap but you're on the fence about the standard view, you have another, perhaps ultimately more attractive, account of transparency to consider.

1. Introduction

Some claim that there is an explanatory gap between the experiential and the physical. On one interpretation, the claim is that experiential truths lack transparent constitutive explanations in non-mental terms. Let's start with the notion of *constitutive* explanation. There is a constitutive sense of the term 'because' in addition to a causal sense (Dasgupta 2017). Suppose you ask me why there was a meeting yesterday. And suppose I answer by saying something about the causal history that led to it (e.g., "There was a meeting then because you sent out invitations for it last week"). Here the causal sense of 'because' is operative—in this case I provide a causal explanation of the timing of the meeting. But I might interpret your question differently and answer instead by describing events that occurred yesterday in virtue of which there was a meeting then (e.g., "There was a meeting then because that was the time we met to discuss matters"). Here the constitutive sense of 'because' is operative—in this case I provide a constitutive explanation of the timing of the meeting. I'm going to take the notion of constitutive explanation for granted. Both causal and constitutive explanation are thoroughly familiar notions we deploy in a variety of contexts, including everyday thinking.[1]

Okay, what are *transparent* explanations? Such explanations are "compelling", possessing a "hard-to-specify clarity" (Sturgeon 1994, 226). We find these explanations "satisfying", even when they contain little information (Melnyk 2025, 16). For such explanations, there is a "discernible connection" between explanans and explanandum (Levine 2001, 77); as such, the

[1] How does constitutive explanation relate to *grounding*? Among proponents of grounding, some identify it with constitutive explanation (e.g., Dasgupta 2017), others with a determination relation that can back or underwrite explanations (e.g., Schaffer 2016). More on grounding later.

link between them is “intelligible” (Levine 2006, 145). On the face of it, the constitutive explanation gestured at above is transparent in the relevant sense. Transparent constitutive explanations are a tacit goal of many reductive projects in the sciences (Chalmers 2012, 305-307).

These comments, together with examples we will soon consider, isolate the relevant notion of transparency. But how should we ultimately understand transparency in this context? On the standard view, a constitutive explanation is transparent just in case its explanans a priori entails its explanandum. Some challenge the extensional adequacy of this view—they claim that it classifies clear cases of transparent constitutive explanation as non-transparent. Rather than try to resolve this matter one way or another, I offer an alternative view, one available to proponents and opponents of the standard view alike. My ecumenical proposal is that a constitutive explanation is transparent just in case it’s *a priori internal* in the sense that its explanans and explanandum together a priori entail that the former constitutively explains the latter. For proponents of the standard view, the idea is that a constitutive explanation is a priori just in case it’s a priori internal. For opponents, the idea is that, while some transparent constitutive explanations are priori and some aren’t, any such explanation is a priori internal, and any a priori internal explanation is transparent. The internalist view gets clear cases of transparent constitutive explanation right. And, given this view, it’s reasonable to maintain that no experiential truth has a transparent constitutive explanation in non-mental terms. An upshot is that proponents and opponents of the standard view can in effect understand the claim that there is an explanatory gap between the experiential and physical in the same way. Relatedly, if you think there is an explanatory gap but you’re on the fence about the standard view, you have another, perhaps ultimately more attractive, account of transparency to consider.

2. A priori explanation

Sentence S_1 a priori entails sentence S_2 just in case the material conditional ‘If S_1 then S_2 ’ is a priori. And a sentence is a priori just in case we’re in a position to know that the sentence is true independently of experience. Following Chalmers (2012, Ch. 2, Seventh Excursus) we can define sentential apriority in terms of thought apriority: a sentence is a priori when we’re in a position to have a thought that constitutes a priori knowledge and is apt to be expressed by that sentence.

Let’s say that an explanation is a priori just in case its explanans a priori entails its explanandum. The standard view is that a constitutive explanation is transparent just in case it’s a priori (Chalmers and Jackson 2010; Chalmers 2012, Twelfth Excursus). This view comports with the orienting characterization of transparency set out above, as well as the idea that a transparent constitutive explanation leaves no “residual open question” concerning why the explanandum is

true (Schaffer 2017, 3), as, once we've spelled out the explanans, the explanandum is "rendered transparent" (Chalmers 2012, 305).[2]

Given the standard view, it's reasonable to maintain that experiential truths lack transparent constitutive explanations in non-mental terms. Consider, for example, the following experiential truth: such-and-such neural states have a visual-experience-as-of-motion phenomenal character. And consider a potential constitutive explanation in non-mental terms of this truth, one whose explanans says that the neural states involve recurrent activity between the V1 and MT/V5 structures. No matter what the relevant activity comes to, this potential constitutive explanation isn't transparent—it's neither compelling nor satisfying, as there is no discernible connection between the explanans and explanandum. And this explanation is a priori just in case the following conditional is a priori:

(1) If neural states involve recurrent activity between the V1 and MT/V5 structures, then they have a visual-experience-as-of-motion phenomenal character.

This conditional clearly isn't a priori. Corresponding reasoning applies to other experiential truths.

However, is the standard view extensionally adequate? Consider the following constitutive explanation: there are genes because there are molecules that play the gene role. This constitutive explanation is transparent—it's compelling and satisfying despite the lack of details, as there is a discernible connection between the explanans and explanandum in virtue of which the link is intelligible (Chalmers 2012, 307). And this explanation is a priori just in case the following conditional is a priori:

(2) If molecules play the gene role, then there are genes.

Functional concepts like GENE have substantive application conditionals. On the face of it, (2) is such a conditional and is therefore a priori.[3] So, the standard view seems to get this case right—it says that the explanation above is transparent. The same goes for other constitutive explanations with explananda that concern functional roles and explanantia that concern role-fillers—they're transparent, and the standard view apparently classifies them as such.

---

[2] Levine in places uses similar language to describe transparency—when we understand a transparent explanation, the question of why its explanandum is true no longer has "substantive content" or "cognitive significance" for us (2001, 83). But Levine rejects the standard view, as we will discuss shortly.

[3] Some argue that that properties can disassociate completely from their causal roles (Schaffer 2005). This view, however, is standardly restricted to fundamental properties. Some non-fundamental properties are *functional* properties, properties whose natures are exhausted by their causal roles. And some functional properties are such that we conceive of them under functional concepts, concepts specifying their causal roles.

However, some argue that the standard view doesn't count all clear cases of transparent constitutive explanation as being transparent. Consider the following constitutive explanation: there is water because there is H_2O that uniquely plays the watery role in our local environment. This explanation is also transparent. And it's a priori just in case the following conditional is a priori:

(3) If there is H_2O that uniquely plays the watery role in our local environment, then there is water.[4]

Proponents of the standard view claim that (3) is an application conditional for WATER and thus a priori. But you might think that we can't conclusively rule out that H_2O is *zombie-H_2O*, H_2O that uniquely plays the watery role in our local environment such that there is no water (Levine 2001, Ch. 2; 2006). In this case, (3) isn't a priori.

This point of controversy should be familiar. But there is another that may be less familiar, one that isn't directly connected to theoretical identities. Let's say that hydrogen and oxygen atoms are *arranged H_2O-wise* just in case two hydrogen atoms are covalently bonded to a central oxygen atom at an angle of 104.45. Sometimes I'll just speak of atoms being arranged H_2O-wise, but I have in mind hydrogen and oxygen atoms in particular. Here's another transparent constitutive explanation: there is H_2O because there are atoms arranged H_2O-wise. And this explanation is a priori just in case the following conditional is a priori:

(4) If there are atoms arranged H_2O-wise, then there is H_2O.

Proponents of the standard view claim that (4) is an application conditional for the concept of H_2O and is thus a priori. But you might think that we can't conclusively rule out that H_2O-wise arrangements of atoms are *zombie-arrangements*—atoms arranged H_2O-wise such that there is no H_2O (Schaffer 2017). As a rationale, you might claim that there are H_2O molecules only if there are individuals composed by hydrogen and oxygen atoms, yet we can't conclusively rule out that, while there are atoms arranged H_2O-wise, they don't compose anything. For that matter, perhaps we can't conclusively rule out that nothing composes anything distinct from itself (Rosen & Dorr 2002). In this case, (4) isn't a priori.

As indicated above, I'm not going to try to adjudicate these issues here. I will, however, note some limitations of a couple lines of response. As for (4), suppose it isn't a priori. Still, we can expand its antecedent in a way that rules out whatever scenarios keep (4) from being a priori. In this case, the resulting conditional will itself be a priori. Now, there is some question about just what adjustments would achieve this result. You might think that all we need to do is add to the antecedent of (4) an answer to the special composition question (Chalmers 2012, Ch. 6). But, while the expanded antecedent in this case a priori entails that there are composite individuals,

[4] Why the condition about uniqueness? The thought is that, if H_2O and *ghost-water* (some occult substance) both play the watery role in our local environment, then there is no water (Block and Stalnaker 1999).

perhaps it doesn't a priori entail that there are H_2O molecules in particular. (Rationale: while something is an H_2O molecule only if it has causal powers, the truth of the expanded antecedent in this case is compatible with epiphenomenalism about composite individuals.) At any rate, it's clear that there are adjustments that will get the job done. For example, we could abandon any reductive aspirations we might have and add a principle directly linking being arranged H_2O-wise to composing H_2O molecules in particular.

Remember, however, that our target is a particular transparent constitutive explanation: there is H_2O because there are atoms arranged H_2O-wise. The worry is that (4) isn't a priori, so the explanation itself isn't a priori. Whether there are *other* transparent constitutive explanations of this macroscopic truth that are a priori (e.g., explanations whose explanantia include the explanans of our target explanation) is neither here nor there.

Still, you might think that there is a simple path forward here. An assumption in the discussion above is that something is a priori only if alternative, incompatible claims can be conclusively ruled out. Apriority is characterized by the strong epistemic standards satisfied by deduction rather than the lower epistemic standards satisfied by induction and abduction, or so the idea goes. The points of controversy described above, then, are about the scope of *conclusive* a priori knowledge. Some claim, however, that there is also a *non-conclusive* form of a priori knowledge (Hawthorne 2002). If there is both conclusive and non-conclusive a priori knowledge, then we can say that a constitutive explanation is a priori just in case its explanans a priori entails, either conclusively or non-conclusively, its explanandum. And you might think that (3) and (4) are at least non-conclusively a priori.

On the face of it, whatever considerations we might draw on to motivative the idea that (3) and (4) are non-conclusively a priori (e.g., abduction), the same reasoning will apply to conditionals like (1) that concern experiential truths. And if conditionals like (1) are non-conclusively a priori, then given the current proposal we lose the explanatory gap between the experiential and physical. So, it may be that the price of adjusting the standard view so that it doesn't misclassify the cases we've discussed is jettisoning the explanatory gap.[5] But I'm interested in accounts of transparency that potentially verify the claim that there is an explanatory gap between the experiential and the physical. And the standard view, at least in the hands of Chalmers and Jackson (2010), seems to target conclusive apriority in any case.[6]

While I'm going to set the standard view understood in a way to include non-conclusive apriority aside, the view makes what I think is a useful suggestion: while transparency ultimately concerns a priori knowledge (rather than, say, pragmatic success in supporting predictions and

[5] Schaffer (2017) discusses a similar issue in the context of Chalmers' (2009; 2012, Ch. 6) "stretched-out" notion of analyticity. Chalmers introduces a notion of analyticity such that all sorts of things count as being analytic, including that H_2O molecules exist. Supposing that (3) and (4) are analytic in this attenuated sense, it's not entirely clear why we shouldn't also think that (1) is analytic in the same sense.

[6] Chalmers (2012, Ch. 4) considers both conclusive and non-conclusive a priori entailment in this context.

interventions), we can bracket the particular points of controversy concerning conclusive a priori knowledge described above. The discussion in the next section is guided by this suggestion.[7]

3. A priori internal explanation

In what follows, by 'a priori' I mean conclusively a priori in particular. I propose that a constitutive explanation is transparent just in case it's *a priori internal* in the sense that its explanans and explanandum together a priori entail that the former constitutively explains the latter. Given this internalist view, the claim that there is an explanatory gap between the experiential and the physical is the claim that no experiential truth has a constitutive explanation in non-mental terms such that it's a priori that if the explanans and explanandum are true, then the former constitutively explains the latter.[8]

As with the standard view (cast in terms of conclusive a priori knowledge), it's reasonable in this case to maintain that experiential truths lack transparent constitutive explanations in non-mental terms. Consider again the following experiential truth: such-and-such neural states have a visual-experience-as-of-motion phenomenal character. And consider again a potential constitutive explanation in non-mental terms of this truth, one whose explanans says that the neural states involve recurrent activity between the V1 and MT/V5 structures. This explanation is a priori internal just in case the following conditional is a priori:

(5) If neural states involve recurrent activity between the V1 and MT/V5 structures and they have a visual-experience-as-of-motion phenomenal character, then they have this phenomenal character because they involve that activity.

This conditional clearly isn't a priori. Corresponding reasoning applies to other experiential truths.

What about the extensional adequacy of the alternative view? Recall the gene case: there are genes because there are molecules that play the gene role. This constitutive explanation is a priori internal just in case the following conditional is a priori:

---

[7] For alternative takes on the explanatory gap sensitive to the apparent epistemic possibility of zombie-arrangements of atoms as characterized above, see Rabin 2019 and Sassarini 2021.

[8] If interest-relativity is built into constitutive explanation (and not merely the *practice* of giving such explanations—see Dasgupta 2017 for discussion of the difference), then the characterization of a priori internal explanation above is elliptical. Here are things fully spelled out: if S_1 constitutively explains S_2 (given our interests), then this explanation is a priori internal just in case it's a priori that if S_1 and S_2 are true, then (given these interests) S_1 constitutively explains S_2. And the claim that there is an explanatory gap between the experiential and the physical is the claim that no experiential truth has a constitutive explanation in non-mental terms (given our interests) such that it's a priori that if the explanans and explanandum are true, then (given these interests) the former constitutively explains the latter.

(6) If molecules play the gene role and there are genes, then there are genes because molecules play the gene role.

While (2) from above is plausibly an application conditional for GENE, (6) is plausibly an application conditional for the constitutive sense of 'because', CONSTITUTIVE EXPLANATION. After all, it's a priori that if there are molecules that play the gene role and there are genes, then playing the gene role *functionally realizes* being a gene in these cases. (If something has *P* and *Q*, then *P* functionally realizes *Q* on this occasion just in case *Q* is the property of having some property or other that plays causal role *R*, and *P* plays *R* on this occasion.) If both (2) and (6) are a priori, then both the standard and internalist views say that the explanation above is transparent.

Let's return to the water and H_2O cases: there is water because H_2O uniquely plays the watery role in our local environment, and there is H_2O because there are atoms arranged H_2O-wise. These explanations are a priori internal just in case the following conditionals are a priori:

(7) If there is H_2O that uniquely plays the watery role in our local environment and there is water, then there is water because there is H_2O that has these features.

(8) If there are atoms arranged H_2O-wise and there is H_2O, then there is H_2O because there are atoms arranged in this way.

On the face of it, (7) and (8) are also application conditionals for CONSTITUTIVE EXPLANATION. After all, it's a priori that if there is H_2O that uniquely plays the watery role in our local environment and there is water, then water is identical to H_2O. And it's a priori that if there are atoms arranged H_2O-wise and H_2O molecules, then the former compose the latter.[9] The status of (7) and (8) as a priori is entirely compatible with it being epistemically possible that H_2O is zombie-H_2O and that H_2O-wise arrangements of atoms are zombie-arrangements. If (7) and (8) are a priori, then the internalist view says that the explanations above are transparent.[10]

Here's one more example of a transparent constitutive explanation: there is a statue on the pedestal because there is material in the same place with such-and-such properties. This explanation is a priori just in case the following conditional is a priori:

---

[9] Following Gilmore (2018), an entity is *simple* just in case it lacks proper parts and *extended* just in case it's spatiotemporal and doesn't have the shape and size of a point. For all we know a priori, there are simple entities with the shape and size of H_2O molecules. And it's a priori open that such extended entities occupy the fusion of the regions occupied by hydrogen and oxygen atoms arranged H_2O-wise. But we can conclusively rule out that any such entity is an H_2O molecule. This is because it's a priori that something is an H_2O molecule only if it's composed by hydrogen and oxygen atoms.

[10] An anonymous referee has suggested that it's a priori open that H_2O uniquely plays the water role and there is water, yet something prevents the former from constitutively explaining the latter. (And, if interest-relativity is built into constitutive explanation—see note 8—the idea as I understand it is that explanation failure is possible in this case even if we have the relevant interests.) But, in my view, this is no more a priori open than, say, an apple being red and something blocking it from being colored.

(9) If the material on the pedestal has such-and-such properties, then there is a statue there.

If (9) is an application conditional for STATUE, then (9) is a priori. But you might think that we can't conclusively rule out that the material on the pedestal is *zombie-material*—material with the relevant properties such that there are no statues. As a rationale, you might claim that there are statues only if there are materially constituted individuals, yet we can't conclusively rule out that, while there is material with the relevant properties, it doesn't constitute anything. For that matter, perhaps we can't conclusively rule out that nothing materially constitutes anything (Unger 1979).[11]

By contrast, the constitutive explanation above is a priori internal just in case the following conditional is a priori:

(10) If the material has such-and-such properties and there is a statue in the same place, then there is a statue there because there is material there with those properties.

On the face of it, (10) is an application conditional for CONSTITUTIVE EXPLANATION. After all, it's a priori that if there is material with such-and-such properties on the pedestal and there is a statue in the same place, then the former materially constitutes the latter. And the status of (10) as a priori is entirely compatible with it being epistemically possible that the material on the pedestal is zombie-material. If (10) is a priori, then the internalist view says that the explanation above is transparent.

In our discussion of how the internalist view handles clear cases of transparent constitutive explanation, I've appealed to various intuitive judgments about what's a priori. However, I think that these judgments have a theoretical backing as well. Let S_1 and S_2 be the explanans and explanandum of either (6), (7), (8), or (10). The argument I have in mind has two key premises. The first concerns how knowledge and understanding interact when it comes to S_1 and S_2 in particular:

The understanding premise: if you know that S_1 and S_2 are true, you're thereby in a position to understand what it is in virtue of which S_2 is true.

Now consider the following principle: you're in a position to understand why something is true only if something explains it and you're in position to know that it does so. The second premise is a special case of this principle, provided that understanding *what* it is in virtue of which something is true is a way of understanding *why* it's true:

[11] You might also think that something is a statue only if it has causal powers, yet we can't conclusively rule out that, while there are materially constituted individuals with the relevant properties, they lack causal powers. For that matter, perhaps we can't conclusively rule out that no materially constituted individual has causal powers.

The explanation premise: you're in a position to understand what it is in virtue of which S_2 is true only if something constitutively explains S_2 and you're in position to know that it does so.

Suppose that the understanding and explanation premises are true. Hence, if you know that S_1 and S_2 are true, then something constitutively explains S_2 and you're in a position to know that it does so. It must be that what you're in a position to know in this case is that S_1 in particular constitutively explains S_2—what else could the relevant explanans be? Hence, if you know that S_1 and S_2 are true, you're thereby in a position to know that S_1 constitutively explains S_2. And, if knowing that S_1 and S_2 are true on its own puts you in a position to know that S_1 constitutively explains S_2, then it's a priori that if S_1 and S_2 are true then S_1 constitutively explains S_2.

The explanation premise seems plausible. At the very least, the general principle on which it's based follows from the "received view" of understanding (Khalifa 2017, 18).[12] What about the understanding premise? I concede that there is little daylight between it and the conclusion of the argument. (Given the non-ampliative nature of the argument, this shouldn't be especially surprising.) Still, I take it that the understanding premise enjoys greater antecedent plausibility than the conclusion. So, the positive argument above moves the needle, though perhaps only slightly. Now, if you reject the understanding premise and maintain that either (6), (7), (8), or (10) isn't a priori, you must think that there are certain epistemic possibilities that I haven't taken into account. So, in the next section I supplement the positive argument above with a negative argument, one consisting of responses to various claims to the effect that there are such possibilities.

4. Alternative, direction, and absence

Again, let S_1 and S_2 be the explanans and explanandum of either (6), (7), (8), or (10). You might think that either Alternative, Direction, or Absence is epistemically possible:

Alternative: S_1 and S_2 are true, but S_1 doesn't constitutively explain S_2; nevertheless, something else does.

Direction: S_1 and S_2 are true, but S_2 constitutively explains S_1 rather than the other way around.

Absence: S_1 and S_2 are true, but nothing constitutively explains S_2.

Note that, if S_2 has a constitutive explanation and S_2 constitutively explains S_1 rather than the other way around, then both Alternative and Direction are true. And, if nothing explains S_2 and

[12] Khalifa (2017) defends a variation of the received view.

S_2 constitutively explains S_1 rather than the other way around, then both Direction and Absence are true.

If either Alternative, Direction, or Absence is epistemically possible, then either (6), (7), (8) or (10) isn't a priori. With respect to Alternative, Direction, and Absence, I propose the *conflation hypothesis*: if they seem epistemically possible, this is because we aren't properly distinguishing their content from that of other claims that are epistemically possible. While I focus below on instances of Alternative, Direction, and Absence that concern atoms arranged H_2O-wise and H_2O in particular, similar considerations apply to other instances of these claims.

What might lead us to not properly distinguish the content of Alternative, Direction, or Absence from that of other claims? There are at least three possibilities: we might conflate either constitutive explanation with *ultimate* constitutive explanation, constitutive explanation with *grounding*, or H_2O with *watery stuff*. For each of Alternative, Direction, and Absence, I'll just consider one of these three possibilities in order to streamline our discussion. So, keep in mind that, relative to these possibilities, the discussion below provides only a sampling of what might go wrong when we consider these claims.

We begin with Alternative: while S_1 and S_2 are true, S_1 doesn't constitutively explain S_2 but something else does. Consider the following instance of Alternative:

(11) There are atoms arranged H_2O-wise and there is H_2O, but the fact that there are atoms arranged H_2O-wise doesn't constitutively explain why there is H_2O; nevertheless, something else does: the fact that the cosmos has such-and-such properties.[13]

Let *priority monism* be the thesis that any sub-cosmic fact has a constitutive explanation in terms of cosmic facts that themselves lack constitutive explanations.[14] Priority monism so understood is epistemically possible, and you might conclude that (11) is epistemically possible on this basis. This, however, would be to conflate constitutive explanation with ultimate constitutive explanation (constitutive explanations whose explanantia themselves lack constitutive explanations), and thereby conflate the content of (11) with that of (12):

(12) There are atoms arranged H_2O-wise and there is H_2O, but the fact that there is H_2O doesn't have an *ultimate* constitutive explanation in terms of H_2O-wise arrangements of

[13] Suppose that, in addition to hydrogen and oxygen atoms arranged H_2O-wise and H_2O, there are alien atoms arranged H_2O-wise. Would this not be a more straightforward instance of Alternative? No—in the relevant scenarios, hydrogen and oxygen atoms compose and constitutively explain the existence of H_2O molecules. And, since the alien atoms arranged H_2O-wise don't compose H_2O molecules (see note 9), the former don't constitutively explain the existence of the latter.

[14] For Schaffer (2010), priority monism in the first instance instead is a thesis about what determines what, as he formulates the view in terms of *grounding* and identifies grounding with a form of determination rather than explanation (see note 1).

atoms; nevertheless, this fact does have such an explanation in terms of the cosmos as a whole.

Recall Direction: S_1 and S_2 are true, yet S_2 constitutively explains S_1 rather than the other way around. Consider the following instance of Direction:

(13) There are atoms arranged H_2O-wise because there is H_2O rather than the other way around.

If (13) seems possible, this may be because we're conflating constitutive explanation with grounding, here understood to be a determination relation that can back or underwrite explanations (see note 1). In this case, we in turn conflate the content of (13) with that of (14):

(14) H_2O molecules ground H_2O-wise arrangements of atoms rather than the other way around.

Now, earlier I suggested that it's a priori that if there are atoms arranged H_2O-wise and H_2O molecules then the former *compose* the latter. Why not think it's also a priori that if there are atoms arranged H_2O-wise and H_2O molecules then the former *ground* the latter? If this grounding claim is a priori then (14) isn't epistemically possible. In this case, if we conflate the content of (13) with that of (14), we aren't conflating something that isn't epistemically possible with something that is.

Grounding is sometimes introduced as "metaphysical causation" and a comparison to causation will be helpful here. Berstein suggests that a "probable" difference between the two concerns the "a priori discoverability of grounding claims versus the a posteriori nature of causal claims" (2016, 33). It seems, however, that there are a priori and a posteriori elements to both causal and grounding claims.

With implicit reference to background conditions, consider the following causal claim: were event e_1 to occur then it would cause event e_2. According to Sober (2011), some claims of this form (including ones drawn from the theory of natural selection) are a priori. It's a priori that if e_1 occurs and the causal claim is true then e_2 occurs. So, if the causal claim is a priori, then it's also a priori that if e_1 occurs then it causes e_2. And if that's a priori, then it's of course a priori that if e_1 occurs *and* e_2 *occurs* then the former causes the latter. This instance of causation is a priori internal—relative to canonical descriptions of the cause and effect, it's a priori that if both occur then the former causes the latter. Still, instances of causation typically aren't a priori internal. For example, suppose at time t_1 we ignite hydrogen and oxygen gases with a spark. And suppose that in doing so we create an H_2O molecule at later time t_2. It's not a priori that, if we manipulate hydrogen and oxygen gases in this way (in the right conditions) at t_1 and the H_2O molecule exists at t_2, then our activity causally determines the H_2O molecule (or anything else). Appeals to experience (beyond whatever experiences are required for the acquisition of the relevant concepts) are necessary here.

Consider the following grounding claim: such-and-such facts *non-factively* ground [φ]. According to Siscoe (2021), some claims of this form (e.g., 'The shirt being navy non-factively grounds that it's blue') are a priori. Suppose it's a priori that if such-and-such facts obtain and non-factively ground [φ] then those facts ground [φ]. So, if the non-factive grounding claim is a priori, then it's also a priori that if those facts obtain then they ground [φ]. And if that's a priori, then it's of course a priori that if those facts obtain *and [φ] obtains* then the former ground the latter. This instance of grounding is a priori internal—relative to canonical descriptions of the ground and grounded, it's a priori that if they both obtain then the former grounds the latter. However, like causation, some instances of grounding aren't a priori internal. Suppose the fact that there are H_2O-wise arrangements of atoms grounds the fact that there are H_2O molecules. It's not a priori relative to canonical descriptions of these facts that if they obtain then the former grounds the latter, as we can't conclusively rule out that the grounding goes in the other direction.[15]

Suppose, then, that (14) is epistemically possible. There are epistemically possible worlds in which H_2O molecules ground H_2O-wise arrangements of atoms. Nevertheless, these worlds are such that constitutive explanation goes in the other direction—the fact that there are atoms arranged H_2O-wise constitutively explains why there is H_2O. Really? I think so. Note that, for all we've said, there are other senses of explanation such that the direction of grounding and explanation coincide in these worlds. For example, let's stipulate that the following is an application conditional for what we might call *ground* explanation: if H_2O molecules ground H_2O-wise arrangements of atoms, then the fact that there is H_2O ground explains why there are atoms arranged H_2O-wise. Constitutive explanation isn't ground explanation. After all, CONSTITUTIVE EXPLANATION is quotidian, while GROUND EXPLANATION is a special purpose notion. While ground explanation might be expressed in some contexts with an everyday word such as 'because', this doesn't mean that we have a pre-theoretical grasp of the notion.[16] The same is true of grounding itself (Kovacs 2023, 302-303). So, if you're resistant to what I say about the direction of constitutive explanation in the worlds at issue above, this may because you aren't properly distinguishing between constitutive explanation and something like ground explanation.[17]

We finally turn to Absence: S_1 and S_2 are true, but nothing constitutively explains S_2. Consider the following instance of Absence:

---

[15] For that matter, Schaffer (2018) takes seriously a principle according to which the existence and features of integrated wholes such as H_2O molecules ground the existence and features of their parts.

[16] As I'm conceiving of things here, an asymmetry between causal and ground explanation is that CAUSAL EXPLANATION is quotidian in nature while GROUND EXPLANATION isn't.

[17] There are independent reasons to think that constitutive explanation and grounding (and grounding explanation) can come apart—Glazier (2017) argues that there are instances of constitutive explanation (what Glazier calls *metaphysical* explanation) that track essential connections rather than grounding connections.

(15) There are atoms arranged H_2O-wise and there is H_2O, but there is no constitutive explanation of why there is H_2O.

(15) may seem epistemically possible if we conflate H_2O with *watery stuff* (roughly, the presentation of H_2O). In this case, we in turn conflate the content of (15) with that of (16):

(16) There are atoms arranged H_2O-wise and there is watery stuff, but there is no constitutive explanation of why there is watery stuff.

(16) is epistemically possible provided that it's epistemically possible that, while there are atoms arranged H_2O-wise and no H_2O, there are *F*s (e.g., extended simples spatially coincident with H_2O-wise arrangements of atoms) that play the water role whose existence resists constitutive explanation.[18]

Having fleshed out the conflation hypothesis, let's turn to a potential problem with appealing to it in defending the internalist view. Suppose for the moment that the hypothesis as I've described it is true. You might think in this case that we should expect that a corresponding hypothesis with respect to zombie arrangements of atoms is also true. Consider the following claim:

(17) There are atoms arranged H_2O-wise but no H_2O.

The corresponding hypothesis: if it seems that (17) is possible, this is because we aren't properly distinguishing its content from that of other claims that are possible. In particular, we conflate the content of (17) with one of the following:

(18) There are atoms arranged H_2O-wise and H_2O, but there is no ultimate constitutive explanation of the latter in terms of the former.
(19) There are atoms arranged H_2O-wise and H_2O, but the former don't ground the latter.
(20) There are atoms arranged H_2O-wise and H_2O, but H_2O doesn't play the water role.

If this is right, then perhaps opponents of the standard view, rather than endorsing the internalist view, should just change course and endorse the standard view. The worry, in a nutshell, is that the conflation hypothesis, if true, ultimately undermines rather than helps motivate the internalist view as an alternative to the standard view.

Note, however, that (17) doesn't say anything about constitutive explanation, so there is no reason to think that we conflate the content of (17) with that of either (18) or (19). (17) does, though, mention H_2O. Hence, a live possibility is that in considering (17) we conflate H_2O with watery stuff and thereby conflate the content of this claim with that of (20). Still, the original

[18] Exercise for reader: work out how things go if we conflate constitutive explanation with grounding or H_2O with watery stuff in considering (11), conflate H_2O with watery stuff in considering (13), or conflate constitutive explanation with either ultimate constitutive explanation or grounding in considering (15).

conflation hypotheses is more plausible than the corresponding hypothesis in so far as the former identifies realistic routes to conflation that the latter doesn't. And that we're more likely to get into trouble in considering instances of Alternative, Direction, and Absence compared to (17) should come as no surprise given the relatively complex conceptual structure of the former compared to the structure of the latter.

Ultimately, I think that a reasonable line can be drawn that lets in zombie-arrangements of atoms but not worlds in which there are atoms arranged H_2O-wise and H_2O but no constitutive explanation of the latter in terms of the former. And it's fine if certain conflation hypotheses—including ones modelled after what I propose above in defending the internalist view—prove useful in defending the standard view in any case. While I suspect that the internalist view is extensionally more adequate than the standard view, remember that I'm not really trying to show that this is so here. Instead, I propose the internalist view in an ecumenical spirit—it's supposed to be a view that you might find attractive whether you're for or against the standard view. Recall that, for proponents of the standard view, the idea is that a constitutive explanation is a priori just in case it's a priori internal. For opponents, the idea is that, while some transparent constitutive explanations are priori and some aren't, any such explanation is a priori internal, and any a priori internal explanation is transparent. I concede, however, that the proposal may ultimately be of more interest to the opponent of the standard view (e.g. Levine) than the proponent (e.g. Chalmers).

5. The explanatory gap challenge to physicalism

I close by considering implications of our discussion for Levine's (2014, 2020) explanatory gap challenge to physicalism. Some explanations are *reductive*. An explanation of an experiential truth in non-mental terms is reductive. According to Levine, there is an explanatory gap between the experiential and the physical in the sense that no experiential truth has a constitutive explanation that is both transparent and reductive. The best explanation of why this is the case is that experiential truths lack reductive constitutive explanations, period. And physicalism requires that experiential truths have reductive constitutive explanations. Therefore, physicalism is false. This is the explanatory gap challenge to physicalism.[19]

[19] We actually shouldn't think of this as *the* explanation gap challenge, as there are other abductive challenges to physicalism worth considering. For example, all concepts are representational, but only phenomenal concepts are *presentational* as well in that they seem to situate phenomena—in this case the contents of experience—"directly before the mind" (Sturgeon 2000, 24). You might claim that, for any phenomenal concept, there is no transparent constitutive explanation in non-mental terms of its presentational nature, and the best explanation of why this is the case is that there is no constitutive explanation in non-mental terms of why the concept has this feature, period. Levine (2020) ultimately rejects physicalism on this basis. But why this version of the challenge rather than the original? Well, you might think that the physicalist can appeal to special features of phenomenal concepts to explain why experiential truths lack constitutive explanations that are both transparent and reductive—this is an instance of the so-called phenomenal concept strategy (Stoljar 2005). But Levine argues that features of these concepts themselves (specifically, their presentational nature) resist a physicalist treatment.

Here's how I understand the rationale for the abductive premise (the premise about the best explanation of the explanatory target). Roughly speaking, *ordinary macroscopic* truths are truths concerning macroscopic phenomena couched in ordinary, non-mental terms. Any ordinary macroscopic truth has a constitutive explanation that is both transparent and reductive. Hence, an ordinary macroscopic truth has a reductive constitutive explanation only if it has a transparent reductive constitutive explanation. Experiential truths and ordinary macroscopic truths are analogous—they both concern macroscopic phenomena and are couched in ordinary terms. So, it stands to reason that, like ordinary macroscopic truths, an experiential truth has a reductive constitutive explanation only if it has a transparent reductive constitutive explanation.

As we've seen, the standard view is that a constitutive explanation is transparent just in case it's a priori. And Chalmers and Jackson (2010) and Chalmers (2012, Ch. 3) argue at length that ordinary macroscopic truths have constitutive explanations that are both a priori and reductive. But consider again the fact that there is H_2O. This counts as an ordinary macroscopic truth in the relevant sense. There is a constitutive explanation of why there is H_2O that is both reductive and transparent only if there is such an explanation whose explanans appeals to H_2O-wise arrangements of atoms characterized in the language of (classical) physics. As we've seen, some claim that zombie-arrangements—H_2O-wise arrangements of atoms such that there is no H_2O—are epistemically possible relative to how we normally describe H_2O-wise arrangements of atoms. If that's right, then presumably zombie-arrangements are epistemically possible relative to descriptions of H_2O-wise arrangements couched in the language of physics as well, no matter what further non-ordinary-macroscopic truths we might appeal to (e.g., other truths that go into Chalmers and Jackson's (2010) PQTI). So, it's not entirely clear that ordinary macroscopic truths have constitutive explanations that are both a priori and reductive. And thus it's not entirely clear that the explanatory gap challenge to physicalism gets off the ground with the standard view of transparency.

I've already noted that Levine (2006) rejects the standard view. So, in advancing the explanatory gap challenge, he doesn't have a priori explanation in mind. Yet he doesn't offer an informative characterization of transparency in its place. In the absence of such a characterization, it's not totally clear what the explanatory target of the abductive argument is in the first place. But we're now in a position to clarify the argument. A constitutive explanation is transparent just in case it's a priori internal. So, our explanatory target is why experiential truths lack constitutive explanations that are both a priori internal and reductive. The best explanation of our target is that experiential truths lack reductive constitutive explanations, period. The rationale for the abductive premise proceeds upon the claim that any ordinary macroscopic truth has a constitutive explanation that is both a priori internal and reductive. This is something that proponents and opponents of the standard view alike might endorse.[20]

---

[20] Thanks to Phil Bricker, Sam Cowling, Louis deRosset, Ed Ferrier, Tim Fuller, Dan Hoek, Andrew Lee, Joe Levine, Kirk Ludwig, Donnchadh O'Conaill, Gabe Rabin, Kevin Richardson, Jonahan Schaffer, Rohan Sud, Zach Thornton, Gene Witmer, anonymous referees, and audiences at the American Philosophical

References

Bernstein, S. 2016. "Grounding is not causation," *Philosophical Perspectives* 30/1: 21-38.

Block, N. and R. Stalnaker. 1999. "Conceptual analysis, dualism, and the explanatory gap," *Philosophical Review* 108/1: 1-46.

Chalmers, D. 2009. "Ontological anti-realism." In R. Wasserman et al (eds.), *Metametaphysics*. Oxford University Press.

Chalmers, D. 2012. *Constructing the World*. Oxford University Press.

Chalmers, D. &. F. Jackson. 2010. "Conceptual analysis and reductive explanation." In Chalmers's *The Character of Consciousness*. Oxford University Press.

Dasgupta, S. 2017. "Constitutive explanation," *Philosophical Issues* 27/1: 74-97.

Gilmore, G. 2018. "Location and mereology." In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.

Glazier, M. 2017. "Essentialist explanation," *Philosophical Studies* 174/11: 2871-2889.

Hawthorne, J. 2002. "Deeply contingent a priori knowledge," *Philosophy and Phenomenological Research* 65/2: 247-269.

Khalifa, K. 2017. *Understanding, Explanation, and Scientific Knowledge*. Cambridge University Press.

Kovacs, D. 2023. "Varieties of grounding skepticism," *Monist*, 106/3: 201-316.

Levine, J. 2001. *Purple Haze*. Oxford University Press.

Levine, J. 2006. "Phenomenal concepts and the materialist constraint." In T. Alter & S. Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge*. Reprinted in Levine's *Quality and Content* (2018): Oxford University Press.

Levine, J. 2014. "Modality, semantics, and consciousness," *Philosophical Studies* 167/3: 775-784.

Association Eastern Division (Jan. 2025), UMass Amherst (Oct. 2024), University of Birmingham (Oct. 2020), and University of Gothenburg (Nov. 2018) for helpful feedback.

Levine, J. 2020. “A posteriori physicalism and the explanatory gap.” In U. Kriegel (ed.), *Oxford Handbook of the Philosophy of Consciousness*. Oxford University Press.

Melnyk, A. 2025. *Phenomenal Properties and the Intuition of Distinctness*. Oxford University Press.

Rabin, G. 2019. “Grounding the gaps or bumping the rug?” *Journal of Consciousness Studies* 26: 191–203.

Rosen, G. and C. Dorr. 2002. “Composition as a fiction.” In R. Gale (ed.), *The Blackwell Companion to Metaphysics*. Blackwell.

Sassarini, E. 2021. “No ground to bridge the gap,” *Synthese* 199/3-4: 7981-7999.

Schaffer, J. 2005. “Quiddistic knowledge,” *Philosophical Studies* 123/1: 1–32

Schaffer, J. 2010. “Monism: the priority of the whole,” *Philosophical Review* 119/1: 31-76.

Schaffer, J. 2016. “Grounding in the image of causation,” *Philosophical Studies* 173/1: 49-100.

Schaffer, J. 2017. “The ground between the gaps,” *Philosophers’ Imprint* 17: 1-26

Schaffer, J. 2018. “Monism.” In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.

Siscoe, R. 2021. “Grounding and a priori epistemology: challenges for conceptualism,” *Synthese* 199/3-4: 11445-11463.

Sober, E. 2011. “A priori causal models of natural selection,” *Australasian Journal of Philosophy* 89: 571-589.

Stoljar, D. 2005. “Physicalism and phenomenal concepts,” *Mind and Language* 20/5: 469-94.

Sturgeon, S. 1994. “The epistemic view of subjectivity,” *Journal of Philosophy* 91/5: 221-235.

Sturgeon, S. 2000. *Matters of Mind: Consciousness, Reason, and Nature*. Oxford University Press.

Unger, P. 1979. “There are no ordinary things,” *Synthese* 41/2: 117-154.