

Transparency and the Explanatory Gap

Kelly Trogon

Draft: Dec. 2021—comments welcome

For *Grounding and Consciousness*, (ed.) G. Rabin, Oxford University Press

1. Introduction

Some explanations make transparent why the explanandum should be true given the truth of the explanans. Transparent explanations possess a “hard-to-specify clarity” that renders them “compelling” (Sturgeon 1994, p. 226). When the explanans of an explanation is transparently linked to its explanandum, their connection is “intelligible” and “non-arbitrary” (Levine, 2006, p. 14). By contrast, when they aren’t so linked, there is a “residual open question” concerning why the explanandum should be true given the truth of the explanans (Schaffer 2017, p. 3; see also Chalmers 2012, p. 305).

Consider a superficial property by which we identify the presence of water in a kettle, such as having a boiling point of 212°F (at sea level). As water is H₂O, there are, of course, H₂O molecules in the kettle that have this property as well. Let BOILING be the proposition that these H₂O molecules have a boiling point of 212°F. Plausibly, there is a collection of propositions concerning the behavior of particular hydrogen and oxygen atoms—let ATOM be this collection—such that ATOM actually explains BOILING. And, according to Levine (1983), the explanans of this explanation is transparently linked to its explanandum.¹

Now consider the phenomenal character of pains—how it feels to be in pain—that underwrites our ability to identify when we are in pain. If pain is physical property *N* (*N* is for ‘neural’, as this property is instantiated by brains or relatively large proper parts of brains), then instances of *N* have the phenomenal character of pains.² Suppose you burn your finger on the kettle and *N* is instantiated on this occasion. Let PAIN be the proposition that this *N*-instance has the phenomenal character of pains. Let a *physical explanation* be any explanation whose explanans consists of propositions representing facts concerning the instantiation of (non-mental) physical properties, perhaps along with other non-mental facts (e.g., indexical facts). Levine suggests that, for any potential explanation of PAIN, if it’s a physical explanation, then its explanans isn’t transparently linked to its explanandum.

More generally, Levine endorses the *contrast claim*: while there are actual physical explanations whose explanantia are transparently linked to their explananda (e.g., the explanation of BOILING in terms of ATOM), for any potential explanation of why an instance of a physical property has the phenomenal character it does (or why it has any

¹ Explanantia and explananda are propositions, where propositions represent facts (where the latter are understood to be non-representational in nature). *Potential* explanations are “live options”, explanatory proposals that pass an initial contextually determined epistemic filter (Lipton 1991, p. 61). *Actual* explanations are potential explanations that meet all the criteria for being explanations, including that their explanantia and explananda are true.

² On my usage of ‘physical’, physical properties/facts include functional properties/facts.

phenomenal character at all), if this explanation is a physical explanation, then its explanans isn't transparently linked to its explanandum. It's this epistemic asymmetry that Levine (1983) has in mind when he claims that theoretical identity claims like 'Pain is *N*' leave special explanatory gaps while theoretical identity claims like 'Water is H₂O' don't.

Contrasters (advocates of the contrast claim) like Levine deem certain explanations transparent and others non-transparent. But, in the absence of a characterization of transparency more substantive than the orienting characterization above, it's not entirely clear why contrasters don't also classify certain familiar explanations as being transparent. Consider, for example, simple (non-covering law) causal explanations, ones whose explanantia cite causes and whose explananda cite effects of these causes (e.g., an explanation expressed by the sentence 'My finger hurts because I touched the hot kettle'). No contraster I'm aware of would say that such explanations are transparent. But, while such explanations do seem to leave residual open questions (e.g., Why should events of this type cause events of that type?), there is a clear sense in which the link between their explanantia and explananda is non-arbitrary and intelligible—after all, events described by the former cause events described by the latter.

Happily, there are three closely related substantive characterizations of transparency in the literature that underwrite the transparency judgments of contrasters. According to Levine (1993), for any explanation of A (a proposition) in terms of Δ (a collection of one or more propositions), Δ is transparently linked to A just in case it's *inconceivable* that Δ is true while A is false.³ According to Chalmers (1996, Ch. 2), Δ is transparently linked to A just in case it's *logically impossible* that Δ is true while A is false. And according to Chalmers and Jackson (2001), Δ is transparently linked to A just in case it's *a priori* that if Δ is true then A is true (see also Chalmers 2012, Twelfth Excursus). Crucially, on each of these characterizations, transparency involves deductive inference. For example, on the a priori proposal the thought is that we know a priori that if Δ is true then A is true, so there is a deductive argument with Δ and this conditional claim as its premises and A as its conclusion. Returning to the issue raised above, on these characterizations simple causal explanations come out non-transparent, while causal explanations whose explanantia include covering laws may come out transparent.

Schaffer (2017) argues that the contrast claim is implausible on each of these characterizations of transparency as follows: on each characterization, the mere coherence of certain metaphysical theses shows that explanations contrasters claim are transparent aren't.⁴ Focusing for now on the a priori characterization of transparency, here's what Schaffer says about the explanation of BOILING in terms of ATOM, which is taken by contrasters to be a paradigm case of a transparent explanation. Let *ghost H₂O* be H₂O that's epiphenomenal. As ghost H₂O lacks causal powers, it obviously doesn't boil at 212°F. While ghost-H₂O (we will suppose) is metaphysically impossible,

³ Levine goes on to reject this characterization of transparency (as well as the other two described above)—see, e.g., Levine 2006, p. 14. We will return to this issue later.

⁴ In an influential paper, Block and Stalnaker (1999) also reject the contrast claim, focusing on different issues. See Chalmers and Jackson (2001) for a response.

mereological zeroism—the thesis that, while there are complex concrete entities such as H₂O molecules, they're all epiphenomenal—is a coherent view. Hence, there is no contradiction or incoherence in supposing that there is ghost H₂O. And, thus, the same goes for supposing that ATOM is true yet the H₂O molecules don't have a boiling point of 212°F. As it's not a priori that if ATOM is true then BOILING is true, the former isn't transparently linked with the latter, given the a priori characterization of transparency.

I'm inclined to agree with Schaffer. Given the a priori characterization of transparency, the idea that ATOM is transparency linked with BOILING seems to overestimate what we know a priori. The claim that H₂O molecules have particular causally relevant properties is highly plausible. But this claim is also substantive in nature. Indeed, it's a substantive matter whether they have any causally relevant properties whatsoever. This ensures that it's a priori open that, while there are H₂O molecules, they lack these features. And I'm also inclined to agree that, paired with the conceivability or logical impossibility characterizations of transparency, the idea that ATOM is transparently linked to BOILING either underestimates what is conceivable or overestimates what is logically impossible.

If contrasters don't get what are held up as paradigm cases of transparent explanation right, then the contrast claim is in trouble. And this in turn potentially undermines the view that there is a special explanatory gap between the physical and experiential, a view that has shaped discussions of phenomenal consciousness for nearly four decades. But I think that there is a substantive characterization of transparency relative to which the explanation of BOILING in terms of ATOM is transparent, while no potential physical explanation of PAIN is transparent. And, more generally speaking, I think that in this case the contrast claim is defensible.⁵ As I will explain, there are three key elements to this characterization of transparency: grounding understood as a relation between facts, fine-grained metaphysical relations between constituents of facts, and inductive inference (specifically inference to the best explanation), in addition to deductive inference. The overall idea is that, for any explanation of A in terms of Δ , Δ is transparently linked with A just in case, given the truth of Δ and A, a particular combination of deductive and inductive inferences—ones concerning metaphysical relations between constituents of the facts represented by Δ and A—license the claim that the facts represented by Δ ground the fact represented by A.⁶

⁵ As for other explanations not involving phenomenal character (e.g., potential explanations of the normative in terms of the natural) that we might compare with the explanation of BOILING in terms of ATOM, I'm happy to let them fall where they may.

⁶ You might think that transparency so conceived is too different from what Levine et al have in mind to count as capturing a sense of 'transparency' by their lights. My primary interest, however, is whether there is a substantive epistemic asymmetry between the explanations at issue in the contrast claim, one that might support the idea that there is a special explanatory gap between the physical and phenomenal. Whether this is best described in terms of the (or a) transparent/non-transparent distinction is of secondary interest, despite how I've framed the discussion above.

2. Other proposals

I'm not the first, however, to respond to Schaffer's challenge by attempting to characterize transparency in way that would render the contrast claim (or something similar to it) defensible—Rabin (2019) and Sassarini (forthcoming) do so as well. In this section I spell out their approaches and point to some limitations. Considering these limitations will clarify certain matters relevant to my own approach to transparency.

2.1. *Being-laid-out-as-if*

Mereological nihilists coherently claim that nothing (non-degenerately) composes anything. While common-sense has it that there are complex entities like H₂O molecules, mereological nihilists maintain that there are none. But they claim that, while BOILING from above is false, there is nonetheless a sense in which it's correct. On one approach to the matter, BOILING is correct in the sense that 'The H₂O molecules have a boiling point of 212°F' (a sentence that expresses BOILING) is relevantly similar to a sentence that expresses a true proposition, specifically 'The particles arranged H₂O-wise have a boiling point of 212°F' (Rosen and Dorr 2002). To borrow a phrase from Merricks, while BOILING is false, it's "nearly as good as true" given the availability of this paraphrase (2001, p. 171). And note that mereological nihilists and proponents of complex objects alike can maintain that the proposition expressed by this paraphrase is true—in this case it's just that the former claim that BOILING is false, while the latter claim that it's true.

Flatworlders coherently claim that nothing grounds anything.⁷ Where F_B is the fact represented by BOILING, if there are no grounded facts, then F_B doesn't obtain, as this fact obtains only if it's grounded. And, if F_B doesn't obtain, then BOILING is false. So, while common sense has it that there are H₂O molecules and that they have various causally relevant features, flatworlders, like mereological nihilists, deny this. But, according to Rabin (2019), flatworlders can reasonably claim that, while BOILING is false, there is nonetheless a sense in which it's correct. Rabin introduces a grounding correlate to the notion of being arranged *F*-wise, the notion of certain facts *being laid out as if* certain facts obtain. This proposition is correct in the sense that 'The H₂O molecules have a boiling point of 212°F' (again, a sentence that expresses BOILING) is relevantly similar to a sentence that expresses a true proposition, namely 'There are facts laid out as if H₂O molecules have a boiling point of 212°F.' So, from the perspective of flatworlders, BOILING is perhaps nearly as good as true. And note that flatworlders and proponents of grounded facts alike can maintain that the proposition expressed by this paraphrase is true—in this case it's just that the former claim that BOILING is false, while the latter claim that it's true.

What, however, is it for some facts to be laid out as if some other fact obtains? Rabin characterizes the notion as follows: some facts are laid out as if some other fact obtains just in case the former facts are such that, were there the "type of grounding relations" that "permit" the latter fact to obtain, then the latter would obtain (p. 195). I think it may be more perspicuous to frame the discussion in terms of grounding *principles*

⁷ The term was coined by Bennett (2011, p. 211).

rather than relations. With this in mind, the idea is that flatworlders can reasonably maintain the following: there are fundamental (i.e., ungrounded) physical facts such that, were the general principles specifying how grounding connections run (e.g., facts concerning parts ground facts concerning wholes) compatible with H₂O molecules having a boiling point of 212°F, then H₂O molecules would have a boiling point of 212°F.

Returning to our main topic, while Rabin doesn't put things in exactly these terms, we can read Rabin as proposing the following: for any explanation of A in terms of Δ , Δ is transparently linked to A just in case it's a priori that, if Δ is true, then the facts Δ represent are laid out as if the fact represented by A obtains. And we can read Rabin as making two further claims. First, there is an actual explanation of BOILING in terms of propositions representing fundamental physical facts such that it's a priori that if the latter are true, then the facts they represent are laid out as if the fact represented by BOILING obtains. Second, for any potential explanation of PAIN in terms of propositions representing fundamental physical facts, it's not a priori that if the latter are true then the facts they represent are laid out as if the fact represented by PAIN obtains. Hence, given Rabin's conception of transparency, the former explanation is transparent, while any explanation along the lines of the latter isn't.

What should we make of this proposal? My sense is that the being-laid-out-as-if notion isn't sufficiently clear for the purposes of characterizing transparency. An indication that this is so is that the notion apparently isn't up to the task of playing the theoretical role that Rabin originally assigns to it. This, you will remember, is the task of enabling flatworlders to specify a substantive sense in which BOILING is correct. Flatworlders can pull this off with recourse to the being-laid-out-as-if notion, as they can reasonably claim that there are facts laid out as if H₂O molecules have a boiling point of 212°F, or so the idea goes. But, given Rabin's characterization of this notion, apparently flatworlders should *deny* this being-laid-out-as-if claim. Supposing that the relevant fundamental physical facts obtain, they're laid out as if H₂O molecules have a boiling point of 212°F only if the nearest worlds to the actual world with grounding principles that don't on their own rule out there being H₂O molecules are worlds in which there are H₂O molecules. But it seems that the nearest worlds to the actual world aren't like this, provided that there actually are no H₂O molecules, as flatworlders maintain.

We might instead replace "permits" with "requires" in Rabin's characterization of being-laid-out-as-if. But any fundamental physical fact is such that it's a priori that if it obtains, then, were the grounding principles to *require* that *N*-instances from above have the phenomenal character of pains, then *N*-instances would have the phenomenal character of pains. In this case we would lose the proposed epistemic asymmetry—there would be an actual physical explanation of PAIN whose explanans is transparently linked to its explanandum.

Rabin might respond that, while his initial characterization of being-laid-out-as-if might misfire, the notion itself is clear enough for the purposes of characterizing transparency, as it's a straightforward extension of the notion of being arranged *F*-wise. But what is it

for particles to be arranged *F*-wise in the first place?⁸ On one approach, particles are arranged *F*-wise just in case the particles have thus-and-so features, and were there particles with these features as well as *F*s, then the particles would compose the *F*s.⁹ Using this characterization of being arranged *F*-wise as a model, the following alternative characterization of Rabin's being-laid-out-as-if notion comes into focus: some facts are laid out as if some fact obtains if and only if, were the latter facts and the former fact to obtain, then the former would ground the latter.

For reasons that will become clear later (see §3), I don't think that transparency understood in terms of Rabin's notion so conceived supports the contrast claim. The problem is that apparently there is no collection of true propositions representing fundamental physical facts such that it's a priori that if they and BOILING are true, then the facts represented by the former ground the fact represented by the latter. But, as we will see, there is something in this proposal that is useful to the project of characterizing transparency with the aim of reintroducing a core epistemic contrast between the relevant cases.

2.2. Ordinary existence claims

Chalmers (2009) draws a distinction (one he takes to be uncontroversial and pre-theoretical in nature) between *ordinary* and *ontological* existence claims. Roughly speaking, ontological existence claims are those made inside the "ontology room", while ordinary existence claims are made in everyday life. One and the same sentence (e.g., 'There are H₂O molecules') can be either an ordinary or ontological existence claim depending on the context. Chalmers suggests that true ordinary existence claims (unconditional claims such as 'There H₂O molecules' as well as ampliative conditional claims such as 'If there are particles arranged H₂O-wise, then there are H₂O molecules) are a priori. By contrast, unconditional and ampliative ontological existence claims lack determinate truth-values.

Chalmers doesn't speak of the ordinary/ontological distinction applied to propositions, but let's say that a proposition is an ordinary existence proposition just in case it's expressed by an ordinary existence claim in Chalmers' sense. While Sassarini (forthcoming) doesn't put things in exactly these terms, we can read Sassarini as proposing the following: for any explanation of A in terms of Δ , Δ is transparently linked to A just in case Δ and A are ordinary existence propositions, and it's a priori that if Δ is true then A is true as well. And we can read Sassarini as making two further claims. First, where BOILING* is the ordinary existence proposition that there are H₂O molecules, there is an actual explanation of BOILING* in terms of ordinary existence propositions concerning the locations of and relations between certain hydrogen and oxygen atoms

⁸ While I won't explore this here, the notion of being arranged *F*-wise itself might be problematic. The reasoning here parallels the discussion above: given extant characterizations of being arranged *F*-wise, it's unclear that appealing to it actually helps mereological nihilists specify a substantive sense in which propositions expressed by sentences like 'The H₂O molecules have a boiling point of 212°F' are correct (Tallant 2014).

⁹ This characterization is modelled after Merricks' (2001, Ch. 1) discussion.

such that it's a priori that if these propositions are true then BOILING* is true. Second, where PAIN* is the ordinary existence proposition that there are pains, for any potential explanation of PAIN* in terms of ordinary existence propositions concerning non-mental entities, it's not a priori that if those propositions are true then PAIN* is true. Hence, given Sassarini's characterization of transparency, the former explanation is transparent, while any explanation along the lines of the latter isn't.

What should we make of this proposal? To begin, note that Sassarini's discussion targets *existence explanation*—explanations aiming to answer existence questions (e.g., 'Why are there H₂O molecules?')—rather than *nature explanation*—explanations aiming to answer questions concerning why something has the features it does (e.g., 'Why does the *N*-instance have the phenomenal character of pains?'). The contrast claim, however, targets nature explanation. Hence, you might think that our focus, at least initially, should be on explanations with PAIN rather than PAIN* as their explanandum. But if the ordinary/ontological distinction applies to existence claims, then it presumably applies to non-existence claims as well. In this case, we can speak of ordinary propositions in general, which include both propositions directly concerning existence like BOILING* and propositions that don't. And we can therefore reformulate Sassarini's proposal in a way that it also applies to ATOM, BOILING, PHYSICAL, and PAIN, provided that these are all ordinary propositions.¹⁰

A deeper concern with the proposal, however, is this: it's unclear that there is a collection of ordinary propositions concerning the locations of and relations between certain hydrogen and oxygen atoms such that it's a priori that if these propositions are true then BOILING* is true. (And the same goes for ATOM and BOILING understood as ordinary propositions.) Why not?

Merricks argues that ordinary existence claims concerning non-living macroscopic objects (save for those concerning living organisms) are false (2001, p. 19). According to Merricks, the view that there are no non-living macroscopic objects, what he calls *eliminativism*, "is striking and surprising simply because—and this is the obvious answer—it contradicts what nearly all of us believe" (2001, p. 163). Eliminativism so understood is a coherent if implausible view. Similarly, mereological nihilism as I understand the view is striking and surprising because it contravenes our ordinary beliefs—a consequence of the thesis so understood is that our ordinary existence claims, including claims about complex microscopic entities such as H₂O molecules—are false.¹¹ To take mereological nihilism to have the consequence that BOILING* understood as an ordinary existence proposition is false doesn't render the view incoherent. So

¹⁰ We will return to the distinction between existence and nature explanation in §3.2 when we distinguish between explanations of why there are *N*-instances (instances of the physical property that have the phenomenal character of pains), and explanations of why these property instances have the phenomenal character of pains.

¹¹ While mereological nihilists so understood claim that ordinary existence claims concerning H₂O molecules are false, they are free to claim that there is a sense in which the claims are correct, recalling our discussion in §2.1. And, provided that they agree there are particles arranged H₂O-wise, they deny that, when people in everyday contexts say 'There are H₂O molecules', what they mean is that there are particles arranged H₂O-wise.

mereological nihilists coherently claim that BOILING* is false. And, were BOILING* to follow a priori from some collection of true ordinary existence propositions concerning the locations of and relations between certain hydrogen and oxygen atoms, it would be incoherent to maintain that BOILING* is false. The point isn't so much that mereological nihilism so understood is in fact the view discussed in the literature called 'mereological nihilism' but that the view is coherent.

3. The hybrid characterization

3.1. The proposal

Three considerations serve as points of departure in characterizing transparency with the aim of reintroducing a core epistemic contrast between the relevant cases. The first is a lesson from Schaffer's discussion, namely that considerations involving the a priori (or conceivability or logical possibility) alone don't capture the relevant asymmetry. The second comes directly from Rabin's discussion: the notion of grounding will be useful to appeal to in this context. And the third is something implicit in the alternative characterization of Rabin's notion of being-laid-out-as-if considered above. This is the idea that key to understanding what it is for an explanans to be transparently linked to an explanandum involves what our epistemic situation is like not just with respect to situations in which the explanans is true (e.g., what follows a priori from ATOM being true) but with respect to situations in which the explanans *and* explanandum are true (e.g., what follows a priori from *both* ATOM and BOILING being true). Interestingly, I take it that Levine would actually agree with the first and third points, given how his conception of the explanatory gap has evolved over time.¹²

As for the first point, Levine ultimately concedes that there is a central notion of conceivability such that it's conceivable that ATOM is true yet BOILING is false (1998; 2001, Ch. 3). On this basis, he concludes that the epistemic difference between the ATOM-BOILING link and the link between the explanans in any potential physical explanation of PAIN and its explanandum isn't captured by what's conceivable and what isn't in these cases. And Levine draws the same conclusion about a priori derivability (2006, p. 14). I take it that he would say the same about logical possibility.

As for the third point, I read Levine as ultimately suggesting an alternative approach to transparency that embraces this idea concerning which sorts of situations we should be considering. According to Levine, ATOM is transparently linked to BOILING in that, were both ATOM and BOILING true and we were to know the underlying chemistry, then certain questions about the connection between these propositions wouldn't have "substantive content" or "cognitive significance" for us (2001, p. 83). One such question would be: "But how are ATOM and BOILING connected such that propositions like the latter are true when propositions like the former are true?" By contrast, PHYSICAL (the explanans of any potential physical explanation of PAIN) isn't transparently linked with PAIN in that, were

¹² In personal correspondence, Levine has expressed skepticism about grounding, so I doubt that he would be sympathetic with the second point above. There is, however, a grounding-free version of my approach to transparency, one that may be attractive to grounding skeptics and agnostics—see notes 22 and 26.

both PHYSICAL and PAIN true and we were to know the underlying neurophysiology, corresponding questions would still have substantive content or cognitive significance for us. What's important for our purposes here isn't so much what the notion of substantive content or cognitive significance comes to as it is that for Levine the transparency of an explanation is related to our epistemic situation with respect to not just situations in which the explanans is true but situations in which both the explanans and explanandum are true.¹³

With these points of departure in mind, we now turn to what I call the *hybrid* characterization of transparency (as it appeals to inductive as well as deductive inference). Roughly put, the idea behind the proposal is this. With respect to certain explanations, it's priori that if the explanantia and the explananda are true, then entities they concern stand in certain special relations. And it's plausible, as an abductive matter, that among these special relations are ones whose instantiation ensure that the facts represented by the explanantia ground the facts represented by the explananda. It's explanations with these two features that are transparent according to the hybrid characterization.

More carefully put, the hybrid characterization says that, for any explanation of A in terms of Δ , Δ is transparently linked to A just in case there is some metaphysical relation, R, that satisfies the following conditions:

The a priori condition: it's a priori that, if Δ and A are true, then Δ -entities are R-related to A-entities in thus-and-so manner; and

The abductive condition: it's plausible on abductive grounds that if Δ -entities are R-related to A-entities in thus-and-so manner, then the facts represented by Δ ground the fact represented by A.¹⁴

Δ -entities are entities that Δ concerns, and A-entities ones that A concerns (e.g., the proposition that Socrates is a philosopher concerns Socrates and the property of being a philosopher). Metaphysical relations—the “special” relations alluded to above—are dyadic, non-causal, irreflexive, asymmetric relations familiar from metaphysics that take constituents of facts rather than facts themselves as their relata. They include (non-degenerate) composition, realization, singleton formation, material constitution, and the determinable-determinate relation. These relations also include principled restrictions of the relations just mentioned (e.g., composition where the composed are mere aggregates as opposed to integrated wholes), inverses of these relations (e.g., decomposition), and principled restrictions of their inverses (e.g., decomposition where we decompose from integrated wholes rather than mere aggregates).

¹³ See Trogdon 2013 for related discussion.

¹⁴ There are constraints on what propositions might be the explanans or explanandum of an explanation, depending on the type of explanation. As for the type of explanation we're interested in above, the propositions expressed by 'The fact immediately grounded by thus-and-so fact' and 'the fact whose constituents stand in metaphysical relation R to thus-and-so fact' aren't candidate explananda.

Consider a potential explanation of A in terms of Δ . Suppose that Δ and A are true and they meet the a priori condition specified above. In this case, the following is a sound deductive argument, where the second premise is knowable a priori:

- (1) Δ and A are true.
- (2) If Δ and A are true, then Δ -entities are *R*-related to A-entities in thus-and-so manner.
- (3) Hence, Δ -entities are *R*-related to A-entities in thus-and-so manner.

And suppose that Δ and A also meet the abductive condition specified above. In this case, the claim that the facts represented by Δ ground the fact represented by A is underdetermined by our evidence (which includes the fact that Δ -entities are *R*-related to A-entities in thus-and-so manner) and the rules of deduction. Still, this grounding claim is entailed by the best explanation, among a competitive list of potential explanations, of some aspect of our evidence related to Δ -entities being *R*-related to A-entities. And the best explanation—roughly the potential explanation that, if correct, would provide the deepest understanding compared to competing potential explanations—is the likeliest explanation (Lipton 1991, Ch. 4). In this case, the non-demonstrative, ampliative judgement that the facts represented by Δ ground the fact represented by A is justified. Hence, we may suppose that the following is a strong inductive argument:

- (4) Δ -entities are *R*-related to A-entities in thus-and-so manner.
- (5) Hence, the facts represented by Δ ground the fact represented by A.

Once we string together the deductive and inductive inferences above, the explanation of A in terms of Δ seems clear and compelling. In this case, there is a clear sense in which the connection between Δ and A is intelligible and non-arbitrary. In contexts in which the intelligibility of their connection is salient, there is a clear sense in which follow up questions as to why A should be true given the truth of Δ are misplaced.

3.2. *The contrast claim revisited*

Now we can return to the contrast claim. I want to suggest that the contrast claim is defensible given the hybrid characterization of transparency. To keep the discussion manageable, I'll focus on providing reasons to think that, given this characterization, the actual explanation of BOILING in terms of ATOM is transparent, while no potential explanation of PAIN is, provided that the explanation is a physical explanation.

Given the hybrid characterization of transparency, the explanation of BOILING in terms of ATOM is such that its explanans is transparently linked to its explanandum just in case there is some metaphysical relation, *R*, that satisfies two conditions. The first is the a priori condition: it's a priori that, if ATOM and BOILING are true, then ATOM-entities are *R*-related to BOILING-entities in thus-and-so manner. There is a metaphysical relation that satisfies this condition—*realization*. Let REALIZATION be the proposition that instances of properties of the hydrogen and oxygen atoms realize an instance of the property of

having a boiling point of 212°F of the H₂O molecules. It's a priori that if ATOM and BOILING are true then REALIZATION is true.

What, however, about *quiddistic contingentism*, the view that properties can disassociate completely from their causal powers? This thesis is not only coherent but widely endorsed (see, e.g., Schaffer 2005 for a defense). The worry is that, for all we know a priori, the property of having a boiling point of 212°F isn't a functional property, a property whose nature is exhausted by its causal role. As only instances of functional properties are realized, this means that it isn't a priori that if ATOM and BOILING are true then REALIZATION is true.

In response, note that quiddistic contingentism is standardly restricted to *sparse* properties (i.e., perfectly natural properties that "carve reality at its joints"). So, the thesis doesn't say that having a boiling point 212°F isn't a functional property; instead, it has the consequence that, provided that this property is a functional property, it's an abundant as opposed to sparse property. And this is a good thing, as it seems a priori that, for any property *X*, *X* = having a boiling point 212°F only if *X* is a functional property.¹⁵

The second condition is the abductive condition: it's plausible on abductive grounds that if REALIZATION is true then the F_{AS} (the facts represented by ATOM) ground F_B (the fact represented by BOILING). Why not think, however, that there is an a priori rationale for this claim linking realization and grounding? The reason is that any substantive claim linking metaphysical relations to grounding is a priori open, as flatworldism—the view we discussed earlier according to which nothing grounds anything—is coherent, as is flatworldism combined with claims to the effect that thus-and-so metaphysical relation is instantiated.¹⁶ Instead, any substantive claim linking metaphysical relations to grounding, if justified, is justified via inference to the best explanation. Of course, the general thought that inference to the best explanation plays an important role in metaphysics (and philosophy in general) isn't new. Returning to the coherence of the exotic theses concerning causal powers and composition discussed earlier, these theses themselves (and their more down to earth competitors) are to be defended or rejected via inference to the best explanation.¹⁷

Okay, so why think it's plausible on abductive grounds that if REALIZATION is true then the F_{AS} ground F_B? While I think there are probably other ways of motivating this idea, here's the proposal I'll focus on. Let NOTHING be a proposition specifying the subclass of

¹⁵ See Mellor and Oliver 1997 for more on the sparse/abundant distinction. Given that quiddistic contingentism is coherent, for any transparent explanation whose explanandum concerns sparse properties, realization isn't the metaphysical relation in virtue of which the explanation is transparent.

¹⁶ While flatworldism is coherent, if what I said above about the a priori condition for transparency is right then the view that nothing stands in any metaphysical relations to anything is incoherent, provided that ATOM and BOILING are true.

¹⁷ See Armstrong 1983, Ch. 5 and Lewis & Lewis 1970 for early discussions of abduction in metaphysics; see Paul 2012 and Schaffer 2021 for recent detailed defenses.

metaphysical relations—the *metaphysical relations**—that carry *nothing-over-and-above-ness*. Plausibly, realization is a metaphysical relation*. While realization is irreflexive (as all metaphysical relations are irreflexive), there is nonetheless a sense in which the realized are “nothing over and above” what realizes them. Intuitively, the idea is one of *containment*—the realized are somehow contained within their realizers. (More on metaphysical relations that aren’t metaphysical relations* shortly.) NOTHING is the aspect of our evidence related to REALIZATION that I wish to focus on.

One way to understand the nothing-over-and-above idea is in terms of ontological cost. The general thought is this: supposing that x is nothing over above y , if you’re already committed to y , then you do not, in acknowledging that x exists, incur a further commitment (Lewis 1991). But to have a “further commitment” here isn’t to have a *distinct* commitment, as in this case nothing-over-and-above-ness would require identity. Instead, to have a further commitment requires that the commitment have an additional cost with respect to your other commitments, which is potentially relevant to the ontological economy of your theory compared to others (Cameron 2014; Schaffer 2015; Trogdon and Witmer 2021).¹⁸

So now our question is this: what is the best explanation of NOTHING? In other words, for each metaphysical relation that is a metaphysical relation*, why does that relation carry nothing-over-and-above-ness?¹⁹ I propose that what best explains NOTHING is what I’ll call CONNECTION, a proposition specifying how each of the metaphysical relations* are yoked with grounding such that instances of the former and latter point in the same direction. Just how the connection runs in any case will depend on the metaphysical relation* at issue, so think of CONNECTION as encoding how things work for each such relation. In the case of realization, the connection runs as follows: if instances of the X s (where x has the X s) realize an instance of Y (where y has Y), then the fact that the x s have the X s grounds the fact that y has Y . If CONNECTION is true then, given REALIZATION, the F_A s (the facts represented by ATOM) ground F_B (the fact represented by BOILING).

Here’s how the explanation works. Plausibly, nothing-over-and-above-ness isn’t restricted to particular ontological categories. Chalmers, for example, focuses in this context on *facts*—he supposes that biological facts are nothing over above microphysical facts in stating that “They may be *different* facts (a fact about elephants is not a microphysical fact), but [the former] are not *further* facts [compared to the latter]” (1996, 41; Chalmers’ emphasis). Now, while grounding is irreflexive, if some facts (fully) ground another fact, then the latter is nothing over above the former.²⁰ And if some fact is nothing over and above some collection of facts, then entities the former concerns

¹⁸ For early discussions of the nothing-over-above-ness, see Wisdom 1931 and later Wiggins 1968.

¹⁹ For any metaphysical relation, R , that is a metaphysical relation*, perhaps it’s a priori that R carries nothing-over-and-above-ness. This, however, is compatible with the idea that there is an explanation of why R has this feature.

²⁰ See Trogdon and Witmer 2021 for relevant discussion. A number of grounding theorists make similar claims, describing facts as being an ‘ontological free lunch’ or ‘constituted by’ the facts that ground them (see, e.g., Fine 2001 & 2012; Schaffer 2009).

are nothing over and above the entities the later concerns.²¹ So, the thought is that CONNECTION, if true, explains NOTHING, as in this case the nothing-over-and-above-ness characteristic of the metaphysical relations* plausibly is inherited, as it were, from the nothing-over-and-above-ness characteristic of grounding. Returning to the realization case, the realized in this case is nothing over and above what realizes it because F_M (a fact concerning the former) is grounding by, and hence nothing over and above, the F_{AS} (facts concerning the latter). Call this explanation of NOTHING in terms of CONNECTION the grounding explanation of NOTHING.²²

Now, I concede that it's not an entirely straightforward matter that the grounding explanation is, all things considered, the *best* explanation of NOTHING. But some initial reflection suggests that it will compare well to alternative proposals. Note that the grounding explanation is more unified than any approach according to which each metaphysical relation* is such that different properties or relations explain why it carries nothing-over-and-above-ness. So, at least compared to this class of potential explanations of NOTHING, the grounding explanation scores well with respect to the theoretical virtues of simplicity and unification.

I suggested above that the explanation of BOILING in terms of ATOM is transparent given the hybrid characterization of transparency. By contrast, it seems that, where PHYSICAL is the explanans of any potential physical explanation of PAIN, the relevant a priori and abductive factors don't work together in the right way for this explanation to be transparent. Given the hybrid characterization of transparency, PHYSICAL is transparently linked to PAIN just in case there is some metaphysical relation, R , that satisfies two conditions. The first is the a priori condition: it's a priori that, if PHYSICAL and PAIN are true, then PHYSICAL-entities are R -related to PAIN-entities in thus-and-so manner. The second is the abductive condition: it's plausible on abductive grounds that if PHYSICAL-

²¹ Here I speak of *facts* concerning entities, while above I speak of *propositions* concerning entities. We can potentially define one of these notions in terms of the other (e.g., while a fact may concern some entity tout court, for a proposition to concern an entity is for it to represent a fact that concerns that entity tout court. As facts are non-representational (see note 1), perhaps we can understand what it is for x to concern y tout court in terms of some non-extensional mode of composition.

²² Let NOTHING* be the proposition that encodes, for each metaphysical relation*, R , information about nothing-over-and-above-ness between facts concerning the relata of R . In the case of realization, rather than saying (as CONNECTION does) that if instances of the Xs (where x has the Xs) realize an instance of Y (where y has Y) then the fact that the x s have the Xs grounds the fact that y has Y , NOTHING* says that, if the Xs realize the Ys , then the fact that y has Y is nothing over and above the fact that the x s have the Xs . Above I in effect argued that CONNECTION explains NOTHING, as CONNECTION explains NOTHING*, and NOTHING* in turn explains NOTHING. This means that, if you're skeptical of grounding, then you might consider excising it from the hybrid characterization in the following way: for any explanation of A in terms of Δ , Δ is transparently linked to A just in case the a priori condition for above is met and it's plausible on abductive grounds that if Δ -entities are R -related to A -entities in thus-and-so manner, then the fact represented by A is nothing over and above the facts represented by Δ . See note 26 for more on this grounding-free version of the hybrid characterization.

entities are *R*-related to PAIN-entities in this manner, then the F_P s (the facts represented by PHYSICAL) ground F_M (the fact represented by PAIN, where '*M*' is for mental). There apparently is no metaphysical relation that meets both of these conditions—any such relation that satisfies one condition fails to satisfy the other.

Let's start with a metaphysical relation that seems to satisfy the abductive condition.²³ Let the *P*s be the physical properties that PHYSICAL concerns, and let *M* be the phenomenal property that PAIN concerns. (*N*, you will remember, is the property at issue in the theoretical identity claim we consider earlier, namely that pain is *N*—PAIN is the proposition that the *N*-instance has *M*. While *N* is instantiated by brains or relatively large proper parts of brains, the *P*s presumably are instantiated by smaller components, ones that perhaps aren't usefully described in neuroscientific terms.) Given our discussion of the explanation of BOILING in terms of ATOM, we should say that it's plausible on abductive grounds that if the *P*-instances realize the *M*-instance, then PHYSICAL grounds PAIN. Realization, however, clearly doesn't satisfy the a priori condition—it's not a priori that if PHYSICAL and PAIN are true, then *P*-instances realize the *M*-instance. Why not?

Recalling our discussion about realization and functional properties above, property instances realize an *M*-instance only if *M* is a functional property, a property whose nature is exhausted by its causal role. Let a functional concept be a concept that characterizes its referent solely in terms of its causal role. Plausibly, if a functional concept refers to some property, then it refers to a functional property in particular. Now consider the following conditional: if PHYSICAL and PAIN are true, then the *P*-instances realize the *M*-instance. Plausibly, this conditional is a priori only if the concept of the phenomenal character of pains that we deploy in this context is a functional concept. But, while our ordinary concept of, say, a gene is a functional concept, the concept of *M* that we deploy in this context (and phenomenal concepts in general) isn't functional in nature (Chalmers 1995, p. 8).²⁴

Now let's turn to a metaphysical relation that apparently satisfies the a priori condition. Let COMPOSITION be the proposition that the *P*-instances compose the *N*-instance. Perhaps it's a priori that if PHYSICAL and PAIN are true then COMPOSITION is true. I want to suggest that abduction doesn't license the claim that if COMPOSITION is true then the F_P s ground F_M . But what about CONNECTION, the proposition that encodes just how each

²³ In considering which metaphysical relations* might satisfy the abductive condition in this case, we need consider only such relations that are plausibly instantiated in our case. As for metaphysical relations* we can safely set aside, set formation is one, as clearly no PHYSICAL-entity stands in the set formation relation to any PAIN-entity.

²⁴ What about the determinable-determinate relation? Let's say that if *Y* is a determinable with respect to the *X*s (so having the *X*s is a way of having *Y*) and the *X*s are instantiated, then the *X*-instances determine* a *Y*-instance on this occasion. It's plausible on abductive grounds that if the *P*-instances determine* the *M*-instance then PHYSICAL grounds PAIN. (The reasoning here is the same as above: provided that the *P*-instances determine* the *M*-instance, CONNECTION entails that PHYSICAL grounds PAIN.) But determination* doesn't satisfy the a priori condition—it's not a priori that if PHYSICAL and PAIN are true, then the *P*-instances determine* the *M*-instance.

metaphysical relation* interacts with grounding? I've already suggested that CONNECTION, if true, is the best explanation of NOTHING. Also, if CONNECTION is true, then, given REALIZATION, the F_{AS} (the facts represented by ATOM) ground F_B (the fact represented by BOILING). Is it not then also the case that, if CONNECTION is true, then, given COMPOSITION, the F_{PS} (the facts represented by PHYSICAL) ground F_M (the fact represented by PAIN)?

You might think that on certain principled restrictions of composition the composed are nothing over and above what composes them (e.g., when the composed are mere aggregates) while on others this isn't the case (e.g., when the composed are integrated wholes). In this case, composition tout court isn't a metaphysical relation*, so CONNECTION and COMPOSITION together don't guarantee that the F_{PS} ground F_M . However, let's grant for the sake of argument that there is some principled restriction on composition, composition*, such that composition* is a metaphysical relation*. And let's suppose that it's a priori that if PHYSICAL and PAIN are true then COMPOSITION* is true, where COMPOSITION* is the proposition that the P -instances compose* the N -instance.²⁵

Still, it's not plausible on abductive grounds that if COMPOSITION* is true then the F_{PS} ground F_M . For think about what CONNECTION will and won't say about the connection between composition* and grounding. In the case of composition*, CONNECTION will say something like this: if the x s compose* y then the fact that the x s exist grounds the fact that y exists. CONNECTION *won't* say that if the x s compose* y then y has the phenomenal character of pains. Composition* is a *principled* restriction on composition; as such, it isn't merely, say, composition such that the composed have a particular phenomenal character. Hence, it won't say that the fact that the x s exist grounds the fact that y has M either. The phenomenal character of pains doesn't somehow figure into the nature of composition*—panpsychist musings aside, while putting things together can be tough, surely no principled restriction on composition is by its nature painful! The upshot is that, while CONNECTION and COMPOSITION* together guarantee that the fact that the P -instances exist ground the fact that the N -instance exists, they don't guarantee that the F_{PS} ground F_M .²⁶

²⁵ Let an entity be *simple* just in case it lacks proper parts and *extended* just in case it's a spatiotemporal entity and doesn't have the shape and size of a point (Gilmore 2018). You might claim that, for all we know a priori, N -instances are extended simples, which means that it isn't a priori that if PHYSICAL and PAIN are true, then the COMPOSITION* is true. In response, I grant that, for all we know a priori, there is something without proper parts with the shape and size of an N -instance. Moreover, it's a priori open that such an object occupies the fusion of the regions occupied by the P -instances. What I suspect is a priori closed, however, is that such an entity is an N -instance. The thought is that it's a priori that, for any, x = an N -instance only if x is complex.

²⁶ Returning to the grounding-free version of the hybrid characterization (see note 22), it's not plausible on abductive grounds that if COMPOSITION* is true then F_M is nothing over and above the F_{PS} . To keep things simple, let's just focus on what might best explain why it is that if the P -instances compose* the N -instance then the N -instance is nothing-over-and-above the P -instances (rather than what best explains NOTHING, a proposition that, together with the claim that the P -instances compose* the N -instance, entails that the N -instance is

Extrapolating from the discussion above, perhaps you agree that, given the hybrid account, any potential physical explanation of PAIN *that has actually been proposed* (e.g., neural theories, the global workspace theory, the integrated information theory) isn't transparent. But you might still wonder: do we have any reason to think that there is no such potential explanation, period?

I think we do. For consider what such an explanation would need to look like. The explanation wouldn't involve any metaphysical relations we're familiar with—familiar metaphysical relations distinct from realization and composition* either aren't plausibly instantiated in this case (see note 23) or they don't satisfy the a priori and abductive conditions (see note 24) or both. Such an explanation would therefore have to involve some know-not-what (i.e., some metaphysical relation that hitherto hasn't played a role in metaphysical theorizing) that is both plausibly instantiated in this case and meets the a priori and abductive conditions. It's of course epistemically possible that there is some relation that satisfies these conditions, but it seems unlikely.

4. Aporetic conclusion

I've argued that there is an epistemic asymmetry between the explanation of BOILING in terms of ATOM and the explanation of PAIN in terms of PHYSICAL—the former is transparent and the latter isn't given the hybrid characterization of transparency. Suppose I'm right about this. And suppose, more generally, that the contrast claim is true given the hybrid characterization. In this case, there is a special explanatory gap between the physical and the experiential in that potential physical explanations of why instances of physical properties have the phenomenal character they do (or any phenomenal character at all) aren't transparent, while other familiar explanations are transparent.

While physicalism (about the mental) hasn't been our focus, what implications might all of this have for the thesis? Unfortunately, I don't have a clear-cut answer to this question to offer here. This is because working out just what the implications are for physicalism requires that we address various complex issues regarding the metaphysics and epistemology of grounding and explanation, issues that go well beyond what I could hope to fully address here. Our work never ends! I will, however, briefly discuss two potential ways of linking up our discussion with physicalism and point to some complications for each of these approaches.

You might think that the truth of the contrast claim given the hybrid characterization of transparency is relevant to physicalism in the following way: physicalism requires there to be a particular sort of explanation of PAIN in terms of PHYSICAL, and a necessary condition for an explanation to be of the relevant sort is that it be transparent in the

nothing over and above the *P*-instances). In this case, we have two competing proposals about what best explains our target: (i) the fact that the *N*s exist is nothing over and above the fact that *P*s exist, and (ii) F_M is nothing over and above the F_P s. The idea that the phenomenal character of pains somehow plays an explanatory role here is pretty obscure—we should clearly go with the first proposal.

sense specified by the hybrid account. Hence, if the link between PHYSICAL and PAIN is indeed non-transparent, then physicalism is false.

What form of explanation might play this role? The most obvious candidate is *reductive* explanation, as Chalmers and others see physicalism as essentially the thesis that there are reductive explanations of propositions like PAIN (Chalmers 1996, Ch. 2; Chalmers and Jackson 2010; Chalmers 2012, p. 305). Any explanation that is reductive in the relevant sense satisfies two conditions. First, it's a *nothing-over-and-above explanation*—the fact represented by its explanandum is nothing over and above the fact represented by its explanans (Chalmers 1996, p. 41). Second, it's non-circular—this means that, for any potential physical explanation of PAIN, the explanans doesn't itself include propositions concerning mental properties (Chalmers 2003, p. 105). Proposal: any non-circular nothing-over-and-above explanation is transparent.

What should we make of this proposal? Compatible with the view is the idea that some explanations that satisfy the nothing-over-and-above condition but not the no circularity condition aren't transparent. Consider essentialist explanations—suppose that ESSENCE explains PAIN, where ESSENCE is the proposition that it's essential to the *N*-instance that it has *M*. Here we have a paradigm case of nothing-over-and-above-ness: the fact represented by PAIN is nothing over and above the fact represented by ESSENCE, as whatever it takes for the former to obtain is already “built into” the obtaining of the latter. This explanation clearly fails to satisfy the no circularity condition. And it's not transparent, as it fails to satisfy the a priori condition for transparency. There just is no metaphysical relation (and, hence, no metaphysical relation*) *R* such that if ESSENCE and PAIN are true then ESSENCE-entities are *R*-related to PAIN-entities. So, of course, there is no metaphysical relation *R* such that it's priori that if ESSENCE and PAIN are true then ESSENCE-entities are *R*-related to PAIN-entities. If the fact represented by ESSENCE grounds the fact represented by PAIN, this is an example of *bare* grounding which, for our purposes here, we can characterize as grounding where it's not the case that entities the grounds concern stand in metaphysical relations to entities the grounded concern (Trogdon 2018).

Also compatible with the proposal is the idea that some explanations that satisfy the no circularity condition but not the nothing-over-above condition are non-transparent. Consider simple causal explanations whose explanantia don't include covering laws. Some of these explanations are non-circular in the relevant sense, such as an explanation that cites a (non-mental) physical cause of a mental effect. The effect in this case is something over and above its cause, and this explanation, like the one above, fails to satisfy the a priori condition necessary for transparency. Again, in this case there just is no metaphysical relation *R* such that if explanans and explanandum are true then entities the former concerns stand in *R* to entities the latter concerns.

On the proposal we're considering, however, there are no explanations that satisfy both the nothing-over-and-above and no circularity conditions that are non-transparent. But, while I think this proposal is worth further thought, you might think that there are such explanations. Consider a grounding-theoretic take on Fine's (1994) account of necessity in terms of essence: for any truth of the form “It's necessary that...”, there is an explanation of it in terms of a truth of the form “It's essential to *x* is that...”, where the

fact represented by the latter grounds the fact represented by the former (Rosen 2010). These explanations satisfy the nothing-over-and-above conditions, as grounded facts are nothing over and above the facts that ground them. And they seem to satisfy the no circularity condition as well—the explanantia are modal-free. But they also seem to be cases of bare grounding. As such, it seems that they don't satisfy the a priori condition for transparency.

To take a different track, you might think that the truth of the contrast claim given the hybrid characterization is relevant to physicalism in the following way: important routes to justifying physicalism are closed off if any potential explanation of why an instance of a physical property has the phenomenal character it does isn't transparent. How might this work?

On a grounding-theoretic take on physicalism, the view is that any mental fact is grounded by (non-mental) physical facts, perhaps together with non-mental, non-physical facts (e.g., indexical facts). In this case, physicalism is true only if particular grounding claims are true, such as the claim that the F_{PS} (the facts represented by PHYSICAL) ground F_M (the fact represented by PAIN). But why think that the F_{PS} ground F_M ? In earlier work (Trogon 2018), I developed a mechanistic approach to the justification of grounding claims. There I argued that one way—and perhaps the best way—to justify a grounding claim is to show how the connection runs between the putative ground and grounded by specifying a “mechanism” that mediates their connection. This involves specifying metaphysical relations* that entities these facts concern stand in to one another. (Compare: one way—and perhaps the best way—to justify a causal claim is to specify a mechanism mediating the putative cause and effect.) Proposal: the mechanistic approach to justifying the claim that the F_{PS} ground F_M is viable only if the link between PHYSICAL and PAIN is transparent.

The thought behind the proposal is that, if it's plausible to maintain that a mechanism mediates the connection between the F_{PS} and F_M in such a way that the former grounds the latter, this is because of two things: there is some metaphysical relation R such that it's a priori that, if Δ and A are true, then Δ -entities are R -related to A -entities in thus-and-so manner; and it's plausible on abductive grounds that, if Δ -entities are R -related to A -entities in thus-and-so manner, then the facts represented by Δ ground the fact represented by A . If this is right, then perhaps the best way of justifying the grounding claims at issue with physicalism are off the table, provided that the contrast claim is true given the hybrid characterization of transparency. In this case, perhaps we should conclude that physicalism itself lacks proper justification.

You might think, however, that, while one way to implement the mechanistic approach to justifying the claim that the F_{PS} grounds F_M is via the mix of a priori and abductive considerations above, this isn't the only way. Suppose, as I argued above, that it's not a priori that if PHYSICAL and PAIN are true then the P -instances realize the M -instance. But this, of course, is compatible with it being plausible on *abductive* grounds that if PHYSICAL and PAIN are true then the P -instances realize the M -instance. Suppose for the moment that we can show that the conditional claim is plausible in this way. And suppose, as I suggested above, that it's also plausible on abductive grounds that if the P -instances realize the M -instance, then the F_{PS} grounds F_M . In this case, we've implemented the

mechanical approach to justifying the claim that the F_P s ground F_M , despite the fact that the link between PHYSICAL and PAIN isn't transparent.

So now the question is this: should we think that the P -instances realize the M -instance on the basis of the fact that that it's entailed by the best explanation, among a competitive list of potential explanations, of some aspect of our evidence related to the truth of PHYSICAL and PAIN? This matter is worth further thought.²⁷

References

*Page numbers in the citations above refer to reprints when applicable.

Armstrong, D.M. 1983. *What is a Law of Nature?* Cambridge University Press.

Bennett, K. 2011. "By Our Bootstraps," *Philosophical Perspectives* 25: 27–41.

Block, N. and R. Stalnaker. 1999. "Conceptual Analysis, Dualism, and the Explanatory Gap," *Philosophical Review* 108: 1–46.

Cameron, R. 2014. "Parts Generate the Whole, But They Are Not Identical to It." In A. J. Cotnoir and D. Baxter (eds.), *Composition as Identity*. Oxford University Press.

Chalmers, D. 1995. "Facing Up to the Problem of Consciousness," *Journal of Consciousness Studies* 2: 200–219. Reprinted in Chalmers 2010, pp. 3–27.

---- 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

---- 2003. "Consciousness and Its Place in Nature." In S. Stich & T. Warfield (eds.), *Blackwell Guide to the Philosophy of Mind*. Blackwell. Reprinted in Chalmers 2010, pp. 103–140.

---- 2009. "Ontological Anti-Realism." In D. Chalmers et al (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.

---- 2010. *The Character of Consciousness*. Oxford University Press.

---- 2012. *Constructing the World*. Oxford University Press.

Chalmers, D. and F. Jackson. 2001. "Conceptual Analysis and Reductive Explanation," *Philosophical Review* 110: 315–361. Reprinted in Chalmers 2010, pp. 207–247.

²⁷ I wish to thank audiences at the University of Birmingham, University of Gothenburg, and Virginia Tech for their helpful comments and suggestions. Thanks to Tim Fuller, Dan Hoek, Joseph Levine, Jonathan Schaffer, Gabe Rabin, and anonymous referees for comments on earlier drafts.

- Fine, K. 1994. "Essence and Modality," *Philosophical Perspectives* 8: 1–16.
- 2001. "The Question of Realism," *Philosophers' Imprint* 1: 1–30.
- 2012. "Guide to Ground." In F. Correia and B. Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge University Press.
- Gilmore, G. 2018. "Location and Mereology." In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
- Levine, J. 1983. "Materialism and Qualia: The Explanatory Gap," *Pacific Philosophical Quarterly* 64: 354–361.
- 1993. "On Leaving Out What It's Like." In M. Davies and G. Humphreys (eds.), *Consciousness: Psychological and Philosophical Essays*. Blackwell.
- 1998. "Conceivability and the Metaphysics of Mind," *Nous* 32: 449–480.
- 2001. *Purple Haze*. Oxford University Press.
- 2006. "Phenomenal Concepts and the Materialist Constraint." In T. Alter & S. Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Reprinted in *Quality and Content* (2018), pp. 13–34. Oxford University Press.
- 2018. *Quality and Content: Essays on Consciousness, Representation, and Modality*. Oxford University Press.
- Lewis, D. 1991. *Parts of Classes*. Oxford: Blackwell.
- Lewis, D. and S. Lewis 1970. "Holes," *Australasian Journal of Philosophy* 48: 206–12.
- Lipton, P. 1991. *Inference to the Best Explanation*. Routledge.
- Mellor, D. H. and A. Oliver. 1997. "Introduction." In D.H. Mellor and A. Oliver (eds.), *Properties*. Oxford University Press.
- Merricks, T. 2001. *Objects and Persons*. Oxford University Press.
- Paul, L. 2012. "Metaphysics and Modeling: The Handmaiden's Tale," *Philosophical Studies* 160: 1–29.
- Rabin, G. 2019. "Grounding the Gaps or Bumping the Rug?" *Journal of Consciousness Studies* 26: 191–203.
- Rosen, G. 2010. "Metaphysical Dependence: Grounding and Reduction." In R. Hale and A. Hoffman (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press.

- Rosen, G. and C. Dorr 2002. "Composition as a Fiction." In R. Gale (ed.), *The Blackwell Companion to Metaphysics*. Blackwell.
- Sassarini, E. Forthcoming. "No Ground to Bridge the Gap," *Synthese*.
- Schaffer, J. 2005. "Quiddistic Knowledge," *Philosophical Studies* 123: 1–32
- 2009. "On What Grounds What." In D. Chalmers et al (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- 2015. "What Not to Multiply without Necessity," *Australasian Journal of Philosophy* 93: 644–664.
- 2017. "The Ground Between the Gaps," *Philosophers' Imprint* 17: 1–26
- 2021. "Grounding Functionalism." In U. Kriegel (ed.), *Oxford Studies in Philosophy of Mind*, Vol. 1. Oxford University Press.
- Sturgeon, S. 1994. "The Epistemic Basis of Subjectivity," *Journal of Philosophy* 91: 221–35.
- Tallant, J. 2014. "Against Mereological Nihilism," *Synthese* 191: 1511–1527.
- Trogon, K. 2013. "Grounding: Necessary or Contingent?" *Pacific Philosophical Quarterly* 94: 465–485.
- 2018. "Grounding-Mechanical Explanation," *Philosophical Studies* 175: 1289–1309.
- Trogon, K. & D. G. Witmer. 2021. "Full and Partial Grounding," *Journal of the American Philosophical Association* 7: 252–271.
- Wiggins, D. 1968. "On Being in the Same Place at the Same Time," *Philosophical Review* 77: 90–95.
- Wisdom, J. 1931. "Logical Constructions (I)," *Mind* 40: 188–216.