

Knowing the Powers

Wolfgang Schwarz*

22 June 2021

1 Introduction

If the world contains primitive modal elements – irreducible laws, powers, potentialities, or propensities –, how could we know about them? Humeans have long worried that we could not. Put bluntly, the worry is that observation and experiment only tell us what *does* happen; they don't directly reveal what *might* or *must* or *would* happen under non-actual circumstances. If modal phenomena are reducible to facts about occurrent events, then it is no surprise that observing occurrent events can give us information about modality. By contrast, knowledge of *primitive* modality seems to require an inexplicable leap from observations of one kind of fact to conclusions about an entirely different kind of fact.

Let's call this the *access problem* for non-Humean accounts of (natural) modality. It is a sub-problem of a more general problem, often highlighted by Lewis. As Lewis pointed out, it is not enough to posit “unHumean whatnots” [Lewis 1994: 239]. We also need a credible story of how these whatnots could play the familiar roles of laws, powers, potentialities, or propensities. Part of that role concerns the methods by which these things can be discovered.

In its blunt formulation, the access problem is easy to resist. Arguably, some modal facts *can* be directly observed: we can see that a surface is slippery, or that its slipperiness causes a fall. In other cases, inference to the best explanation might be hoped to get us from, say, an observation of frequencies to hypotheses about chance.

Like many Humeans, I am not convinced by these replies. I don't think they get to the heart of the access problem. My aim in this essay is to explain why.

In a nutshell, I will argue that if the world has primitive modal elements, then there are different *a priori* conceivable ways in which these elements might be arranged, many of which are compatible with our total history of perceptual experience. If, for example, certain objects in our environment have some primitive power F , then there are conceivable (although perhaps metaphysically impossible) scenarios in which the objects display the same observable behaviour even though they have a different power F^* . Since the two kinds of scenario deserve equal *a priori* credence, and are equally compatible with our perceptual experiences, we can't know that we are in an F scenario.

* Thanks to Jenann Ismael, Sigfried Jaag, Barbara Vetter, and members of the Kollegforschungsgruppe “Human Abilities” at HU Berlin for helpful comments on an earlier version.

I will outline three lines of response. One is to concede that we can never know which unHumean whatnots are present in our surroundings – except in a certain shallow sense of ‘know’. Another is to posit implausible epistemic norms that would allow us to favour F scenarios over F^* scenarios without any relevant evidence. A third option, I will argue, is to re-construe the anti-Humean position as a doctrine about ideology rather than ontology. This helps with some forms of the access problem, but others remain. For example, difficult questions would still arise for our knowledge of objective chance.

2 The access problem

Let me try to explain what I see as the heart of the access problem.

I will assume without argument that we have abundant knowledge of natural modality. We know that nothing can travel faster than the speed of light, that electrons are disposed to repel one another, that ordinary wine glasses are fragile, and that a properly minted coin has a 50% chance of landing heads when tossed.

I will also assume that we know (at least roughly) how we know these facts. We may discover that a glass is fragile by observing its delicate build, or by watching it crack under pressure. We may discover that nothing can travel faster than the speed of light by noticing that this is implied by a theory which in turn is supported by a host of theoretical and empirical considerations.

Our question is therefore not how – by what methods – we can discover modal facts. Rather, the question is whether a certain metaphysics of modality can make sense of these methods.

To illustrate, consider a crazy view which identifies fragility with the property of having been touched by an angel. A glass’s fragility, on this view, does not imply anything about its physical composition or about whether it will break when struck. (There is nothing incoherent about a scenario in which, say, a glass breaks under light pressure even though it has never been touched by an angel.) As a consequence, it becomes mysterious how observing the glass’s build, or its behaviour under stress, could reveal anything about its fragility – that is, about whether it has been touched by an angel.

Suppose we throw the glass against the wall and see it shatter. To get from this observation to the conclusion that the glass has been touched by an angel, we would need “bridge principles” according to which, say, objects that shatter have usually been touched by an angel. But how could we have discovered these principles, without some other means of knowing whether something has been touched by an angel?

It is not a response to say that on the crazy theory in question, having been touched by an angel is identical to fragility, and that surely we can find out that a glass is (or was) fragile by watching it shatter. The problem is that having been touched by an angel *cannot be* identical to fragility, because the assumed identity would make a mystery of

the methods by which we discover fragility. To assume that fragility equals past angelic touch is to presuppose what is in question: that past angelic touch can play the role of fragility.

Now consider a (still somewhat crazy) anti-Humean view on which fragility is a metaphysically primitive property – a special unHumean property with a “dispositional essence”.¹ Again, we may ask why observation of a glass’s shattering should tell us anything about this primitive property. We would, it seems, have to assume a bridge principle linking the primitive property with the observed events. But how could we have discovered that principle?

Again, it is not a response to insist that the posited primitive property is identical to fragility, thereby presupposing what is in question. For the vast majority of properties, you can’t find out whether a glass has them by watching it shatter. An identification of fragility with some property *X* is only plausible if there are *independent* reasons to believe that the way we find out about fragility is a way to find out about *X*, without already assuming that *X* is fragility. If *X* is a primitive whatnot, it is hard to see what these independent reasons might be.

The problem arises not just for friends of primitive powers. Whenever anti-Humeans identify a modal phenomenon with the presence of some unHumean whatnot (an irreducible law, a higher-order universal, or whatever), they owe us an explanation of how our methods for identifying the phenomenon could serve as methods for identifying the unHumean whatnot.

3 Permutations

You may think such an explanation is not hard to find. Suppose we inspect a glass and come to believe that it is fragile. If all goes well, our judgement is true, and appropriately caused by the glass’s fragility. If so, we *see* that the glass is fragile, and thereby come to know that it is fragile. What is needed to ensure our knowledge, then, is that whatever property is identified with fragility stands in the appropriate causal relation to our judgement. And why shouldn’t a primitive power satisfy that condition? What’s the problem?

There are, in fact, several problems. It will take a while to sort them all out.

Let’s briefly set aside modality and think about diagnostics. Imagine a patient consulting a doctor, with symptoms of arthritis. An X-ray confirms this suspicion. Looking at the X-ray image, the doctor can *see* that the patient has arthritis. But this

¹ This is still somewhat crazy because fragility is plausibly not fundamental, even on anti-Humeans accounts. However, the issues I will discuss would also arise for more plausible candidates for primitive dispositions – negative unit charge, perhaps – so I will stick with the familiar case of fragility.

does not mean that any remaining epistemological questions about her knowledge can be delegated to the physiology of visual perception.

For one thing, the doctor's training and background knowledge play a role. An amateur like me could not see that the patient has arthritis by looking at the X-ray. Moreover, the X-ray does not provide *conclusive* evidence for the diagnosis. It is easy to think of scenarios in which the patient does not have arthritis, even though the doctor's perceptual evidence is the just same: the patient might have an unknown disease that looks like arthritis in an X-ray; the X-ray technician might have accidentally swapped the patient's X-ray image with that of an earlier patient; and so on.

We can nonetheless describe the doctor as knowing (and seeing) that the patient has arthritis because – roughly speaking – these “error scenarios” (in which the patient does not have arthritis but the X-ray image looks the same) are (a) in some sense remote, (b) the doctor gives them little credence, and (c) she is right to give them little credence.

If *I* were to look at the X-ray image, condition (a) might still be satisfied, but the corresponding conditions (b) and (c) would fail. I would not be confident that the patient has arthritis, and even if I were, my confidence would be irrational. That's why I couldn't know (or see) that the patient has arthritis just by looking at the X-ray.

Now back to the view that fragility is some primitive property *F*. Let's grant that the presence of *F* (partially) causes our perceptual experience when we inspect the glass, and thereby our judgement that the glass is fragile. However, the experience will not be conclusive evidence for the presence of *F*. We can distinguish different kinds of error scenarios.

First, there are scenarios we would intuitively describe as situations in which it appears to us as if we are inspecting a fragile glass, but in reality something else is going on. Perhaps we are being deceived by an evil demon. Or perhaps we are inspecting a special glass that looks fragile but is actually unbreakable. Even if we've thrown the glass against the wall and saw it shatter, the glass might only have become fragile as it hit the wall, so that we would be wrong to conclude that it was fragile before. In normal situations, scenarios like these can be regarded as remote, and don't stand in the way of knowledge. Let's set them aside.

Here is a different kind of error scenario. Consider a world that is much like the actual world except that the primitive property *F* has been replaced (in all its instances) by a different primitive property *F** – a property that (let's say) nothing has in the actual world. Anything that is, in our world, caused by the glass being *F* is here caused by the glass being *F**. In particular, it is *F**, not *F*, that (partially) causes our experiences as we inspect the glass.

By assumption, the two kinds of scenario do not differ in the light waves that arrive at our retina, in our subsequent brain states, or in the phenomenology of our experience.

Our experience therefore does not put us in a position to rule out the F^* hypothesis.²

We also can't dismiss the F^* scenarios as far-fetched. In an F scenario, the glass has one primitive property; in the corresponding F^* scenarios, the glass has a different primitive property. *A priori*, the two scenarios are on a par. Any *a priori* preference in favour of either scenario would be arbitrary and irrational.

But if there is no *a priori* reason to favour one scenario over another, and our perceptual experiences can't distinguish between the two scenarios – insofar as we have the same (kind of) experiences in both scenarios – then both scenarios deserve roughly equal (posterior) credence. So we can't know that we're in an F scenario.

I find this argument persuasive. If you agree, feel free to skip the next section, in which I will investigate an argument for the opposite conclusion.

4 Deep knowledge and shallow knowledge

The argument for the opposite conclusion begins with a semantic (or conceptual) ascent. Arguably, the word (or concept) 'fragile' functions in a semi-demonstrative way to pick out a certain property in the world around us – a property that causes things to break when struck.³ We can dismiss as far-fetched scenarios in which the term fails to pick out anything. So we can be confident that 'fragile' applies to some things in the world around us. That is, we can be confident that some things around us are fragile.

Now, if 'fragile' picks out F , then to be confident that some things are fragile is to be confident that some things are F . Any scenario in which F is replaced by an alien property F^* is a scenario in which nothing is fragile. So we can be confident that we are not in an F^* scenario.

Let's call this argument *the verbal trick*. To see what's wrong with it, let's return to our doctor. While we were talking about fragility, another patient has arrived whose symptoms are compatible with both arthritis and fibromyalgia. This time, the X-ray image comes back blurry, and the doctor remains unsure if the patient has arthritis or fibromyalgia.

Assume the patient actually has arthritis. The doctor might then use the following trick to identify his ailment. First, she introduces a new name, say 'julitis', for the patient's disease. Having paid attention to this act of baptism, she is rationally confident

² You might object that if F is a more plausible candidate for a primitive property (see footnote 1 above), then swapping F by F^* will make a difference to our brain state and thereby to our experiences. Alternatively, you might suggest that our experience is directly sensitive to what primitive properties are present in our environment. But relations of causal or metaphysical dependence aren't pertinent here. Suppose we were initially unsure about whether our environment contains instances of F or F^* . The question is whether our perceptual experience alone would then put us in a position to conclusively rule out the F^* hypothesis. The answer is no.

³ Remember that 'fragile' is a placeholder for something like 'negative unit charge'.

that the patient has julitis. She can ignore scenarios in which the patient does not have julitis. But ‘julitis’ denotes arthritis (and not fibromyalgia): julitis and arthritis are the same disease. Any scenario in which the patient has fibromyalgia is a scenario in which the patient does not have julitis. So the doctor can be confident that her patient has arthritis, and not fibromyalgia.

We are dealing with a “Frege case”. Arthritis is known to the doctor under two modes of presentation: as the illness called ‘arthritis’ about which she learned in medical school, and as the illness called ‘julitis’ that causes her present patient’s symptoms.

Fregeans hold that these modes of presentation affect the truth-conditions of attitude reports: to know (or be confident) that the patient has *arthritis*, the doctor would have to know (or be confident) that the patient has the illness called ‘arthritis’ about which she learned in medical school. Evidently, the doctor did not acquire any such knowledge through her introduction of the name ‘julitis’. Fregeans would conclude that the doctor may know that the patient has julitis, but not that he has arthritis.

Fregeans will spot the same mistake in the above verbal trick. We may be confident that some things in the world are fragile, but it doesn’t follow that we can be confident that some things are *F*, even on the hypothesis that ‘fragile’ picks out *F*.

The Fregean account of attitude reports is controversial. Let’s assume that it is false, so that we can truly report the doctor as knowing that her patient has arthritis.

Nonetheless, it should be uncontroversial that the doctor has not gained any useful information through her introduction of a new name. She has made no genuine epistemic progress with respect to her patient’s disease. She is in no better position to prescribe a treatment. Medical textbooks rightly do not mention the introduction of new names as a diagnostic method.

I will say that the doctor has *shallow knowledge* that the patient has arthritis, but lacks *deep knowledge* of the same fact. Perhaps the English word ‘knowledge’ means shallow knowledge. I doubt it, but it doesn’t matter. I’m interested in deep knowledge – a kind of knowledge that tracks genuine epistemic progress.

Deep knowledge of an empirical hypothesis *H* requires a history of perceptual experience that favours *H* over its alternatives – in the sense that among scenarios in which you have these experiences, *H* scenarios have (significantly) greater *a priori* credibility than $\neg H$ scenarios.

For example, consider your past and present experiences related to your hands. There are scenarios in which you have the same⁴ history of perceptual experiences even though you don’t have hands. (You might be a brain in a vat.) But these scenarios deserve much lower *a priori* credence than scenarios in which you have the experiences and you also

⁴ What does it mean that your experiences in two scenarios are “the same”? Good question. Roughly, the experiences should have the same phenomenology. The full answer, I think, is complicated – see [Schwarz 2018].

have hands. Provided that the other conditions for knowledge are satisfied, you may therefore have deep knowledge that you have hands.

By comparison, the doctor's history of perceptual experiences does not (significantly) favour scenarios in which her patient has arthritis over scenarios in which he has fibromyalgia. There are scenarios in which she has all the same experiences (her experience of listening to the patient's description of his symptoms, the blurry X-ray image, her introduction of the name 'julius', etc.), but in which the patient has fibromyalgia. These scenarios do not deserve much lower credence than corresponding scenarios in which the patient has arthritis. So the doctor does not have deep knowledge that the patient has arthritis.

When it comes to distinguishing F and F^* , we are in the same position as the doctor with arthritis and fibromyalgia. We may be able to pick out F demonstratively, and we may have a word for it, but we don't have deep knowledge that anything around us is F .

5 Humean knowledge

My permutation argument from section 3 resembles a well-known argument purporting to show that we could never discover the identity of fundamental (categorical) properties in a Lewisian metaphysics – for how could we tell apart scenarios that merely differ by swapping these properties? Some have, in effect, responded that we might still have shallow knowledge about fundamental properties (e.g. [Langton 2004], [Schaffer 2005]). But I think Lewis [2009] was right to accept the argument for deep knowledge.⁵ Like Lewis, I do not find the conclusion especially problematic. It's not like we all thought we knew the relevant facts, and then Lewis tells us that we don't. On Lewis's view, we can't even state or entertain the propositions of which we are ignorant (see [Kelly 2013], [Dasgupta 2015]).

Anti-Humeans might adopt a similar response to my permutation argument. They might concede that we can't have deep knowledge of whether a glass (or electron) has a specific power F . But this is not how anti-Humeans usually present their view. The irreducibly modal facts they posit are supposed to be familiar facts about dispositions, potentialities, or laws – facts that aren't beyond our epistemic reach.

Anti-Humeans might also complain that my constraints on "deep knowledge" are too demanding. Can *Humeans* explain our deep knowledge of the glass's fragility?

Let's see. On a typical Humean analysis, a glass is fragile iff (roughly) it has a material structure which, together with the laws of nature, entails that it is likely to break under moderate stress, where the relevant laws are certain regularities in the history of the universe. One might reasonably wonder how we could find out that a glass satisfies this

⁵ Lewis held that our ordinary concept of knowledge is a concept of deep knowledge: see [Lewis 1996].

condition (of having a material structure etc.) simply by watching it shatter, given that the condition requires suitable regularities in the entire history of the universe.

What would the relevant error scenarios look like? They would be scenarios in which the glass shatters upon being thrown at the wall, but in which this event is not an instance of any general regularity in the history of the world.

Now we have other experiences to draw on. These other experiences suggest that the dynamics of physical systems, at least to the extent that we have observed them, is fairly predictable. Not only have we often seen delicately built objects break under stress, we have noticed that similar physical systems of all sorts generally respond in similar ways when put in similar conditions. It is still *possible* to have all these experiences in a world without relevant dynamical regularities. There are, for instance, scenarios in which the dynamics of the systems we have observed is regular, but the dynamics of unobserved systems is entirely irregular. But scenarios like these deserve little *a priori* credence.⁶

On the Humean account, non-trivial knowledge of natural modality always involves an element of conjecture. But at least our experiences often favour some modal hypotheses over others – assuming that we may treat some scenarios as *a priori* more credible than others. If we can't assume that the regularities in the observed part of the world still hold in the unobserved part⁷, the Humean epistemology is doomed (as [Ismael 2020] notes). But so is everyone else's. I will return to this point in section 10.

6 Ontology and ideology

It is time to revisit a premise of my argument in section 2. A friend of primitive fragility, I claimed, would need “bridge principles” to connect the hypothesis that a glass is F with the observation that it shattered upon being thrown against the wall. Similarly, in section 3, I claimed that when we see the glass shatter, our experience is neutral between scenarios in which the glass is F and scenarios in which it has a different disposition F^* .

This may seem strange. According to our (imaginary) friend of primitive fragility, F is, by its very essence, a disposition to break under stress. Dispositions are not independent of their manifestations. Why, then, should we need bridge principles connecting the disposition with its manifestation? Why should we allow for scenarios in which a different disposition F^* – a disposition to *glow* under stress, perhaps? – has traded places with fragility, causing things to break under stress? Such a scenario makes no sense.

⁶ I am not suggesting that this story mirrors some kind of inference we are supposed to make. On a cognitive level, it may well be that we come to believe that the glass is fragile because we trust our experience, which “presents it” as fragile. See [Beebe 2003] for how Humeans might account for this kind of presentation.

⁷ To make this assumption precise, we would need a criterion for distinguishing gruesome regularities from genuine regularities. Humeans disagree on what that criterion should look like.

I agree that we can't arbitrarily swap an object's dispositions while holding fixed the manifestations. But our question is whether dispositions can be identified with primitive properties. We should not presuppose a positive answer by assuming that these primitive properties can be referred to as 'the disposition to break under stress' or 'the disposition to glow under stress'. That's why I have used the neutral names ' F ' and ' F^* '.

Different versions of the permutation problem arise for different types of anti-Humeanism. On some views (e.g. [Heil 2010], [Williams 2019]), fundamental powers have both a dispositional character and a non-dispositional, qualitative character. A relevant F^* scenario for the permutation argument might then be a scenario in which the qualitative character of F is swapped with the qualitative character of some other power. Scenarios like these may be deemed metaphysically impossible: the pairing between qualitative and dispositional characters is supposed to be metaphysically necessary. But I never said that the F^* scenario is metaphysically possible. I only assumed that it can't be ruled out *a priori*.

Others (e.g. [Bird 2007], [Mumford and Anjum 2011]) deny that fundamental powers have a qualitative aspect. Here the posited fundamental properties are assumed to have an entirely modal essence, an essence that grounds counterfactuals and other modal truths. A relevant F^* scenario might then be a scenario in which a different fundamental property grounds the same modal truths. Perhaps it is a scenario in which a different property F^* shares F 's essence. (Again, we need not assume that such scenarios are metaphysically possible, as long as they can't be ruled out *a priori*.) Alternatively, it could be a scenario in which F is swapped for a property with an entirely different modal essence – an essence, perhaps, that grounds counterfactuals about glowing when struck.

I admit that I have trouble understanding talk about essence and grounding. But I have been told that grounding is an objective relation that need not be epistemically transparent: X can ground Y even if there is no *a priori* connection between X and Y . This suggests that if F^* is a property that grounds counterfactuals about glowing when struck, there can still be *a priori* conceivable scenarios in which F^* is present even though the relevant object wouldn't glow (but rather break) when struck.

To escape the permutation argument, we would need not just a metaphysical, but an *a priori* connection between the presence of F and counterfactuals about breaking. Ideally, we would have a connection that goes both ways, so that the observed breaking of a glass is evidence for the presence of F . This can be made to work, but it might require rethinking the metaphysics of powers.

Remember the Humean project of analysing counterfactuals in ultimately non-modal terms. The project has not been a resounding success. One might reasonably hold that it will never succeed. More strongly, one might hold that there is no way to determine the truth-value of (arbitrary) counterfactuals from suitably different propositions.

Even if counterfactuals are in this sense primitive, the truth-value of a counterfactual

may still be *constrained* by other propositions. For example, it is widely held that a counterfactual $A > C$ entails the corresponding material conditional $A \supset C$ (and so the falsity of $A \supset C$ entails the falsity of $A > C$). Accepting this entailment does not commit us to the Humean reductive project. $A > C$ can entail $A \supset C$ even if $A > C$ is primitive.

When I say that $A > C$ entails $A \supset C$, I don't mean that some opaque metaphysical relation holds between $A > C$ facts and $A \supset C$ facts. I mean that *a priori* reasoning is enough to rule out any putative scenario in which $A > C$ is true and $A \supset C$ false. We don't need empirical bridge principles to infer $A \supset C$ from $A > C$.

Now return to the shattering glass. Let's assume (as before) that when we see the glass shatter, we gain deep knowledge that the glass shatters and that it has been thrown: scenarios in which the glass doesn't shatter or hasn't been thrown can be ignored.⁸ The question is how we can get from here to knowledge of any irreducibly modal facts about the glass.

If counterfactuals are primitive, then one such fact is $Thrown > Shatter$. And it is not hard to see how we could get from $Thrown \wedge Shatter$ to $Thrown > Shatter$.

The inference might be a simple matter of deduction. Some hold that $A \wedge C$ logically entails $A > C$. But the inference might be justified even without that controversial assumption. Suppose before we saw the glass shatter, the live possibilities divided into $Thrown > Shatter$ scenarios and $Thrown > \neg Shatter$ scenarios, with any remaining scenarios deserving little credence. Since $Thrown > \neg Shatter$ entails $Thrown \supset \neg Shatter$, and $Thrown \wedge Shatter$ is logically incompatible with $Thrown \supset \neg Shatter$, observation of $Thrown \wedge Shatter$ then allows us to rule out all $Thrown > \neg Shatter$ scenarios, leaving most of our credence on $Thrown > Shatter$ scenarios.

So the problems I raised for our knowledge of F do not arise for our knowledge of (supposedly) primitive counterfactuals likely $Thrown > Shatter$ – nor do they arise for our knowledge of fragility if that is analysed in terms of such counterfactuals. Watching the glass shatter might give us deep knowledge that it is (or was) fragile.⁹

⁸ This assumption may be too generous. Friends of powers often hold that the glass's shattering is itself just a collection of dispositions. We must then ask how we could tell that *these* dispositions are instantiated. Let's say the shattering of the glass involves, among other things, counterfactuals about how the bits of glass would affect a sheet of paper held in their way. Since no paper is actually held in the way, how do we know that these counterfactuals are true? The answer I'm about to give for our knowledge of $Thrown > Shatter$ does not carry over. This problem is especially acute for "holist" views on which the fundamental powers are individuated by the total nomic profile of all powers. For the glass to have a particular power then requires the truth of many non-trivial counterfactuals about the entire nomic structure of the world.

⁹ I say 'might' because other problems remain. The problems from footnote 8, for example. Also, how do we know that the glass is irreducibly such that it would break under stress as opposed to, say, such that it would break *under stress at room temperature* (but not at other temperatures), or such that it would break *under stress or under UV light*? And how do we know that it does not have an irreducible propensity to randomly break, irrespective of any stress? These alternatives are all

The key to this solution is that the primitive counterfactuals are not posited as *ontological* (or *typological*, see [Busse 2018]) primitives. They are primitive *ideology*. We have assumed (on behalf of the Anti-Humean) that the truth-value of (some) counterfactuals is not settled by any non-counterfactual truths. Informally, if God wanted to give a complete description of reality from which one could in principle infer all truths, she would have to explicitly include counterfactuals. In that sense, the counterfactual operator ‘>’ is a piece of primitive ideology: it has to be used in any complete description of the world.

Treating ‘>’ as primitive ideology is not the same as positing a primitive piece of ontology. We don’t have to assume that counterfactuals are made true by the presence of a special entity.¹⁰ Indeed, we should not, since that would bring back the permutation problem.

7 Generalising

When I stated the access problem, I focused on a particular unHumean whatnot: primitive fragility. I also focused on a particular error possibility: scenarios in which that property is replaced by another primitive power. It might be useful to state the problem in more abstract and general terms.

Anti-Humeans commonly assume that there are metaphysically primitive facts about the presence and distribution of unHumean whatnots. Let M be the totality of these facts. Let H be the totality of all fundamental Humean facts (if any). Finally, let E be the complete truth about the character of our perceptual experience (past, present, and future).

The first premise in my skeptical argument is that $E \wedge H$ is *a priori* consistent with alternative hypotheses about the presence and distribution of unHumean whatnots. If M specifies that some things in our environment have a primitive power F , then $E \wedge H$ is *a priori* compatible with scenarios in which instead these things have F^* . If M specifies that a certain higher-order universal N relates F and G (see [Armstrong 1983]), then $E \wedge H$ is compatible with scenarios in which N instead relates G and H .

Anti-Humeans will typically regard many of these scenarios as metaphysically impossible. If F has an “essence” that “grounds” the truth of counterfactuals like *Thrown > Shatter*, they might say, then any scenario in which a glass has F , is thrown, and yet fails to shatter is metaphysically impossible. Similarly, necessitarians about laws will

compatible with the observed *Thrown* \wedge *Shatter*.

¹⁰By analogy, consider an ontology with individuals and universals. In that picture, a complete description of the world can’t just list the elements in the ontology; one would also need to specify which particulars instantiate which universals. And instantiation is not a further entity, on pain of Bradley’s regress. The copula is a piece of primitive ideology. The counterfactual operator might have the same status.

say that scenarios in which N relates G and H even though some G fails to be H are metaphysically impossible. I do not assume otherwise. I only assume that the relevant scenarios can't be ruled out *a priori*.

My second premise is that scenarios that differ merely by shuffling around unHumean whatnots sometimes deserve equal *a priori* credence. For example, two scenarios that only differ by swapping fundamental powers deserve the same (or roughly the same) *a priori* credence. More specifically, there are many scenarios that agree with respect to E and H , differ substantially in the distribution of unHumean whatnots, but have roughly equal *a priori* credibility.

From these two premises¹¹, it follows that we lack (and will never acquire) deep knowledge about the distribution of unHumean whatnots.

A perhaps familiar special case of this argument involves scenarios in which the unHumean whatnots have all been removed, while holding fixed E and H . The argument claims that we have no reason to favour worlds with unHumean whatnots over corresponding "Hume worlds" with all unHumean elements removed (see e.g. [Earman and Roberts 2005]). I have instead focused on worlds where the unHumean elements are swapped around, because I suspect that friends of primitive powers will maintain that Hume worlds really should be given little *a priori* credence. I disagree, but I don't know how to argue the point. I hope that few will be tempted to maintain that scenarios in which certain things around us have some primitive property F are *a priori* more credible than scenarios in which they have a different primitive property F^* .

Instead of rejecting one of the premises, anti-Humeans might accept the conclusion and admit that we can never know (or even have evidence about) which unHumean whatnots are present in our environment – except in the cheap and shallow sense in which we can know a patient's ailment by giving it a name. I don't think this kind of "modal humility" would be a fatal flaw, but (as I mentioned earlier) it does not fit how anti-Humean views are usually presented.

Another response is to not posit any unHumean whatnots and merely insist that certain modal truths are conceptually basic.¹² This blocks the argument because there can be *a priori* connections between modal and non-modal hypotheses, even if the former are not reducible to the latter. A primitive counterfactual might still entail the corresponding material conditional. Similarly, 'it is a law that all G s are H ' might entail 'all G s are H ', even if the concept of a law is irreducible.

Admittedly, the move to ideology also has its costs. For example, it appears to violate

¹¹ In fact, from the second alone, but it is useful to state the weaker premise 1 as a separate assumption.

¹² [Vetter 2020] also argues that anti-Humeanism is best understood not as an ontological thesis, but as the view that some modal truths cannot be explained in ultimately non-modal terms. However, Vetter's "explanation" relation is metaphysical. To meet the epistemic challenge, we would need an epistemic relation.

an attractive “truthmaker principle” according to which all truths are made true by what fundamental things there are and what fundamental properties and relations they instantiate. Friends of powers, in particular, often endorse the intuition that facts about what a glass would do when struck are grounded in more fundamental facts about the glass’s intrinsic properties. With the move to ideology, they would have to give up this intuition.

One might also argue that inexplicable *a priori* connections are mysterious. On typical Humean accounts, we can explain why, say, a counterfactual entails the corresponding material conditional. No such explanation can be given if counterfactuals are ideologically primitive.¹³

Special problems also remain for our knowledge of chance.

8 Propensities

So far, we have looked at “deterministic” powers. Almost all powers with which we are familiar in science and everyday life are non-deterministic. A fragile glass is not guaranteed to break when thrown against the wall – not even if we specify its exact microstate, along with the precise manner of throwing and the state of the rest of the universe. The evolution of the universe’s wavefunction will give rise to branches on which the glass survives unscathed. Without fixing the microstate of universe, even statistical mechanics tells us that the glass might survive.

I will use ‘propensity’ to refer to non-deterministic dispositions. The canonical expression of a propensity involves objective probability or chance: an object i has probability x to O under C . Sometimes we don’t have a precise numerical representation. We may understand qualitative propensities as coarse-grained versions of numerical propensities. We can also allow for cases where the circumstances C are trivial. Perhaps a radium atom has a certain propensity to decay within a certain time span, under any circumstances whatsoever.

In section 6, I argued that ideological anti-Humeans might explain our knowledge of counterfactuals (and dispositions) by exploiting the entailment between counterfactuals and the corresponding material conditionals. With probabilistic counterfactuals (and dispositions), this move no longer works, because probabilistic counterfactuals do not entail corresponding material conditionals.

Informally, a counterfactual $A > C$ is not “completely modal”. The counterfactual has implications not just for what might or would or could happen, but also for what does happen: $A > C$ entails that either not- A or C . But expressions of non-trivial propensity have no such implications. The fact that a coin has a 50% chance of landing heads

¹³ Thanks here to Siegfried Jaag.

when tossed logically entails nothing about actual outcomes. So how can we discover the propensity by observing outcomes?

As before, this is not quite the right way to ask the question. We know how to discover propensities. The question is whether an anti-Humean interpretation of propensities can make sense of these methods.

Let's look at an example. We want to find out whether a coin is fair or biased towards tails. A good approach is to make a large number of tosses. If the coin lands heads about half the time, it is safe to conclude that it is fair. If we get more tails than heads, we may conclude that it is biased.

We could model this whole process as an inference to the best explanation, following the frequentist approach to statistical inference. But we can get a clearer picture by adopting a Bayesian perspective.

The Bayesian treatment invokes another kind of probability, which earlier I've called 'credibility'. Credibility is an epistemic notion – something like rational degree of belief. (How exactly credibility should be understood doesn't matter for our purposes.) Intuitively, if you keep getting more tails than heads in repeated tosses of a coin, the more *probable* it becomes that the coin is biased. Here 'probable' means credible.

Let's suppose, for simplicity, that there are only two live possibilities for our coin: either the coin is fair (H_F) or it is biased 2:1 towards tails (H_B). Initially, the two possibilities have equal credibility. Then we start tossing. The first four outcomes are heads, tails, heads, heads. Intuitively, this supports the hypothesis that the coin is fair, since the observed outcomes are more probable given H_F than given H_B . More precisely, conditional on H_F , the probability of getting the observed outcomes is $1/2^4 = 1/16$. Conditional on H_B it is $2/81$. By Bayes's Theorem, the credibility of H_F therefore rises from 0.5 to about 0.72.

The crucial step here is the move from an assumption about the coin's propensity to the credibility of certain outcomes. On the assumption that the objective probability of heads on each toss is $1/3$, for example, the credibility of the observed outcomes is $2/81$.

Whether this is plausible depends on what we mean by 'objective probability'. Suppose we identify a coin's objective probability of landing heads with the relative frequency with which an angel has touched that side of the coin. On this interpretation, assumptions about objective probability arguably tell us nothing about how the coin will land. If I told you that an angel has touched the heads side of a coin once and the tails side twice, would you feel rationally compelled to assign credence $2/81$ to the hypothesis that the coin will land heads, tails, heads, heads in the next four tosses? Hardly. The angelic touch theory does not fit our epistemic practice.

Lewis famously argued that the same problem affects anti-Humean interpretations of chance:

I haven't the faintest notion how it might be rational to conform my credences about outcomes to my credences about some mysterious unHumean magnitude. Don't try to take the mystery away by saying that this unHumean magnitude is none other than *chance!* I say that I haven't the faintest notion how an unHumean magnitude can possibly do what it must do to deserve that name [Lewis 1986: xvf.].

Lewis here assumes that the anti-Humean identifies chance with an ontologically primitive "unHumean magnitude". As in the case of deterministic dispositions, it helps to instead treat chance (or propensity) as primitive ideology. On this view, truths about propensities are not entailed by suitably different truths; a complete description of the world would have to explicitly specify the propensities. (We have a somewhat odd primitive of the form '...has probability ...of ... under ...'.) Just as the connection to truth is built into the classical concept of lawhood, so the connection to credibility might be built into the concept of a propensity.

This time, however, the price one has to pay is not just a violation of the truthmaker principle. We don't just get inexplicable modal facts. We also get inexplicable epistemic norms. There is no hope of finding an informative explanation of *why* rational credence should be guided by assumptions about propensities. The bridge principle linking propensities to (rational credence about) outcomes – what Lewis [1980] called the "Principal Principle" – must be accepted as a basic norm of rationality.

I have no objection to basic norms of rationality. I have already committed to basic norms that license inferences from the observed to the unobserved. But the fewer basic norms, the better. We all face the problem of induction: we can't refute the inductive skeptic. Primitivists about propensity face an additional skeptical problem: they can't refute a "propensity skeptic" who denies that we have any reason to be confident that something is going to happen if we know that it has a high chance.

In effect, by assuming a primitive norm linking propensity and credence, the anti-Humean concedes that she has no explanation of why the standard methods of discovering propensities work. In that sense, the access problem has not been solved – although the move to ideology at least makes the position intelligible.

Once again Humeans do not face this problem. On popular Humean interpretations of objective probability, the Principal Principle can be derived from independently plausible epistemic norms (see [Schwarz 2014]). Humeans can refute the propensity skeptic (unless the propensity skeptic is also an inductive skeptic).

9 Resiliency

A noteworthy feature of propensities is that if a system has a non-trivial propensity to O under C , then it is hard to predict whether O will come about in any given instance

of C . It is hard to predict when a radium atom will decay, or how a fair coin will land when tossed.

More precisely, on the assumption that a coin is fair, the credibility of getting heads is $1/2$ even if we take into account all kinds of other information that we could easily obtain – say, about the current time, the weather, the result of previous tosses, and so on. Information about propensities tends to “screen off” other available information. This feature of propensity hypotheses is closely related to what Skyrms [1980] called *resiliency*, and it plays an important role in our epistemic practice.¹⁴

We might think of resiliency as a further link between chance and credence, stating that chance hypotheses screen off a wide range of other propositions. But Lewis arguably intended this to be part of his Principal Principle.

In our terminology, Lewis’s Principle says that the *a priori* credibility of an outcome, conditional on the hypothesis that its objective probability is x , is also x , and it remains x when conditionalising on further “admissible information”. Some have suggested that “admissible information” should be understood as any information that is screened off (from the outcome) by information about objective probability. But this would make the admissibility clause in the Principle redundant. I prefer a more substantive reading of the clause, on which it captures the resiliency aspect of chance hypotheses. On this reading, ‘admissible information’ is a placeholder for the domain of resiliency. Lewis tentatively suggested that information about the past is usually admissible. But the details depend on the relevant chance hypothesis and its theoretical context. In quantum mechanics, information about a system’s present microstate is admissible, in statistical mechanics it is not.

If propensities are primitive, I see no way of deriving the Principal Principle from more basic norms, including the part of it that captures resiliency. But whatever we think about the other part of the Principle, it is hard to believe that the resiliency part is primitive. Surely it is not a brute fact of epistemic normativity that probabilistic hypotheses in quantum-mechanics screen off information about microstates, but those of statistical-mechanics do not!

The best response I can think of for anti-Humeans is to hold that truly primitive propensities only exist in an indeterministic universe, and that (by conceptual necessity)

¹⁴We may have already seen resiliency in action. In the previous section, I claimed that on the assumption that the chance of heads on each toss is $1/2$, the (epistemic) probability of heads, tails, heads, heads is $1/2^4 = 1/16$. This assumes that the individual tosses are (probabilistically) independent. Where did that assumption come from? One might suggest that a fair coin not only has a 50% propensity to land heads when tossed, but also a 25% propensity to land heads-and-then-tails when tossed twice, and so on. I prefer to keep the propensities simple: the coin only has a 50% propensity to land heads when tossed. The independence assumption is then an instance of resiliency: given that the chance of heads on each toss is $1/2$, the credibility of heads on the second toss is $1/2$, even after taking into account the previous outcome.

their domain of resiliency is fixed to include every proposition about the past, or every proposition that carries no information about what happens causally downstream of the relevant circumstances. The propensities associated with coin tosses or the diffusion or gas are somehow derived, in a way that explains their varying domain of resiliency.

I would feel uneasy about this response. It appears unduly opinionated about the domain of resiliency for fundamental propensities, which to me looks like an empirical question. (Indeed, shouldn't we replace 'in the past' by 'in the past lightcone'?) Also, what is the proposed derivation of higher-level propensities? Should we understand them along Humean lines, perhaps as best-system probabilities? They can hardly be derived from fundamental propensities: it is an open question whether our world is deterministic, but it is not an open question whether there are fair coins. Perhaps the best option for anti-Humeans is to base higher-level propensities on properties of the lower-level dynamics, using the "method of arbitrary functions" (see e.g. [Strevens 2003]). But the details would need to be spelled out carefully, and I have not seen that done.

10 Conclusion

If we assume that reality contains facts of a certain kind, we may ask how we could know about these facts. When it comes to facts about natural modality, I have assumed that this is an empirical matter: somehow, our perceptual experiences are supposed to shed light on the world's modal character. But our experiences don't simply reveal the modal facts, settling any doubts or questions one may have had about the presence and distribution of unHumean whatnots.

By this measure, our experiences directly reveal almost nothing. They don't settle all doubts about whether we have hands, or whether the sun will rise tomorrow. In each case, we need bridge principles linking our experiences with hypotheses about the external world. Some of these principles can be learned, but some must be *a priori*. Most of our empirical knowledge rests on the *a priori* assumption that our surroundings are more or less as they perceptually appear to be, and that certain regularities in the observed part of the world continue to hold in the unobserved part.

Anti-Humeans sometimes claim that we are only justified to make this last assumption, about the "uniformity of nature", if we believe in unHumean whatnots, whose presence guarantees the world's regularity. I disagree. I don't understand how unHumean whatnots could guarantee or explain the world's regularity in the first place. Moreover, worlds with unHumean whatnots are not automatically regular. We would need a further *a priori* assumption that the unHumean whatnots themselves are simple, unchanging, and come in regular patterns. There are conceivable scenarios in which things constantly change their powers, or in which things are constantly replaced by other things with different powers (see [Tugby 2016]). There are also scenarios in which the primitive

powers are utterly gerrymandered – in which electrons have a disposition to repel one another on Tuesdays afternoons in the vicinity of apple trees, but to attract one another on Wednesday evenings at the beach. Scenarios like these will have to be deemed *a priori* improbable. So we all need *a priori* assumptions about the simplicity and uniformity of nature.

Standard forms of anti-Humeanism seem to require further assumptions, and much stranger ones. Take any complete hypothesis about the presence and distribution of unHumean whatnots. The hypothesis might specify which things have a certain fundamental power F , which first-order universals are related by the necessitation relation N , or something along these lines. If this hypothesis – this scenario – is compatible with our total history of perceptual experience, then so are many other scenarios in which the unHumean whatnots have been rearranged: F has been swapped with F^* , N relates different first-order universals, and so on. Some of these permuted scenarios may be too gerrymandered to deserve significant credence. But many are just as simple and uniform as the original scenario. If we are meant to have deep knowledge of natural modality, we would need *a priori* principles favouring some of these scenarios over equally simple permutations. We might need a principle according to which, in the absence of any relevant evidence, you should be confident that your surroundings contain F rather than F^* .

I fear some Anti-Humeans would be willing to go down that route. But they don't have to. I have outlined two alternatives.

One is to concede that we can't have deep knowledge about the presence and distribution of unHumean whatnots. It would, I think, be a serious cost to say that we can't have deep knowledge of whether a glass is fragile or whether radium-226 can decay into radon. But perhaps these propositions could be somehow analysed as “structural”, so that their truth does not depend on the identity of any unHumean whatnot (just as it does not depend on the identity of fundamental properties in Lewis's metaphysics).

The other alternative is to do away with unHumean whatnots and construe the anti-Humean position as one about ideology rather than ontology (or typology). Many of the troublesome permutations then become *a priori* impossible.

Either way, the access problem is still not fully resolved. In particular, something like the Principal Principle, with its curious restriction to “admissible” information, will still be needed as a basic norm.

References

David M. Armstrong [1983]: *What Is a Law of Nature?*. Cambridge: Cambridge University Press

- Helen Beebe [2003]: “Seeing Causing”. *Proceedings of the Aristotelian Society*, 103(1): 257–280
- Alexander Bird [2007]: *Nature’s Metaphysics*. Oxford University Press
- Ralf Busse [2018]: “The Adequacy of Resemblance Nominalism about Perfect Naturalness”. *Philosophy and Phenomenological Research*, 96(2): 443–469
- Shamik Dasgupta [2015]: “Inexpressible Ignorance”. *The Philosophical Review*, 124(4): 441–480
- John Earman and John T. Roberts [2005]: “Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part II: The Epistemological Argument for Humean Supervenience”. *Philosophy and Phenomenological Research*, 71: 253–286
- John Heil [2010]: “Powerful Qualities”. In Anna Marmodoro (Ed.) *The Metaphysics of Powers: Their Grounding and Their Manifestations*, Routledge
- Jenann Ismael [2020]: “(Ms) Humean Disillusionism”
- Alexander Kelly [2013]: “Ramseyan Humility, Scepticism and Grasp”. *Philosophical Studies*, 164(3): 705–726
- Rae Langton [2004]: “Elusive Knowledge of Things in Themselves”. In *Whatever*, 130–136
- David Lewis [1980]: “A Subjectivist’s Guide to Objective Chance”. In Richard Jeffrey (Ed.), *Studies in Inductive Logic and Probability* Vol. 2, University of California Press, Berkeley.
- [1986]: *Philosophical Papers II*. New York, Oxford: Oxford University Press
- [1994]: “Humean Supervenience Debugged”. *Mind*, 103: 473–490
- [1996]: “Elusive Knowledge”. *Australasian Journal of Philosophy*, 74: 549–567. Reprinted in [Lewis 1999]
- [1999]: *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press
- [2009]: “Ramseyan Humility”. In D. Braddon-Mitchell and R. Nola (Eds.) *Conceptual Analysis and Philosophical Naturalism*, Cambridge (Mass.): MIT Press, 203–222
- Stephen Mumford and Rani Lill Anjum [2011]: *Getting Causes From Powers*. Oxford University Press

- Jonathan Schaffer [2005]: “Quiddistic Knowledge”. *Philosophical Studies*, 123(1-2): 1–32
- Wolfgang Schwarz [2014]: “Proving the Principal Principle”. In A. Wilson (Ed.) *Chance and Temporal Asymmetry*, Oxford: Oxford University Press, 81–99
- [2018]: “Imaginary Foundations”. *Ergo*, 29: 764–789
- Brian Skyrms [1980]: *Causal Necessity. A Pragmatic Investigation of the Necessity of Laws*. New Haven: Yale University Press
- Michael Strevens [2003]: *Bigger than Chaos*. Cambridge (Mass.): Harvard University Press
- Matthew Tugby [2016]: “The Problem of Retention”. *Synthese*
- Barbara Vetter [2020]: “(Ms) Explanatory Dispositionalism”
- Neil E. Williams [2019]: *The Powers Metaphysic*. Oxford University Press