DOI: 10.1017/psa.2022.85

This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

RESEARCH ARTICLE

Conditional Probabilities and Symmetric Grounding

Andrew Brenner

Department of Religion and Philosophy, Hong Kong Baptist University, Hong Kong
Department of Philosophy, Linguistics, and Theory of Science, University of Gothenburg, Gothenburg, Sweden
Corresponding author. Email: andrew.t.brenner@gmail.com

Abstract
I present new counterexamples to the asymmetry of grounding: we have *prima facie* reason to think that some conditional probabilities partially ground their inverse conditional probabilities, and *vice versa*. These new counterexamples may require that we reject the asymmetry of grounding, or alternatively may require that we reject one or more of the assumptions (regarding, e.g., the correct interpretation of probability) which enable the counterexamples. Either way, by reflecting on these purported counterexamples to grounding asymmetry we learn something important, either about the formal properties of grounding, or about the nature of probability (or both).

1. Introduction

Some facts obtain in virtue of other facts. For example: this paper is a *philosophy* paper in virtue of its philosophical content; the disjunction "grass is blue or the sky is blue" is true in virtue of the fact that the sky is blue; it is wrong to kill me in virtue of the fact that killing me would deprive me of my future. In these cases we might say that some fact(s) is *grounded* in some other fact(s). In recent metaphysics grounding has received a great deal of attention. One dispute regarding grounding concerns grounding's formal properties. While grounding is generally characterized as being transitive, irreflexive, and asymmetric, each of these alleged properties of grounding has come under attack. In this paper I present arguments to the effect that we have *prima facie* reason for thinking that the grounding relations entered into by some conditional probabilities violate the asymmetry of grounding. The basic idea is that there seem to be at least some cases in which the conditional probability of some proposition (i.e., $P(A \mid B)$ – the probability of A, given B) is partially grounded in its inverse conditional probability (i.e., $P(B \mid A)$), while the inverse conditional probability is in turn partially grounded in the original conditional probability. When I say that one conditional probability is partially grounded in some other conditional probability, I am writing of *rational* probability assignments, or *objective* probabilities, rather than, say, our subjective degrees of belief. Perhaps there are no such things as objective probabilities. This is not an assumption I can defend here.

Note that I am *not* suggesting that, for some conditional probabilities $P(A \mid B)$ and $P(B \mid A)$, the truth or obtaining of A is grounded in or explained by the truth or obtaining of B, or vice versa. I am simply talking about conditional probabilities. So, for example, take A to be "Toby gets lung cancer" and B to be "Toby smokes." If Toby does develop lung cancer, then his getting lung cancer might very well be (partially) explained by his smoking. But it may not be the case that his smoking is even partially explained by his getting lung cancer. It is a separate matter whether the *probability* that Toby develops lung cancer, conditional on Toby's smoking (i.e., $P(A \mid B)$) is partially grounded in the probability that Toby smokes, conditional on Toby's developing lung cancer (i.e., $P(B \mid A)$) (or vice versa).

Before I proceed I would like to motivate the discussion of this issue. Other philosophers have also presented what they take to be counterexamples to the asymmetry of grounding.[1] Why care about yet another purported counterexample to grounding asymmetry? First, my new counterexamples to grounding asymmetry may shed some light on the nature of probability. This is true even if we reject the proposed counterexamples to grounding asymmetry, as it may lead us to reject some component of my presentation of these new counterexamples to grounding asymmetry (e.g., the assumption that there are objective probabilities). These new counterexamples to grounding asymmetry may also convince you that grounding is not asymmetric, even if you are unmoved by other proposed counterexamples. This is a matter of some significance, for several reasons. Grounding is generally taken to be a primitive notion.[2] One way in which philosophers generally aim to convey the concept of grounding is by way of an enumeration of its formal properties.[3] Disagreement about the formal properties of grounding might lend support to grounding skeptics, especially those who think that talk of grounding is unintelligible or otherwise confused.[4] Disagreement about the formal properties of grounding may also lend support to those who claim that there are multiple grounding relations, rather than a unified big-G "Grounding" relation,[5] if we ultimately conclude from purported counterexamples to the asymmetry of grounding that there are some grounding relations which are asymmetric, while there are others which are not. Whether or not grounding is asymmetric will also have implications for whether or not grounding is well-founded (i.e., such that all facts are either fundamental or fully grounded in some fundamental fact(s)). One way for grounding to fail to be well-founded is if there are symmetric grounding relations. Whether grounding can fail to be asymmetric will also have implications for issues outside of metaphysics which we might care about. For example, it has been argued that quantum entanglement should be interpreted in terms involving symmetric grounding relations.[6] Some religious doctrines involve symmetric grounding. This includes some interpretations of the Buddhist doctrine of interdependence, as represented by, e.g., the metaphor of the Net of Indra.[7] Similarly, the Christian doctrine of the Trinity can be interpreted in

[1] See, e.g., Jenkins (2011), Bliss (2014, 2018), Correia (2014), Rodriguez-Pereyra (2015), Thompson (2016, 2018), Nolan (2018), Woods (2018).

[2] Schaffer (2009, 364), Rosen (2010, 113), Audi (2012, 686).

[3] Rosen (2010, §5), Daly (2012, 82).

[4] Hofweber (2009), Daly (2012).

[5] Wilson (2014), Koslicki (2015).

[6] Calosi and Morganti (2021).

[7] Bliss and Priest (2018).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

terms according to which the persons of the Trinity stand in symmetric grounding relations to one another.[8] Our having other counterexamples to the asymmetry of grounding might lend some indirect support to those interpretations of entanglement, interdependence, and the Trinity which involve symmetric grounding relations. (At the very least, our having these other counterexamples to the asymmetry of grounding would indicate that these interpretations should not be rejected simply because they would involve the violation of grounding asymmetry.) Whether grounding is asymmetric is also connected with what we should think about the related notion of metaphysical explanation. If grounding is not asymmetric, this might show that metaphysical explanation also need not be asymmetric. Alternatively, if we are convinced that metaphysical explanation must be asymmetric, then showing that grounding is not asymmetric will show that grounding and metaphysical explanation are not as closely linked as some grounding theorists maintain. For example, some philosophers ("separatists") maintain that grounding merely *backs* metaphysical explanation, while other philosophers ("unionists") maintain that grounding *is* (a type of) metaphysical explanation. But if grounding and metaphysical explanation have different formal properties (e.g., one of them is asymmetric while the other one is not), then we would presumably have grounds for thinking that, at most, grounding backs metaphysical explanation.[9]

Here's the plan for the remainder of this paper. In §2 I argue that there plausibly are grounding relations linking some conditional probabilities with their inverse conditional probabilities. In §3 I argue that some such grounding relations are symmetric – i.e., such that, for some conditional probabilities and their inverse conditional probabilities, the one conditional probability partially grounds the other, and vice versa.

2. Grounding Relations Between Conditional Probabilities and Their Inverse Conditional Probabilities

In order to argue that there are *symmetric* grounding relations between some conditional probabilities and their inverse conditional probabilities I must first argue that there are grounding relations, symmetric or otherwise, linking some conditional probabilities and their inverse conditional probabilities. So, why should we think that these sorts of grounding relations obtain?

I give two arguments. The first argument is from the intuitive plausibility of there being these grounding relations, an intuitive plausibility we can recognize in the case of various particular conditional probabilities. The second argument is that the grounding relations linking some conditional probabilities with their inverse conditional probabilities are needed to account for the systematic covariation between the values of those conditional probabilities.

2.1. Argument 1: Intuitive Plausibility

We often intuitively assume that there are grounding relations linking some conditional probabilities and their inverse conditional probabilities when we consider the

[8] Cotnoir (2017).

[9] For related concerns regarding the relationship between grounding and metaphysical explanation see Thompson (2016), Maurin (2019).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

relation between some hypothesis and some piece of evidence. So, consider a case where $P(B \mid A)$ is thought of as the degree to which some hypothesis A can account for or predict some evidence B. For example, where our hypothesis A is that "Toby develops lung cancer" and our evidence B is "Toby smokes," to say that $P(B \mid A)$ is high is to say that our hypothesis that Toby develops lung cancer does a good job predicting our evidence, that Toby smokes – i.e., Toby's smoking is precisely what we would expect, given the supposition that Toby *does* develop lung cancer. Conversely, if the hypothesis fails to account for some evidence – i.e., if the evidence is particularly *unlikely* given the supposition that the hypothesis is correct – then $P(B \mid A)$ will be low.

It is often natural to think that some hypothesis is probable or improbable (partially) *because* or *in virtue of the fact that* it does a good, or a poor, job accounting for our evidence. In other words, it is often natural to think that $P(A \mid B)$ is high, or that it is low, precisely *because* $P(B \mid A)$ is high, or precisely *because* or *in virtue of the fact that* $P(B \mid A)$ is low. For example, in response to the question "why is hypothesis A, conditional on B, so improbable?" we might sensibly respond "*because* or *in virtue of the fact that* A does a poor job accounting for B, the evidence conditionalized on" (i.e., $P(B \mid A)$ is low). More concrete examples are easy to come by, and are ubiquitous anywhere we form judgments regarding the values of conditional probabilities, including both science and everyday life. Why is it so probable that Toby ate the chocolate, given that Toby's face is covered in chocolate? (Partially) *because* it is highly probable that Toby's face would be covered in chocolate if he ate the chocolate. Why is it so improbable that Toby ate the cyanide, given that Toby is alive? (Partially) *because* it is very improbable that Toby would be alive if he ate the cyanide. Why is it highly improbable that Toby is the murderer given that Toby has an alibi? (Partially) *because* it is highly improbable that Toby would have an alibi if he was the murderer. Why is it highly probable that Toby was the murderer, given that Toby's fingerprints are on the murder weapon? (Partially) *because* it is highly probable that Toby's fingerprints would be on the murder weapon if Toby was the murderer.

One way to recognize that there are these grounding relations between the conditional probabilities in question is to see the explanatory connections between the conditional probabilities, explanatory connections which must be underwritten by some non-causal dependence relation.[10] As I noted in §1, grounding is closely connected to metaphysical explanation. And in the case of the conditional probabilities I have been discussing, there clearly seems to be the relevant sorts of metaphysical explanations linking the values of the conditional probabilities with their inverse conditional probabilities. These metaphysical explanations in turn help us understand why the conditional probabilities take the values they do. For example, learning that the conditional probability that Toby is alive given that he has eaten cyanide is very low helps me understand why its inverse conditional probability – the probability that Toby has eaten the cyanide given that he is alive – is also very low. Similarly, learning that the conditional probability that Toby has an alibi given that Toby is the murderer is low helps me understand why its inverse conditional probability – the probability that Toby is the murderer given that he has an alibi – is also low. These explanatory connections are markers of grounding.

[10]Cf. Audi (2012), who notes that we can often recognize cases of grounding by seeing that grounding is needed to underwrite non-causal explanations.

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

You might be tempted to think that in these cases all that's really happening is that we learn that $P(A \mid B)$ is high (or low) by learning that $P(B \mid A)$ is high (or low). If that's the right way to think of things, then it would undermine my claim that $P(A \mid B)$ is (partially) grounded in $P(B \mid A)$. In response I would note that in the examples cited above it seems as if $P(A \mid B)$ has a high or low value because $P(B \mid A)$ has a high or low value, and this is the case whether or not we have any beliefs about the values of the probabilities in question. For example, it is highly improbable that Toby ate the cyanide, given that he is alive, *because* it is highly improbable that he would be alive if he had eaten the cyanide. And this seems to be the case even if we form no beliefs regarding the probabilities in question, and even if we do not form a belief regarding the one probability on the basis of a belief regarding the other probability. So, again, what we seem to have here is a case where a conditional probability takes a certain value in virtue of the fact that its inverse conditional probability takes a certain value, and it isn't just that we simply *learn* the value of the conditional probability by learning the value of its inverse conditional probability.

Of course, $P(A \mid B)$ may be grounded in the values of probabilities other than $P(B \mid A)$. For example, $P(A \mid B)$ may also be partially grounded in the prior probability of A (i.e., $P(A)$). If, for example, there is a low prior probability that Toby would eat the chocolate (since, say, it is very probable that he hates chocolate, given prior evidence conditionalized on), then this may render the conditional probability that he ate the chocolate, given that his face is covered in chocolate, low, even if the probability that his face is covered in chocolate, given the hypothesis that he ate the chocolate, is high – in other words, $P(A)$'s being low might make $P(A \mid B)$ low, even if $P(B \mid A)$ is high. So I don't want to suggest that $P(A \mid B)$ is high (or low) *only* in virtue of the fact that $P(B \mid A)$ is high (or low). My point is just that it seems plausible that there is a certain sort of explanatory relationship between some conditional probabilities and their inverse conditional probabilities. We recognize this explanatory relationship when we talk of the probability, or degree of evidential support (or whatever) that some piece of evidence "confers" on some hypothesis. What's more, the explanations in question don't seem to be *causal* explanations. For example, where A denotes "Toby ate the cyanide" and B denotes "Toby is alive," it does not seem as if $P(B \mid A)$'s being low *causes* $P(A \mid B)$ to be low. Rather, the explanatory relation between the two conditional probabilities seems to be a non-causal "in virtue of" explanatory relation, grounding.[11]

Why grounding, and not some other non-causal explanatory relation? I would say that the non-causal explanatory relation here seems to be sufficiently analogous to other paradigmatic cases of grounding, including the examples of grounding I mentioned at the beginning of this paper. What's more, there don't obviously seem to be any other non-causal explanatory relations to do the job. As we saw in §1, some philosophers maintain that there isn't some single big-G "Grounding" relation which the various

[11] Strictly speaking, in conformity with the common assumption that grounding relates facts, perhaps we should say that the fact that $P(A \mid B)$ has such-and-such a value is partially grounded in the fact that $P(B \mid A)$ has so-and-so a value. I will continue to speak loosely of one conditional probability grounding another conditional probability, or of one conditional probability's being high (or low) grounding another conditional probability's being high (or low). In fact, I aim to remain neutral in this paper regarding whether grounding invariably relates facts (as in Rosen (2010)), relata of any sort (as in Schaffer (2009)), or whether grounding is most perspicuously expressed as a sentential operator rather than a relation (as in Fine (2001)).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

purported cases of grounding have in common, even if there *are* a number of different relations (e.g., parthood, set membership) which fall under the "grounding" label.[12] Even if this is correct, it seems plausible to me that the grounding relations which obtain between some conditional probabilities and their inverse conditional probabilities would be *varieties* of grounding relations, even if they are not instances of *the* single big-G "Grounding" relation. In any case, if you think that the non-causal explanatory relations which obtain between some conditional probabilities and their inverse conditional probabilities should not be called "grounding" relations, you should at any rate think that we have *prima facie* reason to think that the non-causal explanatory relations in question are not asymmetric, for reasons I discuss later in this paper. And that by itself would be as noteworthy (and controversial) a conclusion as the conclusion that *grounding* can fail to be asymmetric.

In this section I have argued, more or less, that we can directly intuit or apprehend the presence of grounding relations between some conditional probabilities and their inverse conditional probabilities. Is it objectionably mysterious that we would be able to directly intuit the presence of grounding relations between conditional probabilities and their inverse conditional probabilities? Perhaps it is. But I have several points I would like to make. First, our ability to recognize cases of grounding is no more objectionably mysterious than our ability to identify causal relations, above and beyond the mere regularities which result from those causal relations. (Hume, of course, thought that it would be objectionable to suppose that we can identify causal relations in this manner. But I think that Hume was wrong about that.) Second, our ability to recognize the grounding relations which, I claim, hold between some conditional probabilities and their inverse conditional probabilities is no more objectionable than our having intuitive direct apprehension of other grounding relations. For example, when we witness an act of great cruelty we can see not only that the act is morally wrong, but we can see that it is morally wrong in virtue of the fact that it is an act of great cruelty. Similarly, if we see a beautiful piece of art, we can see not only that the art is beautiful, but that it is beautiful in virtue of its instantiating certain non-aesthetic properties (e.g., certain colors or shapes). Or, if moral and aesthetic epistemology is too contentious, note that we can see that some object is red in virtue of its being crimson, some disjunction is true in virtue of its true disjunct, and some belief is epistemically justified in virtue of the fact that it satisfies some epistemic criteria (e.g., it is formed on the basis of strong evidence). Or, consider the fact that if you are not convinced by any of the arguments so far presented in this paper, you will on reflection note that the arguments are bad arguments *in virtue of* their instantiating certain bad-making features of arguments, or their failing to instantiate certain good-making features of arguments. To suppose that, as a matter of principle, we are unable to directly intuit or apprehend instances of grounding in the case of conditional probabilities may commit one to the much less plausible thesis that we are also unable to directly intuit or apprehend these other cases of grounding.

2.2. Argument 2: Systematic Covariation

There is a second major consideration I would like to cite in favor of there being these grounding relations, one which may appeal to those who are suspicious of the idea

[12] Wilson (2014), Koslicki (2015).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

that we can directly intuit or apprehend the presence of grounding relations: the systematic covariation between the values of conditional probabilities and the values of their inverse conditional probabilities. The basic idea here is modeled after one of the main ways we infer the presence of causal relations: when there is systematic covariation between two sorts of events this will often give us grounds for inferring that one sort of event causes the other. So, for example, there is a striking positive correlation between smoking rates and deaths from lung cancer. In the United States when smoking rates increased then deaths from lung cancer increased roughly 30 years later, and when smoking rates decreased then deaths from lung cancer decreased roughly 30 years late.[13] This systematic covariation between smoking and lung cancer has been one of the major indicators that smoking causes lung cancer. Of course, systematic covariation is not a foolproof indicator of direct causal connections. For example, systematic covariation between two sorts of events may be a result of some common cause. But in cases where we are able to rule out these sorts of confounding variables, we can often justifiably infer the presence of a causal connection between the two sorts of events. But even when systematic covariation allows us to infer that one sort of event causes another sort of event, further investigation may be needed to determine whether the causal effects are mediated or unmediated, deterministic or probabilistic, or whether they run in one direction rather than the other (e.g., from smoking to lung cancer or from lung cancer to smoking).

We can employ a similar methodology in cases of grounding.[14] Where there is systematic covariation between two sorts of facts, this can serve as an indication that there are grounding relations linking those facts. For example, for some act of great cruelty we can see both that the act is morally wrong *and* that if the act had not been cruel then, other things remaining the same, the act would not have been morally wrong. If some disjunction is true in virtue of its true disjunct, we can see that, other things remaining the same, the disjunction would not have been true had that disjunct not been true. If my arguments are bad arguments in virtue of their instantiating bad-making features, then we can see that, other things remaining the same, the arguments would not have been bad if they had not instantiated those bad-making features. In all these cases grounding brings with it systematic covariation, and in principle we could infer that the relevant grounding relations obtain by noting the systematic covariation between the facts in question (although in practice we may just directly see that the grounding relations obtain, and not infer their presence from the systematic covariation).

Just this sort of systematic covariation obtains in the case of some conditional probabilities and their inverse conditional probabilities. Conditional probabilities are governed by Bayes' Theorem. Where "$P(A)$" denotes the probability of some proposition A, and "$P(A \mid B)$" denotes the probability of some proposition A conditional on some proposition B, the conditional probability of any given proposition is, per Bayes' Theorem, expressed in the following equation:

$$P(A \mid B) = \frac{P(B|A)\,P(A)}{P(B)}$$

Note also that $P(B \mid A)$ can be expressed by Bayes' Theorem:

[13] Pearl and Mackenzie (2018, 172).
[14] Cf. Schaffer (2016); Ismael and Schaffer (2020).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

$P(B \mid A) = \frac{P(A|B)\,P(B)}{P(A)}$

Here I suppress any background beliefs which might enter into the determination of $P(A)$ or $P(B)$, since the purported cases of symmetric grounding which interest me here specifically involve the relationship between the conditional probabilities $P(A \mid B)$ and $P(B \mid A)$.

Bayes' Theorem draws our attention to an interesting covariation or correlation between the values of $P(A \mid B)$ and its inverse conditional probability $P(B \mid A)$, insofar as $P(B \mid A)$ is cited on the right side of Bayes' Theorem when we consider the conditional probability $P(A \mid B)$, while $P(A \mid B)$ is cited on the right side of Bayes' Theorem when we consider the conditional probability $P(B \mid A)$.

In the case of many conditional probabilities, $P(A \mid B)$ and its inverse conditional probability $P(B \mid A)$ are such that had one of them had a higher (or lower) value then, other things remaining the same, the other would have had a higher (or lower) value. (This will hold true only of conditional probabilities which are not probabilistically independent of one another (i.e., such that $P(A \mid B) = P(A)$, and $P(B \mid A) = P(B)$). Where the conditional probabilities are probabilistically independent of one another, raising or lowering the value of $P(A \mid B)$ will not raise or lower the value of $P(B \mid A)$, and vice versa.) Just as systematic covariation between smoking and lung cancer can allow us to infer that there is a causal connection between the two, so too the systematic covariation between the values of conditional probabilities and their inverses can allow us to infer that there is a non-causal connection between the two, namely grounding.

In the case of causal inference we can establish that systematic covariation between two sorts of events is likely the result of causal relationships between them by ruling out alternative explanations of the covariation. We can do something similar in the case of grounding. When it comes to the systematic covariation between many conditional probabilities and their inverse conditional probabilities, the chief alternative explanations of the covariation are: 1.there is some common ground or explanation for the values of the conditional probabilities; 2.there are separate grounds or explanations for each of the conditional probabilities; or 3.the mathematical relationship between the conditional probabilities explains their coviariation. I'll now argue that none of these alternative explanations of the covariation is plausible.

2.2.1. A Common Ground?

Start with the proposal that the values of $P(A \mid B)$ and $P(B \mid A)$ systematically covary because they have a common ground or explanation. This proposal is modeled after the confounding variables which sometimes account for systematic covariation between two sorts of events. When, for example, it was noted that there is a striking positive correlation between smoking and developing lung cancer, in order to determine whether smoking was the cause of lung cancer researchers first had to rule out there being some common cause for both smoking and developing lung cancer (e.g., some genetic characteristic which made people predisposed both to smoke and to develop lung cancer).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

As for a potential common ground for $P(A \mid B)$ and $P(B \mid A)$, two possibilities have been suggested to me. First: $P(A \mid B)$ and $P(B \mid A)$ are grounded in $P(A)$, $P(B)$, and $P(A\&B)$.[15] We can derive $P(A \mid B)$ and $P(B \mid A)$ from $P(A)$, $P(B)$, and $P(A\&B)$:

$$P(A \mid B) = \frac{P(A\&B)}{P(B)}$$

$$P(B \mid A) = \frac{P(A\&B)}{P(A)}$$

The second potential common ground or explanation is one in which the conditional probabilities in question are both grounded in or explained by certain sorts of causal relations, or laws governing causal relations.[16] So, for example, consider one of the examples noted above, where we consider the probability that someone is alive given that they have eaten cyanide ($P(Alive \mid Cyanide)$), and the inverse conditional probability that someone has eaten cyanide, given that they are alive ($P(Cyanide \mid Alive)$). Both of the conditional probabilities are low, and above I claimed that one of these conditional probabilities is low because the other conditional probability is low. But perhaps they are both low simply because cyanide causes death, or because the laws of nature are such that eating cyanide tends to cause death.

It's important to note that the conditional probabilities' having a common ground or explanation is compatible with them being grounded in the way I suggest (i.e., is compatible with $P(A \mid B)$ partially grounding $P(B \mid A)$, and vice versa). Still, it might be claimed that the availability of a common ground or explanation for the covaration between $P(A \mid B)$ and $P(B \mid A)$ would undermine our motivation for postulating grounding relations between those conditional probabilities. So, it's important to see that the proposed common grounds on offer cannot account for the systematic covariation in question.

Start with the first proposal, that the covariation between $P(A \mid B)$ and $P(B \mid A)$ is entirely the result of the fact that those conditional probabilities are grounded in $P(A)$, $P(B)$, and $P(A\&B)$. Note that the proposal is that $P(A)$, $P(B)$, and $P(A\&B)$ *jointly* ground $P(A \mid B)$ and $P(B \mid A)$, as none of the former probabilities by themselves are capable of grounding the latter conditional probabilities.

The claim of systematic covariation is that, other things being equal, if you raise or lower $P(A \mid B)$ then $P(B \mid A)$ will also be raised or lowered. Now, suppose we keep these probabilities $P(A)$, $P(B)$, and $P(A\&B)$ fixed. Is it possible to do that, and see if raising or lowering $P(A \mid B)$ will raise or lower $P(B \mid A)$? If so, the systematic covariation remains. This is how we can often rule out confounding factors in causal inference. For example, we notice a positive correlation between smoking and developing lung cancer, and we are investigating whether this correlation obtains (in part) because smoking causes lung cancer, or whether it obtains because some genetic factor causes both smoking and lung cancer. We can test this, in principle, by holding fixed the relevant genetic factors and seeing if smoking still systematically covaries with lung cancer. For example, we might see if identical twins who smoke develop lung cancer at higher rates than their genetically identical twins who do not smoke. If, holding fixed genetic traits, smoking still

[15]This has been independently suggested to me by several people: Rebecca Chan, Jakob Koscholke, Moritz Schulz, and Maximilian Zachrau.

[16]This possibility was suggested to me by Roman Heil and an anonymous referee.

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

systematically covaries with one's developing lung cancer, this will be an indication that the covariation is not merely the result of a genetic common cause.

But what about cases where the confounding factors we wish to rule out are *joint* causes? For example, suppose we know that the smoking/cancer link cannot be accounted for solely by genetic factors or solely by environmental factors, but we are wondering if it could be entirely accounted for in terms of genetic and environmental factors together. We could test this by holding both genetic and environmental factors fixed, and seeing if there remains a positive correlation between smoking and developing lung cancer. But if for whatever reason it's not feasible to hold both genetic and environmental factors fixed, we could instead hold fixed *either* genetic or environmental factors, and see if there remains a positive correlation between smoking and developing lung cancer. For example, we might see if genetically identical identical twins raised in different environments are still such that their smoking is positively correlated with their developing lung cancer. If it is, then the correlation could still be explained by environmental factors, since the difference in environments between the twins might cause a difference both in smoking rates and rates of developing lung cancer. But we should conclude that the correlation between smoking and developing lung cancer is probably not explained jointly by both genetic and environmental factors, since even when holding fixed genetic factors the correlation between smoking and lung cancer remains.

We can say something similar in the case of the grounding relations which, I claim, hold between some conditional probabilities and their inverse conditional probabilities. We wonder whether the covariation between $P(A \mid B)$ and $P(B \mid A)$ is merely the result of a common ground, $P(A)$, $P(B)$, and $P(A\&B)$. We can test this by holding fixed $P(A)$, $P(B)$, and $P(A\&B)$ and seeing if $P(A \mid B)$ and $P(B \mid A)$ are still such that if one probability is raised so is the other, and if one probability is lowered so is the other. The result is a bit messy. Recall that I am only interested in conditional probabilities which are not probabilistically independent of one another (i.e., such that $P(A \mid B) = P(A)$, and $P(B \mid A) = P(B)$). So, let's just focus on conditional probabilities which satisfy this condition. Let's also assume that $P(A)$ is not equal to 0, as then $P(B \mid A)$ would not be well-defined, and let's also assume that $P(B)$ is not equal to 0, for a similar reason. Given these assumptions, if we hold fixed $P(A)$ and $P(B)$ then raising or lowering $P(A \mid B)$ will indeed raise or lower $P(B \mid A)$, but it will also raise or lower $P(A\&B)$. By contrast, if we hold fixed $P(A\&B)$ then raising $P(A \mid B)$ will raise $P(B \mid A)$ only if both $P(A)$ and $P(B)$ are lowered. Similarly, if we hold fixed $P(A\&B)$ then lowering $P(A \mid B)$ will lower $P(B \mid A)$ only if both $P(A)$ and $P(B)$ are raised.[17] But it is mathematically impossible to hold fixed $P(A)$, $P(B)$, *and* $P(A\&B)$ while raising or lowering $P(A \mid B)$ or $P(B \mid A)$. By the same token, it's mathematically impossible to hold fixed $P(A)$, $P(B)$, *and* $P(A\&B)$ while $P(A \mid B)$ and $P(B \mid A)$ remain such that raising or lowering one of them will raise or lower the other. Even so, the fact that this covariation between the conditional probabilities remains even when *some* of $P(A)$, $P(B)$, and $P(A\&B)$ are held fixed is an indication that the covariation is not merely the result of the conditional probabilities being grounded in $P(A)$, $P(B)$, and $P(A\&B)$. This is similar to the point I made earlier about smoking and lung cancer: the correlation between smoking and developing lung cancer is probably not explained jointly by both genetic and environmental factors,

[17] This is because $P(A \mid B)\, P(B) = P(B \mid A)\, P(A) = P(A\&B)$.

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

since even when holding fixed genetic factors the correlation between smoking and lung cancer remains.

Let's turn to the second proposed common ground: that $P(A \mid B)$ and $P(B \mid A)$ have a common ground or explanation in the causal relations or laws which obtain. For example, $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$ are both low because consuming cyanide causes death, or because the causal laws are such that consuming cyanide causes death.

This proposal doesn't work. The proposal is that the probabilities have certain objective values depending on what real-world conditions obtain (e.g., what contingent causal laws obtain). But if this were the right way of thinking of the conditional probabilities, then we should presumably be led to the conclusion that the conditional probabilities would all have values of either 1 or 0, as, given the real-world conditions, the propositions whose probabilities interest us are either true or false. For example, Toby's being alive conditional on his having consumed cyanide will have the same probability as Toby's being alive conditional on his not having consumed cyanide, namely 1 if he is alive and 0 if he is not alive. After all, the values of the conditional probabilities are determined by which real-world conditions obtain – e.g., which contingent causal laws obtain, or Toby's being alive. This seems like the wrong result. At the very least, the probability that Toby is alive, conditional on his having consumed cyanide, should be lower than the probability that Toby is alive, conditional on his not having consumed cyanide.

What we should say is that $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$ will both be low *conditional* on the proposition that consuming cyanide causes death, or that the laws of nature are such that consuming cyanide causes death, or whatever.[18] And conditional on the proposition that consuming cyanide causes death, $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$ may both be low even if as a matter of fact consuming cyanide does not cause death. Remember, we are dealing with objective conditional probabilities, and so the objective degree of evidential support conferred on some propositions by some other propositions. These objective degrees of evidential support obtain regardless of which contingent causal relations or laws obtain. What I am suggesting, in effect, is that we can hold fixed which contingent causal relations or laws obtain, and the systematic covariation between the conditional probabilities which interest us remain, so long as the conditional probabilities in question involve conditionalization on the obtaining of the relevant causal relations or laws (e.g., those causal laws which account for the fact that consuming cyanide causes death). So, those contingent causal relations or laws cannot entirely account for the systematic covariation between those conditional probabilities.

[18] When I introduced the example involving $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$, I said that both probabilities were low. But I did not explicitly conditionalize on the contingent causal relations or laws being such that consuming cyanide causes death. So, why did I assume that $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$ were both low? The reason is because I *implicitly* conditionalized on the obtaining of the relevant causal relations or laws. For example, when above I said that $P(Alive \mid Cyanide)$ and $P(Cyanide \mid Alive)$ are both low, what I really meant is that those conditional probabilities are low *given* either a tacit or explicit conditionalization on the proposition that consuming cyanide causes death. Much background information about how the world works is often left implicit when we discuss conditional probabilities. We have to do this, as in practice we cannot explicitly state all of the relevant background information.

2.2.2. Separate Explanations For The Conditional Probabilities?

Perhaps $P(A \mid B)$ and $P(B \mid A)$ are not grounded in one another, and they do not have a common ground or explanation, but rather they have the values they do because of *separate* grounds or explanations.

It's important to note that the conditional probabilities' having these sorts of grounds or explanations is compatible with them also being grounded in the way I suggest (i.e., is compatible with $P(A \mid B)$ partially grounding $P(B \mid A)$, and vice versa). What's more, my own proposal is preferable to this one, as this current proposal leaves the systematic covariation between $P(A \mid B)$ and $P(B \mid A)$ unexplained. There being separate grounds or explanations for the values of $P(A \mid B)$ and $P(B \mid A)$ may account for why they have the values they do (e.g., may account for the fact that both $P(A \mid B)$ and $P(B \mid A)$ are high). But it could not account for the fact that, other things remaining the same, raising or lowering one conditional probability will lower or raise its inverse conditional probability. If, for example, $P(A \mid B)$ is high because of some ground or explanation, and $P(B \mid A)$ is high for an entirely separate reason, then raising or lowering the one conditional probability should not be expected to raise or lower the other conditional probability.

2.2.3. Does The Mathematical Relationship Between The Conditional Probabilities Explain Their Covariation?

Here is another proposed explanation for the covariation between $P(A \mid B)$ and $P(B \mid A)$.[19] Perhaps the covariation between $P(A \mid B)$ and $P(B \mid A)$ can be entirely accounted for in terms of their mathematical relationship, e.g., the mathematical relationship described in Bayes' Theorem. Given this mathematical relationship, we can see that if $P(A \mid B)$ is raised or lowered then, other things remaining the same, $P(B \mid A)$ will be raised or lowered as well (assuming, again, that the conditional probabilities are not probabilistically independent of one another).

Here is my response to this proposal. I don't think that the mathematical relationship between $P(A \mid B)$ and $P(B \mid A)$ really offers an explanation for their covariation. The mathematics – Bayes' Theorem, for example – *describes* the covariation, but does not explain the covariation, or show that the covariation need not be underwritten by grounding relations between the conditional probabilities. Compare: we might have a mathematical description of the covariation between two objects' masses and their gravitational attraction, but this mathematical description does not *explain* the covariation, nor does it show that the objects' having greater mass does not explain why they have greater gravitational attraction. And we could say something similar about grounding. We may have some true generalization which describes the covariation between two sorts of facts. For example, supposing utilitarianism is correct, we may have some true generalization to the effect that actions are morally obligatory iff they maximize utility. With this true generalization in hand we can, from some action's being such that it maximizes utility, derive that the action is morally obligatory. Similarly, from some action's being morally obligatory we can derive that the action maximizes utility. But simply because this true generalization *describes* the covariation between an action's maximizing utility and its being morally obligatory, and so allows us to derive one fact

[19] Suggested to me by an anonymous referee.

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

from the other, it does not follow that the true generalization *explains* why the covariation obtains, and nor does it show that one fact does not ground the other (i.e., it does not show that an action's maximizing utility does not ground its being morally obligatory).

3. *Symmetric* Grounding Relations Between Conditional Probabilities and Their Inverse Conditional Probabilities

So far I've argued that there are grounding relations linking some conditional probabilities and their inverse conditional probabilities. I first appealed to the intuitive plausibility of this thesis, and then I argued that we should believe in the grounding relations between the conditional probabilities because they best account for the systematic covariation between the values of those conditional probabilities.

This brings us to the second step of my overall argument: why should we think there are any cases in which *symmetric* grounding obtains? Why should we think that there are propositions A and B which are such that $P(A \mid B)$ is partially grounded in $P(B \mid A)$, *and* $P(B \mid A)$ is partially grounded in $P(A \mid B)$?

It seems to me that the grounding relations involved here should be thought to sometimes be symmetric because there doesn't seem to be any principled reason to think that the order of grounding runs in one direction rather than the other. Consider a particular person, who was randomly selected from some population. We might wonder whether this randomly selected person smokes, and we might also wonder whether they have lung cancer. Consider the probabilistic relationship between the proposition <the randomly selected member of such-and-such population smokes> and the proposition <the randomly selected member of such-and-such population has lung cancer>. Why is it (relatively) highly probable that the randomly selected member of the population smokes, conditional on their having lung cancer? Well, because it is (relatively) highly probable that the randomly selected member of the population has lung cancer, conditional on their smoking. In other words, $P(Smokes \mid Cancer)$ is relatively high, I claim, (partially) because $P(Cancer \mid Smokes)$ is relatively high. That seems to me to be plausible, just as the similar examples discussed above seemed plausible. But I don't see any principled reason to think that $P(Cancer \mid Smokes)$'s being relatively high would partially ground $P(Smokes \mid Cancer)$'s being relatively high, while $P(Smokes \mid Cancer)$'s being relatively high does not partially ground $P(Cancer \mid Smokes)$'s being relatively high. What could account for the asymmetry here?

Presumably the asymmetry would be a result of some asymmetry involving the propositions expressed by "Smokes" (i.e., the proposition <the randomly selected member of such-and-such population smokes>) and "Cancer" (i.e., the proposition <the randomly selected member of such-and-such population has lung cancer>). The most natural way of developing this idea is in terms of there being an asymmetric causal explanatory relation between smoking and developing lung cancer – smoking tends to cause lung cancer (but not vice versa), and *that's* why $P(Cancer \mid Smokes)$'s being relatively high partially grounds $P(Smokes \mid Cancer)$'s being relatively high (but not vice versa).[20]

It's not obvious that an asymmetric causal relation between A and B (in this case, smoking and developing lung cancer) would result in the grounding relations between

[20]Nevin Climenhaga (2020) proposes a similar account of the structure of the grounding relations entered into by conditional probabilities (although not in response to concerns about grounding asymmetry). The

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

the conditional probabilities in question being asymmetric. I will, only for the sake of argument, assume that it does, since even with this concession in place I can modify the example so that the grounding relations between the conditional probabilities in question would still be symmetric.

While smoking does normally cause lung cancer, and lung cancer does not normally cause one to smoke, there are no doubt individual cases where one's developing lung cancer *does* cause one to smoke. In the 1950s the statistician (and tobacco industry consultant) Ronald Fisher speculated that developing lung cancer might cause some people to smoke. He speculated that chronic inflammation of the lungs, acting as a precursor to the development of lung cancer, might cause some people to smoke in order to offset the slight discomfort caused by the inflammation.[21] Fisher was trying to explain away cases in which smoking seems to have caused people to develop lung cancer. But we need not endorse dubious speculation of this sort in order to note that there probably *are* cases where someone's developing lung cancer causes them to smoke. Imagine someone who has never smoked, develops lung cancer, and thinks "well, the damage has already been done, I might as well smoke." No doubt, this is not the best response to the news that one has lung cancer. But it is a response we can imagine someone having, and I would be surprised if noone has ever engaged in this sort of thought process, and smoked as a result of their developing lung cancer.

Now let us take some population which is such that all of its members smoke and have lung cancer, but half of its members have lung cancer as a result of their smoking, while half of its members smoke as a result of their having lung cancer. This population is no doubt much smaller than the total population of people who smoke and develop lung cancer, but that is beside the point: as long as the population has some members, it will serve my purposes. Earlier we considered a particular person who was randomly selected from some population. We can now stipulate that they were randomly selected from *this* population I have described in this paragraph. Now we can ask, of this randomly selected person, what is the probability that they have lung cancer, conditional on their smoking? Similarly, we can ask of this randomly selected person, what is the probability that they smoke, conditional on their having lung cancer?

Let "Population" denote the proposition that the population from which we draw our randomly selected individual is the population described above. Note that "Population" does not specify that the members of the population all smoke and have lung cancer – rather, "Population" simply specifies which individuals make up the population, where the population in question also happens to be such that all of its members smoke and have lung cancer, but half of its members have lung cancer as a result of their smoking, while half of its members smoke as a result of their having lung cancer. I claim that, for the reasons discussed above, we plausibly have a case of symmetric grounding, insofar as the conditional probability $P(Cancer\&Population \mid Smokes\&Population)$ grounds its inverse conditional probability $P(Smokes\&Population \mid Cancer\&Population)$, and vice versa. One cannot respond in this case that the grounding relation runs in one direction rather than the other, in virtue of an asymmetry involved in the propositions involved in the conditional probabilities. In particular, one cannot claim that any such asymmetry

basic idea is that whether $P(A \mid B)$ grounds $P(B \mid A)$ or $P(B \mid A)$ grounds $P(A \mid B)$ depends on what explanatory relations obtain between the propositions A and B. If A is explanatorily prior to B, then $P(B \mid A)$ grounds $P(A \mid B)$, and if B is explanatorily prior to A, then $P(A \mid B)$ grounds $P(B \mid A)$.

[21] Fisher (1958, 162-163).

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press

results from the fact that in the population in question smoking normally causes lung cancer, while lung cancer does not normally cause smoking.

References

Audi, P. (2012). A clarification and defense of the notion of grounding . In Fabrice Correia (2012). pp. 101–121.

Audi, P. (2012). Grounding: Toward a theory of the in-virtue-of relation. *Journal of Philosophy*. *109*(12), 685–711.

Bliss, R. (2014). Viciousness and circles of ground. *Metaphilosophy*. *45*(2), 245–256.

Bliss, R. (2018). Grounding and reflexivity In *Reality and its Structure: Essays in Fundamentality*, Bliss, R. and Priest, G. (eds). Oxford University Press. Oxford. pp. 70–90.

Bliss, R. and Priest, G. (2018). Metaphysical dependence, east and west In *Buddhist Philosophy: A Comparative Approach*, Emmanuel, S. M. (eds). Wiley Blackwell. Hoboken.

Calosi, C. and Morganti, M. (2021). Interpreting quantum entanglement: Steps towards coherentist quantum mechanics. *The British Journal for the Philosophy of Science*. *72*(3), 865–891.

Climenhaga, N. (2020). The structure of epistemic probabilities. *Philosophical Studies*. *177*, 3213–3242.

Correia, F. (2014). Logical grounds. *The Review of Symbolic Logic*. *7*(1), 31–59.

Cotnoir, A. J. (2017). Mutual indwelling. *Faith and Philosophy*. *34*(2), 123–151.

Daly, C. (2012). Scepticism about grounding . In Fabrice Correia (2012). pp. 81–100.

Fabrice Correia, B. S. (2012). Grounding: an opinionated introduction . In *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge University Press. Cambridge. pp. 1–36.

Fine, K. (2001). The question of realism. *Philosophers' Imprint*. *1*, 1–30.

Fisher, R. (1958). Cigarettes, cancer, and statistics. *Centennial Review*. *2*, 151–166.

Hofweber, T. (2009). Ambitious, yet modest, metaphysics In *Metametaphysics*, Chalmers, D. J., Manley, D., and Wasserman, R. (eds). Oxford University Press. Oxford. pp. 260–89.

Ismael, J. and Schaffer, J. (2020). Quantum holism: Nonseparability as common ground. *Synthese*. *197*(10), 4131–4160.

Jenkins, C. S. (2011). Is metaphysical dependence irreflexive? *The Monist*. *94*, 267–76.

Koslicki, K. (2015). The coarse-grainedness of grounding In *Oxford Studies in Metaphysics*, Karen Bennett, D. W. Z. (eds). vol. 9. Oxford University Press. Oxford. pp. 306–344.

Maurin, A.-S. (2019). Grounding and metaphysical explanation: it's complicated. *Philosophical Studies*. *176*(6), 1573–1594.

Nolan, D. (2018). Cosmic loops In *Reality and its Structure: Essays in Fundamentality*, Bliss, R. and Priest, G. (eds). Oxford University Press. Oxford. pp. 91–106.

Pearl, J. and Mackenzie, D. (2018). *The Book of Why: The New Science of Cause and Effect*. Basic Books. New York.

Rodriguez-Pereyra, G. (2015). Grounding is not a strict order. *Journal of the American Philosophical Association*. *1*(3), 517–534.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction In *Modality: Metaphysics, Logic, and Epistemology*, Hale, B. and Hoffmann, A. (eds). Oxford University Press. Oxford. pp. 109–36.

Schaffer, J. (2009). On what grounds what In *Metametaphysics*, Chalmers, D. J., Manley, D., and Wasserman, R. (eds). Oxford University Press. Oxford. pp. 347–83.

Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies*. *173*(1), 49–100.

Thompson, N. (2016). Grounding and metaphysical explanation. *Proceedings of the Aristotelian Society*. *116*(1), 395–402.

Thompson, N. (2016). Metaphysical interdependence In *Reality Making*, Jago, M. (eds). Oxford University Press. Oxford. pp. 38–56.

Thompson, N. (2018). Metaphysical interdependence, epistemic coherentism, and holistic explanation In *Reality and its Structure: Essays in Fundamentality*, Bliss, R. and Priest, G. (eds). Oxford University Press. Oxford. pp. 107–125.

Wilson, J. (2014). No work for a theory of grounding. *Inquiry*. *57*(5-6), 535–579.

Woods, J. (2018). Emptying a paradox of ground. *Journal of Philosophical Logic*. *47*(4), 631–648.

https://doi.org/10.1017/psa.2022.85 Published online by Cambridge University Press