Grounding: it's (probably) all in the head

Author(s): Kristie Miller and James Norton

Source: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, December 2017, Vol. 174, No. 12 (December 2017), pp. 3059-3081

Published by: Springer Nature

Stable URL: https://www.jstor.org/stable/45094031

Philos Stud (2017) 174:3059–3081
DOI 10.1007/s11098-016-0846-5

Grounding: it's (probably) all in the head

Kristie Miller[1] · James Norton[2]

Published online: 21 December 2016

Abstract In this paper we provide a psychological explanation for 'grounding observations'—observations that are thought to provide evidence that there exists a relation of ground. Our explanation does not appeal to the presence of any such relation. Instead, it appeals to certain evolved cognitive mechanisms, along with the traditional modal relations of supervenience, necessitation and entailment. We then consider what, if any, metaphysical conclusions we can draw from the obtaining of such an explanation, and, in particular, if it tells us anything about whether we ought to posit a relation of ground.

Keywords Grounding · Explanation · Metaphysical explanation

1 Introduction

Since Schaffer's 'On what grounds what', a large literature has been spawned that argues that we need to posit a new relation—what we will call, following its proponents, *grounding*—that captures ontological dependencies between objects, facts, or properties. We intend to be quite liberal in our use of 'grounding'. If you think that there is a special kind of explanation, *metaphysical explanation*, which is backed by the existence of a single relation that is distinct from the traditional modal relations of supervenience, necessitation and entailment, then, for us, you count as

✉ Kristie Miller
Kristie_miller@yahoo.com

James Norton
james.norton@sydney.edu.au

1 Department of Philosophy, The Centre for Time, The University of Sydney, Camperdown, NSW 2006, Australia

2 Department of Philosophy, The University of Sydney, Camperdown, NSW 2006, Australia

thinking that there are relations of ground. Defenders of grounding, thus understood, include those who posit a primitive, asymmetric, irreflexive relation that obtains between objects, properties or facts,[1] but also include those who offer a reductive account of grounding where the reductive base includes more than the traditional modal relations (and set theory).

So, for instance, Alastair Wilson (forthcoming) offers a reductive account of grounding in terms of the truths of certain counterpossibles. In a similar vein, Tallant (2015) has recently argued that we can reduce grounding to some combination of Lowe's (2010) relations of Rigid Existential Dependence and Identity Dependence. Given our use of the term, both Wilson and Tallant count as defenders of grounding. However, defenders of a view according to which 'grounding' does not pick out a single relation, but instead, picks out a host of instances of different relations (mereological fusion, supervenience, entailment, set membership, necessitation, and so on)[2] do not count as defenders of grounding in the sense in which we use the term.

Defenders of grounding, understood in this way, point to the existence of an array of what we call *grounding observations*, and suggest that these observations are best explained by the existence of grounding relations. Grounding observations are not (despite the phrase) observations of grounding, since we assume that even if there are grounding relations they cannot be directly observed. Rather, they are observations that defenders of grounding take to be evidence that such a relation exists. These include, first, our observation that a range of objects, properties and facts are *non-diachronically correlated*[3]—for instance, we observe that certain arrangements of bicycle parts frequently accompany the existence of bicycles—and, second, our observation that there exists an array of widely shared judgements about certain cases—for instance, we observe that many people judge the following to be true[4]:

A. The flower is red *because*[5] the flower is maroon.
B. The bicycle exists *because* of the existence and arrangement of the wheels, spokes, handlebars, etc.
C. <a man exists>[6] is true *because* Pythagoras exists.
D. {Pythagoras} exists *because* Pythagoras exists.
E. <Pythagoras exists> is true *because* Pythagoras exists.
F. God loves X *because* X is good.

[1] Such as, amongst others, Schaffer (2009), Raven (2012) and Audi (2012).

[2] See for instance Jessica Wilson (2014).

[3] We suppose that objects, properties or facts of type x and type y are non-diachronically correlated if instances of x and y are correlated, and said instances do not occur at different times.

[4] We will use [square brackets] to indicate facts, which we take to be structured entities comprised of objects, properties and relations.

[5] We use 'because' as a neutral way of expressing these claims (i.e. a way that does not commit one to thinking there are grounding relations). Those who think that we need to posit a relation of grounds to explain these (and other) cases, will rearrange the relevant sub-sentential phrases and read 'because' as 'grounds'.

[6] We use <P> to indicate the proposition that P.

G. [Pythagoras exists] obtains *because* Pythagoras exists.
H. $2 + 2 = 4$ *because* 2 exists and 4 exists.

whilst judging the following to be false:

(a) The flower is maroon *because* the flower is red.
(b) The existence and arrangement of the wheels, spokes, handlebars, etc. exist *because* the bicycle exists.
(c) Pythagoras exists *because* <a man exists> is true.
(d) Pythagoras exists *because* {Pythagoras} exists.
(e) Pythagoras exists *because* <Pythagoras exists> is true.
(f) X is good *because* God loves X.[7]
(g) Pythagoras exists *because* [Pythagoras exists] obtains.
(h) 2 exists and 4 exists *because* $2 + 2 = 4$.

We call the combined judgements that (A) through (H) are true, and (a) through (h) are false, our *grounding relevant judgements.* We take it that in order to explain our grounding relevant judgements, we must explain both (a) why there is an appearance of asymmetry in these cases [i.e. we judge that (A) is true and (a) is false] and (b) why it seems to us that there is an explanatory connection between what lies to the left, and what lies to the right, of the 'because'.

We think that our grounding observations are exhausted by the following two kinds of observation:

(1) We observe that certain objects, facts, or properties, are non-diachronically correlated.
(2) We observe that there are widely shared grounding relevant judgements.

That is, we think that there is no other source of observational evidence for the presence of grounding relations than that adduced by (1) and (2).[8]

Our aim is to provide an explanation for (1) and (2) that does not appeal to the presence of grounding relations. In particular, we defend the following thesis:

> EDG: Our best explanation for our grounding observations appeals to the functioning of certain psychological mechanisms, and makes no mention of the presence of any relation of ground.

[7] Whether (F) is true or (f) is true is the focus of the *Euthyphro* dialogue (Plato 2002).

[8] One might think that there are jobs for grounding that go beyond explaining our grounding observations. Perhaps grounding is needed as a finer grained notion of dependence than modal dependence (Schaffer 2009), to back metaphysical explanations (Audi 2012), or to frame metaphysical positions (Raven 2012). Thanks to an anonymous referee for pointing this out. We think that, if our preferred explanation of the grounding observations is correct, the need for grounding to do these jobs is substantially undermined. That is, if our grounding relevant judgments are best accounted for in terms of overgeneralising psychological mechanisms, it is far less clear that we do in fact need a finer grained notion of dependence. Moreover, we argue in our [blanked] that the explanation we provide here naturally lends itself to a psychologistic theory of metaphysical explanation, which can account for the truth and falsity of metaphysical explanations, and frame metaphysical positions perfectly well, without buying into a grounding-based ontology.

We call this the *Explanatory Dispensability of Grounding*, since we take it to show that appealing to grounding relations is explanatorily dispensable. The bulk of the paper will be dedicated to providing our preferred explanation. The first two sections (Sects. 2 and 3) set out our explanation of our grounding observations in terms of the functioning of a set of two inter-connected cognitive mechanisms: a correlation detection mechanism and a causal detection mechanism.

Though our principal focus lies in defending EDG, we hope that doing so can lead to conclusions of more metaphysical import. It is this possibility we investigate in Sect. 4. We there consider three ways in which one might use our defence of EDG. The first two of these utilise EDG in arguments whose conclusion is that we should not posit a relation of ground: the first is an argument from the explanatory dispensability of grounding relations, and the second is a debunking argument against grounding. Our aim is not to defend either of these arguments (though we are partial to the first) but rather, to show how such arguments would proceed, and which premises their proponents would need to defend. Finally, we consider an alternative way in which one might use EDG. We suggest that the defender of grounding could use our explanation as an account of how we developed a mechanism to track grounding relations. We conclude by suggesting that whichever of these uses of EDG takes one's fancy, its defence should be of interest to the friend, and foe, of grounding alike.

2 General schema of an explanation

Recall that our aim is to defend EDG. That, in turn, involves explaining, without appealing to grounding, why:

(1) We observe that certain objects, facts, or properties, are non-diachronically correlated.
(2) We observe that there are widely shared grounding relevant judgements.

Before we lay out the general form our explanation will take, a few terminological clarifications are in order [following Russell's *The Principles of Mathematics* (1903)]. Symmetric relations are ones in which, for any x, y, if x R y, then y R x. Asymmetric relations are ones in which, for any x, y, if x R y, then it is not the case that y R x. And non-symmetric relations are ones in which, for some x, y, if x R y, then it is not the case that y R x. Hence, all asymmetric relations are non-symmetric, but not vice versa. In what follows there will be a need for us to extend this terminology to apply to instances of relations. Accordingly, we will call an instance of a relation in which x R y, and y R x, a *symmetric instance* of that relation, or, alternatively, we will say that the relation *obtains symmetrically*. Likewise, we will call an instance of a relation in which x R y and it is not the case that y R x, a *non-symmetric instance* of the relation, or, alternatively, we will say that the relation *obtains non-symmetrically*. Thus, a symmetric relation will only have symmetric instances, an asymmetric relation will only have non-symmetric instances, and a non-symmetric relation will have at least one non-symmetric instance. Given this terminology, there is no such thing as an asymmetric instance of a relation:

asymmetry is a property of relations, not instances. We hope that this terminology draws attention to the fact that non-symmetric *instances* of a relation share the same formal features, whether they are instances of a non-symmetric relation or an asymmetric relation. That will subsequently become important.

Traditional modal relations are non-symmetric: they have both symmetric and non-symmetric instances. It can be that A necessitates B, but B does not necessitate A. But it can also be that X necessitates Y, and Y necessitates X. *Mutatis mutandis* for supervenience and entailment. This is in contrast to both grounding and causation, which are taken to be asymmetric.[9] Indeed, it has become commonplace to suppose that one feature of dependence relations is their asymmetry. If that is right, then none of the traditional modal relations is a dependence relation, while causation and grounding are both kinds of dependence relation.[10]

With this in mind, we aim to explain our grounding observations in terms of the functioning of two cognitive mechanisms: a correlation detection mechanism and a causal detection mechanism, where the latter acts as a filter on the outputs of the former. We introduce the correlation detection mechanism in Sect. 2.1. Then in Sect. 2.2 we argue that this mechanism detects both diachronic and non-diachronic correlations: the working of this mechanism explains (1) above. In Sect. 2.3 we argue that the causal detection mechanism acts as a filter on the correlation detection mechanism, paving the way for our suggestion, in Sect. 2.4, that the causal detection mechanism has been co-opted to filter the non-diachronic correlations detected by the correlation detection mechanism. We argue that in detecting non-diachronic *correlations* we sometimes thereby detect what we will call non-diachronic *relations*, some[11] of which, in turn, are modal relations. (Henceforth we use 'non-diachronic relations' to pick out this total set of relations, and 'modal relations' to pick out the sub-set of these that are modal relations.)

When the causal detection mechanism filters the non-diachronic correlations, it does so by seeking the same sorts of cues it uses to separate causal relations from mere diachronic correlations. Since any instance of causation is non-symmetric, the causal detection mechanism searches for instances of non-symmetrical relations amongst the correlations; in doing so it detects instances of causation. When it

[9] Though see, e.g., Rodriguez-Pereyra (2015) for a defence of some putative symmetric instances of grounding. If Rodriguez-Pereyra is right, then grounding is a non-symmetric relation. However, if grounding is non-symmetric, it is less well placed to explain why we have no symmetrical grounding relevant judgements. Thus, in what follows, we will assume that grounding is asymmetric, so as to criticise the strongest version of our opponent's position.

[10] Perhaps some non-symmetric relations are dependence relations. Defenders of grounding suppose that they are not, though not everyone agrees. For present purposes we will grant to the defender of grounding that only asymmetric relations are dependence relations, but nothing we say hangs on this. If some non-symmetric relations are dependence relations it is plausible that the traditional modal relations are dependence relations. In that case, our view is that the grounding observations can be accounted for in terms of traditional modal dependence relations and evolved psychological mechanisms, without any mention of a *further* asymmetric dependence relation in the form of grounding.

[11] We say 'some' here because there are many actual non-diachronic correlations that are not instances of modal correlation: for instance, the correlation between the set containing me, and the set containing you. Yet, correlations such as this are not the kind that attracts the attention of our correlation detector. Those that do tend to be indicative of modal correlations.

applies the same procedure in the case of non-diachronic correlations, it filters non-symmetric instances of non-diachronic relations from symmetric instances. Where it does so successfully, our grounding relevant judgements are explained by us successfully detecting non-symmetric instances of a non-diachronic relation: a modal relation [this is how we explain cases (A) and (B)]. We explain the appearance of *asymmetry* expressed by our grounding relevant judgements by noting that the relation in question holds *non-symmetrically* in such cases, and a non-symmetric instance might be an instance of either a non-symmetric[12] relation, or an asymmetric relation. Since the mechanism is tuned to detect causal relations (which are asymmetric), it is unsurprising that we infer that the instances thus filtered are instances of an asymmetric, rather than a non-symmetric, relation.

Further, we explain the appearance of *explanation* in these cases as the result of a trigger produced by the causal detection mechanism. That mechanism evolved to signal the presence of a dependence relation (causation), which *does* back explanation, and, in the process, to trigger a phenomenology as of there being an explanation present. When the same mechanism detects the presence of a non-symmetric instance of a modal relation, that same phenomenology is triggered. Thus, with respect to (A) and (B), we have explained (2), above.

In Sect. 2.5 we extend our explanation to appeal to empirical evidence which shows that our causal detection mechanism sometimes overgeneralises, signalling the presence of a non-symmetric instance of a relation—causation—where no such instance obtains. It does so because failing to detect an instance of causation that obtains, is more costly than signalling the presence of causal relations where there are none. If such mistakes occur when filtering the diachronic correlations, there is every reason to suppose they also occur when filtering the non-diachronic correlations, since the latter filtration utilises the same mechanisms as the former, co-opted for a slightly different purpose. So we should expect cases in which our causal detection mechanism mistakenly signals, amongst the non-diachronic relations, the presence of a non-symmetric instance where no such instance obtains.

Given this, we argue there are environmental conditions under which we should expect the causal detection mechanism to signal that an instance of a non-diachronic relation is non-symmetric. In Sect. 3 we argue that these conditions are those we find obtaining with respect to the objects, properties or facts mentioned in (C) through (H). So we should *expect* that, having detected the correlations in those cases, the causal detection mechanism will signal that those correlations imply the obtaining of *non-symmetric instances* of a modal relation. And that is exactly what we find. But we would expect this even on the supposition that the instances of those modal relations are symmetric. So we can explain our judgements regarding cases (C) through (H), by appealing to nothing more than the functioning of some of our cognitive mechanisms, and the existence of the traditional modal relations.

We explain the appearance of *asymmetry* expressed by our grounding relevant judgements in the same way it is explained in cases (A) and (B), except in (D) through (H) the relevant mechanism *mistakenly* signals the presence of a non-

[12] But not asymmetric.

symmetric instance of a relation (rather than correctly signalling its presence). With respect to (C) the mechanism gets the right answer—it detects a non-symmetric instance of a modal relation—but it does so for the wrong reasons, namely, because it is sensitive to cues that would have resulted it in signalling the presence of such a relation even if one had not been present. Finally, we explain the appearance of *explanation* in the same way as it was explained for cases (A) and (B), as the result of a trigger produced by the causal detection mechanism. By the end of Sect. 3 we will have explained both (1) and (2) above. So if our explanation is right, then EDG is true.

2.1 The correlation detection mechanism

There is a good deal of evidence that we have a correlation detection mechanism, though exactly how that mechanism functions (and which brain processes subserve it) is more controversial. What is agreed is that this mechanism allows us to distinguish information bearing patterns from random patterns, and to quantify the information bearing patterns. It functions by taking in inputs, namely, the frequency of the presence, or absence, of certain features in the environment. In the literature, these frequencies are represented by what is known as a contingency table—a matrix that displays the frequency distribution of variables. A simple version of a contingency table is below (Fig. 1).

Attempts to understand the correlation detection mechanism focus on determining which heuristic we use to detect correlations. The assumption is that we represent (perhaps sub-personally) something like a contingency table of data, then use a heuristic to determine whether a correlation obtains between the relevant data. Investigators have identified four candidate heuristics. The first (known as the Cell A rule) focuses entirely on the frequencies in Cell A (see, e.g., Smedluns 1963; Nisbett and Ross 1980). The second, called the A minus B rule, holds that we are sensitive to frequencies in Cells A and B: the more the frequencies in cells A and B

Fig. 1 From Arkes and Harkness (1983)

		Factor 2	
		Present	Absent
Factor 1	Present	A	B
	Absent	C	D

diverge the higher the correlation is judged to be (Shaklee and Mims 1982). The third rule is known as the sum of diagonals. Here, we compare A + D, with B + C (Shaklee and Tucker 1980). The strength of the correlation signalled depends upon the extent to which these sums differ. According to the fourth rule, we use conditional probabilities (a Bayesian calculation) to evaluate two competing hypotheses: H_1 is the hypothesis that the data was produced by a random process, and H_2 is the hypothesis that the data was produced by some systematic process. The mechanism then uses Bayes' rule to combine prior beliefs about these hypotheses, with the evidence the data provides. More recently, studies suggest that subjects use all of these rules flexibly (Arkes and Harkness 1983).

In what follows we argue that the correlation detection mechanism is the basis of our ability to detect causal relations, and that it is the mechanism responsible for the detection of non-diachronic correlations. For the former claim to be plausible, it needs to be that the correlation detection mechanism is *good enough* at detecting correlations that, when filtered, it will generate the observed causal judgements. We think that very plausible. Research on the correlation detection mechanism often focuses on why the mechanism is so *inaccurate*. This inaccuracy, however, takes two forms. The first is the presence of false positives: cases where the mechanism signals the presence of correlations where none exist (known as illusory correlations; see Redelmeier and Tversky 1996). The second are not false positives per se; rather, they involve the correct detection of a correlation, but an over-estimation of the strength of the correlation (Chapman and Chapman 1967). It is easy to see why misjudgements of strength will sometimes occur if we use either the Cell A heuristic or the A minus B heuristic, since in either case we are ignoring important data (that contained in cells C and D).

Notice, though, that what matters for our purposes is that the correlation detection mechanism typically detects correlations that are there, not that it is always accurate in detecting their strength. If the causal detection mechanism filters the outputs of the correlation detection mechanism, then it is a virtue if the latter system is highly sensitive—if it is more inclined to produce false positives than false negatives—since false positives can be filtered out by the causal detection mechanism. One way in which the mechanism is thought to be highly sensitive is that (at some sub-personal level) we deploy Bayes' theorem (see Williams and Griffiths 2013). Since the likelihood of most data sets is higher on the hypothesis that the data is non-random, than that it is random, the mechanism tends to yield false positives. Another way to put this is that some kind of pattern can be detected in almost any data set, and the probability of that data set conditional on its being the result of some structure in the world, is typically higher than the probability of that data set conditional on it not being the result of some structure in the world. What this means is that since the mechanism is highly sensitive to possible patterns in data, and will readily signal the presence of those patterns. This feature (as we will now argue) makes it ideal for detecting both diachronic and non-diachronic correlations.

2.2 Detecting non-diachronic correlations

That the correlation detection mechanism is very sensitive to patterns gives us good reason to suppose that it will be successful at detecting both diachronic and non-diachronic correlations. After all, if the correlation detection mechanism works by looking at the frequencies of certain events/properties/objects, then it ought not matter whether those events (etc.) are temporally separated (in the diachronic case) or not (in the non-diachronic case). The only respect in which non-diachronically correlated relata differ from diachronically correlated relata is that in some cases of non-diachronic correlation one relatum is unobservable. It is worth noting that these are the minority of cases. We can observe instances of blue and azure. We can observe instances of bicycles, and parts arranged bicycle-wise. In *some* sense we can observe that a sentence is true, and that a particular state of the world obtains. But we cannot observe sets, and we cannot observe numbers. So if correlation detection proceeds via us representing something like a contingency table, and if we always detect frequencies by observation, then it will be impossible for our correlation detection mechanism to detect correlations between entities, where (at least one) of those is unobservable.

But, we think, there is little evidence that frequencies must *always* be detected by observation. If some entities or properties are unobservable, then we must come to know that they exist through methods other than observation (whatever these might be). Assuming that we can come to know of the existence of such entities in *some* way, there is no reason to suppose that we cannot come to know frequency data through the same method. Given this, we argue, there is good reason to think that our correlation detection mechanism detects both diachronic and non-diachronic correlations, particularly given that there is utility in detecting both. For instance, once we notice that there is a synchronic correlation between the existence of the chair and the existence of a set of parts arranged in a certain way, the possibility opens up of intervening on the chair by intervening on the parts. Thus we have an explanation for (1): our observation that some entities are non-diachronically correlated. That only leaves us needing to explain (2), above. To do so, we appeal to the functioning of our causal detection mechanism, which we describe in the following section.

2.3 The causal detection mechanism

There is overwhelming evidence that we have cognitive mechanisms adapted to identify and track *causal dependencies*, as well as mechanisms that produce sophisticated *causal reasoning*. It seems to us that the best way to understand such claims is to suppose that our causal detection mechanism operates by filtering the outputs of our correlation detection mechanism. The latter signals the presence of (*inter alia*, diachronic) correlations. The causal detection mechanism then searches for cues to filter out those that are *mere* correlations, leaving those which are indicative of a causal relation. It does so by searching for *non-symmetries* amongst the correlations. We take it that searching for non-symmetries crucially involves searching for correlations in which: (a) changing X is likely to change Y and b) changing Y does not change X and (c) X occurs before Y (Sloman 2005). The idea, here, is that the mechanism searches for non-symmetries amongst the correlations because any instance of the causal relation is non-symmetrical. So the mechanism aims to detect non-symmetries amongst the correlations, as a way of detecting causation.

There is a range of empirical data that supports the idea that the causal detection mechanism seeks out non-symmetries via various environmental cues. One such cue is the way in which the environment responds to an intervention. Since more than one causal model is consistent with any observed correlation, the only way to discover which causal model is the right one is for agents to perform an intervention.[13] This is because interventions cut off the thing upon which one intervenes, from any prior causes, but not from any later effects and thus have the capacity to reveal asymmetric dependencies.[14] As Hagmeyer et al. put it:

> Interventions often enable us to differentiate amongst the different causal structures that are compatible with an observation. If we manipulate an event A and nothing happens, then A cannot be the cause of event B, but if a manipulation of event B leads to a change in A, then we know that B is a cause of A, although there might be other causes of A as well. (2007:87)

The process of intervention is, in effect, a process that aims to determine whether there is a non-symmetry present. If intervening on A intervenes on B, and not the converse, there is a non-symmetry present. Interventions, then, are one cue that the causal detection mechanism uses, to filter correlations. Where there is a non-symmetry detected, via intervention, this is a cue that the correlation is associated with a non-symmetric instance of some relation: in this case causation. A second cue is temporal order (Sloman 2005: 6). Here again, the aim is to detect a non-symmetric instance of a relation, among the correlations. If x occurs before y, then y does not occur before x. This non-symmetry is a cue that there is a non-symmetric instance of a relation present: causation. Finally, there are other cues that, among the correlations detected, some are backed by non-symmetric instances of a relation; namely prior knowledge, and an existing hypothesis about causal structure (Waldmann and Hagmayer 2013:745). Where some (or all) of these cues are present, the causal detection system signals the presence of a non-symmetric instance of a relation amongst the diachronic correlations. The non-symmetric instance in question is, in each case, an instance of causation.

2.4 Co-opting the causal detection mechanism

The next step of our explanation is to argue that the causal detection mechanism has been co-opted to filter the *non*-diachronic correlations detected by the correlation detection mechanism. First, however, we suggest that in detecting non-diachronic *correlations* we are sometimes detecting non-diachronic *relations*, of which the

[13] See Gopnik et al. (2004), Kushnir et al. (2010), Lagnado and Sloman (2004), Styvers et al. (2003).

[14] Sloman (2005).

modal relations are a subset. Just as sometimes our detecting of a diachronic correlation is, *ipso facto*, detecting a causal relation, and just as detecting a particular determinable is, *ipso facto*, detecting a particular determinate, so too sometimes detecting instances of a non-diachronic correlation is, *ipso facto*, detecting a non-diachronic relation, and hence, in some cases, detecting a modal relation. If that is right, then our correlation detection mechanism is a mechanism that allows us to detect modal correlations. So we have an explanation for how it is that we track those relations (or at least, their actual world instances).

In what follows we argue that our causal detection mechanism filters the non-diachronic correlations. If what we have just said is right, then in filtering these correlations it thereby filters instances of modal relations into those that are symmetrical and those that are non-symmetrical. Now, one might object, causation is an *asymmetric* relation. How could a mechanism evolved to track causal relations do the work we are suggesting? Well, notice that our detection mechanisms detect, and filter, *instances* of relations. So the causal detection mechanism is really a mechanism evolved to detect non-symmetric instances of a diachronic relation. By reliably detecting these non-symmetric instances, the mechanism is thereby detecting an asymmetric relation: causation. Since that mechanism is sensitive to the formal features of *non-symmetry*, features shared by both diachronic and non-diachronic relations, it is easy to see how it could be co-opted to track non-symmetrical non-diachronic relations, of which the relevant relations, for our purposes, are the modal relations. Now, it might be that detecting non-symmetrical instances of modal relations is an adaptation of the causal detection mechanism. At worst, we think, our capacity to filter non-diachronic correlations is an *exaptation* of our causal detection mechanism. Exaptations are traits that are a by-product of adaptive selection (they are not selected *for*) but which nonetheless come to be useful to the organism (Gould 1991: 43).[15]

How does the causal detection mechanism filter the non-diachronic correlations? Well, the sorts of environmental cues that allow it to discern whether there is an underlying non-symmetric instance of a relation are, by and large, the same sorts of cues for both diachronic and non-diachronic correlations. As we briefly noted previously, there are four important environmental cues to which the causal detection mechanism is sensitive.

(i) Temporal order
(ii) The result of intervention
(iii) Prior knowledge
(iv) An existing hypothesis about causal structure

Clearly (i) will only apply in the case of diachronic correlations. The remaining three cues, however, are relevant. Consider, first, interventions. Return to our example of the chair. Upon noticing that there is a correlation between the chair and

[15] There has been an attempt to explain our tendency towards religious belief as a spandrel (an exaptation that is not useful): it arises not because it is adaptive but as a by-product of a host of other cognitive processes (Atran 2002; Barrett 2004; Boyer 2001; Pyysiäinen 2001; Pyysiäinen and Anttonen 2002; Gould 1991:58). Notice that if this explanatory strategy is right, we do not need to suppose that we are good deity-trackers and that there are deities.

its parts we can engage in an intervention. We can notice that there is no way of intervening on the chair without intervening on its parts. We can see this by trying to wiggle one thing (the chair) *without* wiggling the other (its parts) and by trying to wiggle one thing (the chair) *by* wiggling the other (its parts). So the response of the environment to interventions can serve as a cue to the co-opted causal detection mechanism. So too, presumably, can prior knowledge. Just as one might have causal knowledge that one can bring to bear in determining whether a diachronic correlation is due to causation, so too one might have knowledge one could bring to bear in determining whether a particular non-diachronic correlation is due to a non-symmetrical instance of a modal relation. Perhaps once I see the non-symmetric relation between chairs and their parts, I find it easier to see the non-symmetric relations between other sorts of objects and their parts. Finally, an analogue of (iv) might be relevant. It might be that an existing hypothesis about *modal* structure serves as a cue to help filter the non-diachronic correlations. For instance, I might have a prior hypothesis that token mental events are identical to token physical events; or I might have a prior hypothesis that any physical event like this one, will be correlated with a mental event like this one, but not vice versa. Either of these existing hypotheses might serve as a cue to the causal detection mechanism.

So far, then, we have argued that in detecting non-diachronic correlations, we thereby (sometimes) detect modal relations, and, in filtering those non-diachronic correlations we thereby identify non-symmetric instances of modal relations. Thus we are on our way to explaining (2), our observation that there are widely shared grounding relevant judgements. In particular, we can now explain some of those judgments: namely (A) and (B). These are cases in which there is a non-diachronic correlation between the events, properties, or facts, in question. Features of these correlations correctly cue the co-opted causal detection mechanism to signal the presence of a non-symmetric instance of a modal relation.[16]

Consider (B). At every world where those bicycle parts exist and are arranged appropriately, there is a bicycle, but there are worlds where the bicycle is composed of different parts, so the existence of the bicycle does not guarantee that *those* parts exist in *that* arrangement.[17] Plausibly, our causal detection mechanism is cued to this non-symmetry via the result of interventions (some counterfactual) since the only way of intervening on the bicycle is by intervening on its parts.

Now consider (A). Every possible maroon flower is a red flower, yet there are red flowers that are not maroon (crimson flowers, for example). Thus necessitation obtains non-symmetrically. The only way to intervene on redness is to intervene on a determinate property. Of course, not every way of intervening on a determinate property will change whether redness obtains (it will just change which shade obtains). But since the only way to intervene on redness is to intervene on a determinate property, the interventionist information will cue the presence of a non-symmetric instance of a modal relation.

[16] We assume that there is nothing problematic in appealing to traditional modal relations in the absence of grounding. However, one could object that grounding is required to account for the patterns of modal co-variation described by these relations. We give this objection an extended treatment in our [blanked].

[17] We are assuming that the fact that the bicycle exists does not rigidly designate that particular bicycle.

In each case we explain the apparent *asymmetry* by noting that each instance of the modal relation in question is a non-symmetric instance. Since it is non-symmetric instances of causation that the mechanism evolved to track, and since causation is an asymmetric relation, it can hardly be surprising that we mistakenly conclude that the relation we are tracking in these cases is itself asymmetric (when in fact it is non-symmetric, not asymmetric).

What explains our belief that there is an explanatory relation obtaining between the relata is that the signal from the causal detection mechanism triggers, or is otherwise associated with, a phenomenology as of such an explanatory relation obtaining. It is associated with that phenomenology either because it directly causes it, or because what cues the signal is a common cause of the signal and the phenomenology. It makes good sense for the output of the mechanism to trigger that phenomenology because the mechanism evolved to track causal dependencies, and those dependencies are genuinely explanatory. But the signal triggers the phenomenology even when it is filtering non-diachronic correlations, and so we experience a phenomenology as of the obtaining of an explanation when the causal detection mechanism cues us to the presence of a non-symmetric instance of a relation.

What remains to be explained, then, are judgements (C) through (H). In what follows we argue that what explains the apparent asymmetry, and the phenomenology of explanation, is the same for these cases as in (A) and (B): the functioning of the causal detection mechanism. The difference lies in the fact that in cases (D) through (H) the mechanism *mistakenly* signals the presence of a non-symmetric instance of a relation, where no such instance obtains. In case (C), the mechanism correctly identifies a non-symmetric instance of a modal relation, but, as we will show, it does so for the wrong reasons. We arrive at these mistaken judgements because, as we will now argue, the causal detection mechanism overgeneralises.

2.5 A filter that overgeneralises

The next component of our explanation appeals to empirical evidence to show that our causal detection mechanism sometimes *overgeneralises*. This can hardly be surprising. Evidence suggests that causal reasoning—that is, reasoning in terms of causal models—is typically very successful: it affords agents a good deal of predictive and explanatory power, and is often fast, automatic and unreflective (Sloman 2005: 77–78, 80). According to some, the use of causal models is fundamental to how we understand the world (Schafer 1996).

Here is something we know. In general, we ought to expect that where the costs of a false negative significantly outweigh the costs of a false positive, we typically develop cognitive systems that set the threshold for detecting that stimulus quite low.[18] Consider life as an animal that is predated upon. The cost of failing to detect a predator could well be death. The cost of misidentifying something as a predator is

[18] Notice that we are making a very general claim here: we are not defending anything like the thesis that we should expect cognitive systems to be optimal in the manner in which they make these trade-offs. Clearly there are developmental constraints on the ways in which cognitive systems can solve problems which mean that systems often are not optimal.

not nil (since you might run away, thus using up energy) but it is lower than failing to detect a predator. For this reason, prey animals have predator detection systems that have a very low threshold for detecting predators. That is, the features that something in the environment needs to have in order to set off the predator detection system are relatively minimal. In cases such as these we will say that the cognitive systems in question tend to *overgeneralise*.

There are plenty of examples of overgeneralisation in the human cognitive system. Detecting faces is important. So we have very sensitive facial detection systems that can be triggered by something as simple as an arrangement of three dots in a particular configuration (roughly a triangular configuration).[19] The same can be said, *mutatis mutandis*, for our agency detection module (Guthrie 1993).[20]

There is strong evidence that our causal detection mechanism overgeneralises in this way. We sometimes impose causal structure where none exists. For instance, we seem to perceive causation where there is none. Subjects reliably describe the interaction of two moving dots on a screen in terms of one dot causing the other dot to act in certain ways (Michotte 1956). People impose beliefs about the causal structure of the world onto the correlational data they are trying to understand (Waldman 1996) even when imposing a causal framework distorts their representation of the world. For instance, mathematical equations are symmetric, in the sense that any variable can appear on either side of an equality.[21] Despite this, subjects find certain ways of expressing an equation more natural than others: namely, those ways that fit best with their causal model (Sloman 2005: 72). It has also been argued that the overgeneralisation of our causal detection mechanism leads us to misjudge probabilities. For instance, suppose subjects are asked which of the following is more probable:

(1) A man has a history of domestic violence if his father has a history of domestic violence.
(2) A man has a history of domestic violence if his son has a history of domestic violence.

Subjects report that (1) is more probable than (2), even though they are equally probable. It is thought that (1) appears more probable because the direction of causation goes from father to son, not son to father. The causal detection mechanism responds to certain cues present in the presentation of this data, and signals that these probabilities are non-symmetric, when they are not.

What sorts of cues might lead our causal detection mechanism astray in these cases? Recall the four cues to which the mechanism is sensitive.

---

19 Rigdon et al. (2009), for example, showed how such an arrangement of dots significantly increased participants' giving behaviour. They hypothesised that this is because the dots are sufficiently face-like to cue people to act as though someone is watching to see how generous they are.

20 Indeed, there are those who think that our tendency towards religious belief is to be explained by an overactive agency detection model which 'detects' agency where there is none, leading us to posit supernatural agents (Barrett 2004).

21 Algebra is all about how to permute the order of the variables without changing the relations the equation expresses.

(i) Temporal order
(ii) The result of intervention
(iii) Prior knowledge
(iv) An existing hypothesis about causal (modal) structure

Consider (i). The event of the son engaging in domestic violence occurs after the event of the father engaging in domestic violence (at least in most cases). So temporal order will cue the mechanism. Consider (ii). Evidence shows we are adept at doing counterfactual interventions (Sloman 2005: 80). Were we to imagine performing an intervention in this case, we would likely conclude that if we want to change whether a son is violent, we intervene upon whether his father is, not vice versa. So interventionist information will cue the mechanism. Consider (iii): prior knowledge. Some people will know that one way in which people become abusive is by witnessing it in the home, and that prior knowledge might feed into (iv) and create an existing hypothesis about the causal structure—namely that the father's being abusive causes the son to be abusive. All of these cues result in the causal detection mechanism signalling the presence of a non-symmetric instance of a relation, and the salience of that instance overrides the symmetrical correlation that is relevant in making probability judgements, leading to mistaken judgements.

Prima facie, then, if the causal detection mechanism overgeneralises in signalling the presence of non-symmetric instances of a diachronic relation, we should expect it to similarly overgeneralise in signalling the presence of non-symmetric instances of non-diachronic relations. The next step of our explanation builds on this idea by showing that there are certain environmental conditions under which we should expect the causal detection mechanism to signal the presence of a non-symmetric instance of a modal relation amongst the non-diachronic correlations whether or not such an instance obtains—and these conditions are precisely those associated with cases (C) through (H).

3 Fooling the co-opted causal detection mechanism

Cases (D) through (H) are ones in which a *symmetric* instance of a modal relation obtains between the relevant objects, properties, or facts. (For example, every world in which Pythagoras exists is a world in which {Pythagoras} exists, and vice versa). Though we are inclined to say that an asymmetric explanatory relation obtains, it cannot be that we are correctly tracking non-symmetric instances of modal relations, since no such instances obtain. By contrast, in case (C) the relevant modal relation obtains non-symmetrically. Every world in which Pythagoras exists is a world in which <a man exists> is true. Yet there are worlds in which Pythagoras does not exist, but <a man exists> is true nonetheless, due to the presence of some other man. However, it doesn't seem right to say that we are 'successfully tracking' this non-symmetrical instance, since (C) through (H) have features that we should expect to cue the causal detection mechanism and result in it signalling the presence of a non-symmetric instance of a modal relation *whether or not one obtains.* To see why this is so, we will go through each case.

Consider cases (C) (D) and (E).

C. The proposition <a man exists> is true *because* Pythagoras exists.
D. {Pythagoras} exists *because* Pythagoras exists.
E. The proposition <Pythagoras exists> is true *because* Pythagoras exists.

Consider the kinds of cues that can trigger the filtering system. (i), temporal order, is irrelevant. Consider (ii): The result of intervention. We have already seen that the result of actual (or counterfactual) interventions can help trigger the causal detection mechanism (recall the case of the chair and its parts). Interventions on non-diachronic correlations are, however, not always straightforward. With respect to some of those correlations, one of the correlated properties, facts or objects, cannot be intervened upon. This is what we find in cases (C) through (E), where one relatum is an abstract object, and is, therefore, an object we cannot intervene on.

But consider how interventions work to provide cues to the causal detection mechanism. In the case of diachronic correlations between, say, events of kind x and kind y, we are able, on one occasion, to intervene on a token x and see whether this wiggles a token y, and on another occasion to wiggle a token y, and see if this wiggles a token x. It is the result of this pair of interventions that acts as a cue for the causal detection mechanism. In particular, if in wiggling x we can wiggle y, but not vice versa, this tends to cue the mechanism to signal that there is an underlying non-symmetric instance of a relation.

In cases (C) through (E) it is clear enough that we can intervene upon whether Pythagoras exists, and thereby wiggle whether {Pythagoras} exists or <Pythagoras exists> is true or <a man exists> is true.[22] Furthermore, it is false that we can intervene upon whether {Pythagoras} exists or <Pythagoras exists> is true or <a man exists> is true and thereby wiggle whether Pythagoras exists. This is because propositions and sets are abstract objects upon which we cannot intervene. We can't wiggle whether {Pythagoras} exists in order to wiggle whether Pythagoras exists because we can't wiggle whether {Pythagoras} exists at all! In this way, the causal detection mechanism gets the same non-symmetric feedback about interventions that it does when there are causal relations underpinning a diachronic correlation, but for a very different reason.

Causal relations do not obtain between abstract relata, and thus the causal detection mechanism has not evolved to distinguish between a situation where wiggling y does not wiggle x and a situation where y simply cannot be wiggled. Thus, triggered by an apparent non-symmetry in the results of counterfactual interventions, the mechanism signals that the correlations in these cases are indicative of an underlying non-symmetric instance of a modal relation. In cases (D) and (E) the signal is mistaken: the modal relation obtains symmetrically. Fortuitously, in (C) it correctly signals the presence of a non-symmetric instance of a modal relation. It gets things right, in that case, but for the wrong reasons: it would signal the presence of non-symmetry even if no such non-symmetry were present.

[22] Not every intervention on Pythagoras will wiggle <a man exists> but one way to wiggle <a man exists> is to make it the case that Pythagoras exists.

For that reason we will not describe this as a case of the mechanism successfully tracking a non-symmetric instance of a modal relation.

Let us now consider (F), the Euthyphro case:

F. God loves X *because* X is good.

We are asked to imagine that God and goodness co-exist in all worlds, and to wonder whether things are good because God loves them, or He loves them because they are good. This is a case in which there are no cues of kind (i) or (ii) (there are no cues of kind B because we cannot intervene on *either* relatum). But consider (iii) and (iv):

(iii) Prior knowledge
(iv) An existing hypothesis about modal structure

Both of these kinds of cues could be expected to play a role in determining the output of the causal detection mechanism. Plausibly, we should expect prior causal knowledge to influence a subject's existing hypothesis about modal structure. Each of us is familiar with agents' intentional states depending on the way the world is. (We typically hope our beliefs are like this.) If this is the most salient piece of a subject's prior knowledge, it might lead her to have a hypothesis about modal structure, according to which just as the causal direction of fit goes from the world, to the mind so too does the modal direction of fit. Subjects will have a prior hypothesis that God's attitudes depend on the distribution of goodness, and in the absence of any other cue to the contrary, this will cue the causal detection mechanism to signal the presence of a non-symmetric instance of a modal relation.

On the other hand, subjects are also familiar with cases in which an agent's intentional states cause the world to be a certain way (we typically hope our desires are like this, when appropriately conjoined with our beliefs). Where *this* causal knowledge is most salient, we can expect it to lead to a hypothesis about modal structure according to which, just as the direction of causal fit goes from God's attitudes to goodness, so too does the direction of modal fit. In the absence of any other cues to the contrary, this hypothesis about modal structure also cues the causal detection mechanism to signal the presence of a non-symmetric instance of a modal relation.

Notably, actual judgements about this case vary; we think this could have been predicted, based on the fact that neither (i) nor (ii) offer us any useful cues, and that people could be expected to form very different modal hypotheses given different salient background causal knowledge, when confronted with the description of the case.[23]

Now consider (H):

H. $2 + 2 = 4$ *because* 2 exists and 4 exists.

This is another case in which we cannot intervene on either relatum. So (i) and (ii) provide no cues. But, again, we might expect prior knowledge to lead to the

[23] In this respect, our explanation does better than the grounding-based explanation of our grounding relevant judgements. For, if our intuitions about such cases are to be explained in terms of our successfully tracking grounding relations, we would not predict the kind of disagreement we see regarding the Euthyphro case.

triggering of the causal detection mechanism. We have prior knowledge of relations and relata, itself stemming from experience with interventions. We typically know (even if only implicitly) that in order to intervene on some relation, we need to intervene on the relata. If we want to intervene on the relation of 'being next to' obtaining between Bill and Ben, we can only do this by moving either Bill, or Ben. We can't directly intervene on the 'being next to' relation. So prior knowledge tells us that, in general, if you want to intervene on the obtaining of a relation, then you need to intervene on the relata. Addition is a relation. So prior knowledge suggests that in order to intervene on whether it obtains between 2 and 2, one would need to intervene on the relata of the addition function. Thus in the absence of being able to discern any other cues, this prior knowledge leads to a hypothesis about modal structure which cues the causal detection mechanism, which indicates that there is a non-symmetric instance of a relation here, and that (H) is true.

This cue, however, is relatively weak. After all, cues of kind (i) and (ii) are absent. So it is relatively easily overturned. We can witness that by the fact that mathematical structuralists find (h), rather than (H), intuitive. Their prior knowledge has altered the cues that the causal detection mechanism receives, and changed the output so that although it signals a non-symmetric instance of a modal relation, the 'direction' of the non-symmetry is taken to be the reverse of (H).

Moving on, let's revisit the relationship between a fact and the fact that that fact obtains:

G. [Pythagoras exists] obtains *because* Pythagoras exists.[24]

We believe there is an intuitive pull to suppose that one would intervene on the fact on the right, in order to intervene on the fact on the left, but not vice versa. Here is why we might expect to have that intuition. The fact on the right is a *constituent* of the fact on the left. Talk about facts is sophisticated, and almost certainly something we come to after we have learned about more mundane things such as, saliently, parts and wholes. Prior knowledge of wholes and parts tells us that we intervene on the parts to intervene on the whole. The logical form, if you will, of (G) strongly suggests that we ought to intervene on the fact on the right, in order to intervene on the fact on the left. Since we have little other experience at intervening on facts, the logical form of (G) is a salient cue, which leads us to make a particular hypothesis about modal structure, and ultimately cues the causal detection mechanism to categorise (G) as more than mere correlation.[25]

[24] We take the expressions flanking the 'because' here to pick out facts. This can be made clear by expressing what we take to be the equivalent claim: [[Pythagoras exists] obtains] *because* [Pythagoras exists].

[25] Similar considerations apply, we think, to the putative explanation of a 'conjunctive fact' in terms of its conjuncts (see Raven 2012). That is, prior knowledge of wholes and parts tells us that we should intervene upon a conjunct—a 'part'—in order to intervene on the conjunction—the 'whole'.

3.1 Summing up

So far we have shown why we should expect our causal detection mechanism to signal the presence of a non-symmetric instance of a modal relation when presented with non-diachronic correlations in certain conditions; those conditions we find obtaining between the objects, properties, and facts, mentioned in (A) through (H). In cases (A) and (B) this is (in part) because there really is a non-symmetric instance of a modal relation. In cases (C) through (H) by contrast, we should expect our causal detection mechanism to signal the presence of non-symmetry whether or not it obtains [and it does obtain, but only in (C)]. On the assumption that these non-symmetric instances do not obtain, however, this overgeneralisation is largely cost free in the following sense: feedback that can be gained via interventions cannot reveal that the system is generating false judgements because the relevant interventions that would reveal this cannot be performed. Thus the environment can never provide feedback that would allow us to correct that judgement. So there is, effectively, no cost associated with the causal detection mechanism erroneously being triggered by these correlations.

Thus, our appeal to the sorts of cues the environment sends the causal detection mechanism has a dual role. On the one hand, it explains why we should expect to have the intuitions we do, regardless of whether there are non-symmetric instances underpinning these correlations. But examination of those cues also reveals why there is no cost to being misled in this way, on the assumption that we are, indeed, being misled (that is, on the assumption that there is no grounding relation that we are tracking): namely, that we should also expect this overgeneralisation in such cases, because there is simply no cost to being wrong.

So our defence of EDG comes to a close. We hope we have explained our grounding observations in a way that appeals, in an indispensable manner, to the existence of the traditional modal relations and certain psychological mechanisms, but does *not* appeal, in an indispensable manner, to the existence of any relation of ground. In what follows we briefly consider the implications for grounding, if one accepts EDG.[26]

4 Where does that leave grounding?

Suppose you buy our argument so far. What, if anything, does that tell us about grounding? That depends. In what follows we first outline two ways in which one might deploy EDG to argue that there is no relation of ground. The aim is not to

[26] One might worry that, unlike other cases where the causal detection mechanism overgeneralises, our grounding observations in cases (D)–(H) are highly resistant to change on the basis of reflection. So, for example, making salient the probabilistic symmetry between abusive fathers and sons effectively constrains our tendency to overgeneralise. One difference is that in our explanation of (D)–(H), there are not two relations, one symmetric and one non-symmetric, such that the salience of one swamps the salience of the other. Instead, a single, symmetric relation, is misclassified by a cognitive system as non-symmetric. We know that the output of some sub-personal cognitive systems are (largely) immune to change on the basis of personal-level reflection. No amount of reasoning makes the lines look the same length in the Muller-Lyre illusion. We think the outputs of the causal detection mechanism are like this (or at least, lie towards this end of the spectrum).

defend the arguments, but simply to show what else one would need to buy, in addition to EDG, to reach that conclusion. We then outline a rather different way in which one might appeal to EDG in the service of the epistemology of grounding.

4.1 The explanatory dispensability argument

If true, EDG tells us that we can explain our grounding observations without appealing to a relation of ground. One way in which one might marshal EDG, then, is in conjunction with something like the explanatory criterion of ontological commitment. Then one might offer the following argument:

The Explanatory Dispensability Argument

1. One ought (epistemically) to be ontologically committed only to those entities that are indispensable to the best explanation of our observations.
2. EDG is true.
3. If EDG is true, then grounding relations are not indispensable to the best explanation of our observations.
4. Therefore, we should not be ontologically committed to grounding relations.

Many find (1)—the explanatory criterion of ontological commitment—attractive.[27] If one accepts (1) and (2), and, in addition, supposes that were there grounding, it would be a primitive relation, then (3) is true and so is (4). Matters are less straightforward if one thinks that, if there is a relation of ground, it reduces to something else. For then one might reject (3). One might concede that grounding relations are dispensable to our explanation of *grounding observations* ((2) is true). But perhaps the reductive base of grounding is indispensable to the explanation of some other observations. For instance, perhaps positing impossible worlds is indispensable to explaining how it is that our mental states have certain representational content. In that case, if grounding can be reduced to the truth of certain counterpossibles, (à la Alastair Wilson, forthcoming) then one might argue that (3) is false even though (2) is true. Of course, the reductionist about grounding would need to show that the reductive base of grounding is indeed indispensable to explaining some of our observations and that might prove difficult. So the argument from explanatory indispensability might go through even if one is a reductionist about grounding.

A second option for deploying EDG to reach a metaphysical conclusion lies in mounting what we will call a debunking grounding argument in the style of debunking arguments found in ethics. That argument proceeds as follows.

4.2 The debunking grounding argument

The Debunking Grounding Argument

1. EDG is true.

[27] Harman (1977: 6), Sayre-McCord (1988: 441), Colyvan (2000, 2001).

2. If EDG is true, then our grounding relevant judgements issue from an evolved cognitive mechanism, the causal detection mechanism.
3. The selective pressures that led to the evolution of the causal detection mechanism are such that, even if there were grounding relations, we would not expect that mechanism to successfully track those relations.
4. So if there are grounding relations, we should be sceptical that our grounding relevant judgements are any guide to the truths about grounding.

Debunking arguments attempt to show that, given the way a cognitive mechanism evolved, if there were certain phenomena in the world we would have good reason to be doubtful that our cognitive mechanism tracks those phenomena. So we should be sceptical that our judgements, issuing from those cognitive mechanisms, are any guide to the relevant truths. We think there is interesting work to be done in pursuing the debunking grounding argument; work that goes well beyond showing that EDG is true. For, as yet, nothing that we have said suggests that if there were grounding relations, the causal detection mechanism would have evolved in such a way that it would likely fail to track them.[28] Perhaps such a case can be made and we would be interested to see any such attempt.

4.3 An epistemology for the friend of grounding

Finally, one might reject both the explanatory indispensability argument and the grounding debunking argument. If one thinks that (1) of the explanatory indispensability argument is false, and if one thinks that the cognitive mechanism that we describe is well suited to track grounding relations, if there are any, then one could attempt to use our psychological story to explain how it is that we track grounding relations, and hence how we come to know what grounds what. For the friend of grounding surely needs *some* story about how, if there are grounding relations, we come to detect them, and it seems to us that what we say here has as much to recommend it as does any other story. We assume that the friend of grounding will say something like the following. The mechanisms we have pointed to have evolved to detect non-symmetric instances of relations from amongst the non-diachronic correlations. In doing so, sometimes these mechanisms detect modal relations, and sometimes they detect grounding relations. Indeed, sometimes when the mechanism detects some non-symmetry amongst the non-diachronic correlations, this is, *ipso facto,* to detect *both* a non-symmetric instance of a non-symmetric modal relation (say, necessitation) as well as a non-symmetric instance of an asymmetric grounding relation.

Having said all that, a good deal of work would need to be done in order to marshal the story we offer, here, as an epistemic story. In order for such a story to succeed, the friend of grounding needs not only to show that our mechanistic story does not debunk grounding (that given the way these mechanisms evolved, we

[28] Though, if Rodriguez-Pereyra (2015) is right that grounding is not asymmetric, then our account implies that the causation detection mechanism would likely do a bad job of identifying its symmetrical instances.

should expect them to accurately track grounding), she also needs to have some account of how we distinguish correct from incorrect grounding claims—that is, how we distinguish cases in when the mechanism in question issues in correct judgements, and cases in which it does not. None of this is straightforward, and getting clear on this story would represent a considerable research project for the friend of grounding.

5 Conclusion

For what it's worth, we are drawn to the explanatory dispensability argument, and so are inclined to think that our defence of EDG constitutes good reason not to posit grounding relations; but others may disagree and make different use of EDG. Whichever of these uses of EDG takes one's fancy, its defence should be of interest to the friend, and foe, of grounding alike.

Acknowledgements Funding was provided by Australian Research Council (Grant No. 110100486). With thanks to the following who provided feedback on an earlier draft of this paper: Antony Eagle, Mike Raven, Alastair Wilson, Dana Goswick, and Dan Marshall, as well as the audience of the Explananza 2015 for helpful discussion of these issues, and finally, to David Norton for his meticulous proof reading.

References

Arkes, H. R., & Harkness, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General, 112,* 117–135.

Atran, S. (2002). *In gods we trust: The evolutionary landscape of religion.* Oxford: Oxford University Press.

Audi, P. (2012). A clarification and defense of the notion of grounding. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding: Understanding the structure of reality* (pp. 101–121). New York: Cambridge University Press.

Barrett, J. L. (2004). *Why would anyone believe in God?.* Plymouth: Alta Mira Press.

Boyer, P. (2001). *Religion explained.* London: Vintage.

Chapman, L. J., & Chapman, J. P. (1967). Genesis of popular but erroneous psychodiagnostic observations. *Journal of Abnormal Psychology, 72,* 193–204.

Colyvan, M. (2000). Conceptual contingency and abstract existence. *Philosophical Quarterly, 50*(198), 87–91.

Colyvan, M. (2001). *The indispensability of mathematics.* New York: Oxford University Press.

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111,* 1–30.

Gould, S. J. (1991). Exaptation: A crucial tool for an evolutionary psychology. *Journal of Social Issues, 47*(3), 43–65.

Guthrie, S. (1993). *Faces in the clouds.* Oxford: Oxford University Press.

Hagmayer, Y., Sloman, S., Lagnado, D., & Waldman, M. (2007). Causal reasoning through intervention. In A. Gponik & L. Schultz (Eds.), *Causal learning: Psychology, philosophy and computation.* Oxford: Oxford University Press.

Harman, G. (1977). *The nature of morality.* New York: Oxford University Press.

Kushnir, T., Gopnik, A., Lucas, C., & Schulz, L. (2010). Inferring hidden causal structure. *Cognitive Science, 34*(1), 148–160.

Lagnado, D. A., & Sloman, S. A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 856–876.

Lowe, E. J. (2010). Ontological dependence. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2010 ed.). http://plato.stanford.edu/archives/spr2010/entries/dependence-ontological/.

Michotte. A. (1956) La Perception de la causalite/Louvain: Institut superieur de la Philosophier. *The perception of causality* (T. R. Miles, Trans.). London: Metheun.

Nisbett, R., & Ross, L. (1980). *Human inference; strategies and shortcomings of social judgement*. Englewood Cliffs, NJ: Prentice-Hall.

Plato. (2002). *Five dialogues: Euthyphro, apology, crito, meno, phaedo* (G. M. A. Grube, Trans) (2nd ed.). Revised by John M. Cooper. Indianapolis: Hacket Publishing.

Pyysiäinen, I. (2001). *How religion works: Towards a new cognitive science of religion. Cognition and culture book series, 1*. Leiden: Brill.

Pyysiäinen, I., & Anttonen, V. (2002). *Current approaches in the cognitive science of religion*. London: Continuum.

Raven, M. J. (2012). In defence of ground. *Australasian Journal of Philosophy, 90*(4), 687–701.

Redelmeier, D. A., & Tversky, A. (1996). On the belief that arthritis pain is related to the weather. *Proceedings of the National Academy of Science, 93*(7), 2895–2896.

Rigdon, M., Ishii, K., Watabe, M., & Kitayama, S. (2009). Minimal social cues in the dictator game. *Journal of Economic Psychology, 30*, 358–367.

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

Russell, B. (1903). *The principles of mathematics*. London: George Allen & Unwin.

Sayre-McCord, G. (1988) Moral theory and explanatory impotence. In *Midwest studies in philosophy* (Vol. 12, pp. 433–457) University of Minnesota Press.

Schafer, G. (1996). *The art of causal conjecture*. Cambridge: MIT Press.

Schaffer, J. (2009). On what grounds what. In D. Manley, D. Chalmers, & R. Wasserman (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 347–383). Oxford: Oxford University Press.

Shaklee, H., & Mims, M. (1982). Sources of error in judging event covariation: Effects of memory demands. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 208–224.

Shaklee, H., & Tucker, D. (1980). A rule analysis of judgements of covariation between events. *Memory and Cognition, 8*, 459–467.

Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. Oxford: Oxford University Press.

Smedlund, J. (1963). The concept of correlation in adults. *Scandinavian Journal of Psychology, 4*, 165–173.

Styvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science, 27*(3), 453–489.

Tallant, J. (2015). Ontological dependence in a spacetime-world. *Philosophical Studies, 172*(11), 3101–3118. doi:10.1007/s11098-015-0459-4.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The psychology of learning and motivation* (Vol. 34, pp. 47–77)., Causal learning San Diego: Academic Press.

Waldmann, M. R., & Hagmayer, Y. (2013). Causal reasoning. In D. Reisman (Ed.), *The Oxford handbook of cognitive psychology*. Oxford: Oxford University Press.

Williams, J. J., & Griffiths, T. L. (2013). Why are people so bad at detecting randomness? A statistical argument. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 39*(5), 1473–1490.

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

Wilson, A. (forthcoming). Metaphysical causation. *Nous*.