

When Are Structural Equation Models Apt? Causation Versus Grounding

Lina Jansson¹

University of Nottingham

December 1, 2016

1. Introduction

The notion of ground has made a prominent rise in contemporary metaphysics. While much about the notion of ground remains under debate, one feature has reached near consensus.² Ground is closely connected to a form of explanation. As Jenkins [2013] points out, we often use explanatory locutions such as “because” and “in virtue of” when discussing ground. The association to explanation is often explicit as in Dasgupta’s discussion of ground below.

Imagine you are at a conference, and imagine asking why a conference is occurring. A causal explanation might describe events during the preceding year that led up to the conference: someone thought that a meeting of minds would be valuable, sent invitations, etc. But a different explanation would say what goes on make the event count as a conference in the first place. Someone in search of this second explanation recognizes that conferences are not *sui generis*, so that there must be some underlying facts about event *in virtue of which* it counts as being a conference, rather than (say) a football match. Presumably it has something to do with how the participants are acting, for example that some are giving papers, others are commenting, and so on. An answer of this second kind is a statement of what *grounds* the fact that a conference is occurring. [Dasgupta, 2014: 3]

Even if we accept that the notion of ground is tied to the notion of explanation, there are substantive open questions about how exactly we should understand the connection to explanation.³

Recently, Schaffer [2016] and Wilson [2016, forthcoming] have argued that ground should be understood as an explanation-backing relation akin—or identical—to causation. Just as causal relationships can back a particular type of explanation, causal explanation, so grounding relationships can back a particular type of explanation, grounding explanation. Part of their argument is based on the claim that both notions can be productively treated using the tools of structural equations and directed graphs. The use of directed graphs and structural equations in an account of explanation is most closely associated with Woodward’s [2003] interventionist counterfactual account of causal explanation. If we can use a similar framework in the grounding case, then we could unify the structure of grounding explanations and that of causal explanations, thereby adding weight to the idea that ground should be

¹ Thank you to the Nottingham Metaphysics and Epistemology Reading Group, Juha Saatsi, Alex Reutlinger, and to an anonymous referee for comments on an earlier version of this paper. Thank you also to Jonathan Tallant and Al Wilson.

² See Wilson [2014: 555-556] for a dissenting view.

³ See Bliss and Trogon [2014] for a delineation of several of these issues.

treated as an explanation-backing relation akin to causation. In section 2 I will briefly review how structural equation modelling works in the causal case before showing how the formal framework can be extended to the grounding case.

In section 3 I will argue that it is only the formal framework that carries over to the grounding case. In particular, the *seeming* unification of the structural equations approach to explanation disappears once we take into account what it takes for a structural equations model to have appropriately captured the situation that we are modelling. As Schaffer [2016] and Blanchard and Schaffer [forthcoming] emphasise, structural equation modelling is a type of *modelling*. Once we are given a model of some scenario or system, the obvious question to ask is whether the model is any good. That is, is the model an apt or fitting one (for the purpose at hand)? For a model to be a good one for the purposes at hand it has to contain appropriate (whatever that turns out to mean) variables and appropriately (whatever that turns out to mean) represent the relations of causal or grounding relevance.

I will show that there are important differences between the causal case and the grounding case in how we can judge whether or not a model is an appropriate representation of a scenario. Once we take into consideration the aptness conditions associated with causal explanations we find that we rely on non-trivial a posteriori knowledge about the nature of the specific causal processes relevant to the model. This allows us to rule out some models as failing to adequately represent the causal dependences of interest. In the grounding case we do not have access to such specific a posteriori knowledge of the grounding processes relevant to the grounding models. This means that we cannot draw on the same resources to rule out certain models as failing to appropriately capture the grounding dependences of interest.

Before I set out my argument, let me address a natural initial objection. The question of whether a model is an apt one is a question with epistemic and pragmatic overtones. That such concerns should be part of the modelling criteria for ground and grounding explanation at all might seem surprising given that ground is supposed to be a metaphysical notion. It is tempting to try to avoid epistemic concerns by simply stipulating that an apt model is one that correctly captures the grounding structure.⁴ However, there is a high cost to going down this route.

The connection to explanation is often viewed as the strongest reason to accept the existence of a grounding relation that is conceived as some kind of determination relation along the lines of causation: whether this determination relation is a productive relation, a generative relation, or a dependence relation.⁵ Audi [2012: 105] formulates this line of argument particularly clearly.

(1) If one fact explains another, then the one plays some role in determining the other.

(2) There are explanations in which the explaining fact plays no causal role with respect to the explained fact.

⁴ I take Wilson (2016) to choose this option. Schaffer (2016: 68-69) discusses the worry of the lack of a uniquely appropriate model but does not offer a solution.

⁵ See, for example, Schaffer [2016: 58] and, an early example, Kim [1974]. Kovacs [2016] calls the argument from explanation to ground the “master argument” for grounding.

(3) Therefore, there is a non-causal relation of determination.

This argument runs from the existence of what is supposed to be clear cases of metaphysical, non-causal, explanation to the postulation of a relation of ground. In order for such an argument to work, we have to recognise the explanations in question as explanations. So, whatever the explanation-backing relation (if any) is, it should be a relation that we can recognise as obtaining. If it is not a relation that we have epistemic access to—at least in the sense of having good reason to think that the relation obtains—then we must either take ourselves to not have good reason to think that we have an explanation in the cases under consideration, or take it that the relation obtaining is not crucial for having those explanations after all. Either option would undermine the argument from the recognition of these explanations to the postulation of a relation of ground that is backing the explanations. Given these background assumptions, if we expect ground to be a relation that we can capture using structural equation models, then we should also expect to be able to use these models to recognise certain specific explanations. The models cannot simply be apt, but must be such that we can recognise that they are apt.⁶ This means that epistemic concerns about how we assess the aptness of grounding models cannot be simply set aside.

2. The formalism of structural equations models

In this section I will introduce the formal machinery associated with structural equations models. I will first present the causal case before showing how Schaffer and Wilson modify the notions associated with the causal case in order to apply a structural equations framework to the grounding case.

2.1 Structural equations modelling of causal scenarios

In Woodward's [2003] framework, causal relationships are represented using directed graphs, where a directed edge from one variable to another represents a direct causal relationship. In addition to the directed graph we also specify a collection of structural equations that specify exactly how the values of the dependent variables are functions of their direct causes.

I make use of two devices to represent causal relationships. A *directed graph* is an ordered pair $\langle \mathbf{V}, \mathbf{E} \rangle$ where \mathbf{V} is a set of vertices that serve as the variables representing the relata of the causal relation and \mathbf{E} a set of directed edges connecting these vertices. A directed edge from vertex or variable X to

⁶ This already raises a worry about the argumentative strategy at hand. In so far as we are concerned with grounding as a pure metaphysical relation—unadulterated by epistemic and pragmatic constraints—we have reasons to be concerned about whether structural equation models and the associated account of explanation are suited to provide this. We should be unsurprised to find that that relations modelled with an eye to explanations available to us are a mixture of ontic and epistemic considerations (See Illari (2013) for an argument for this in relation to mechanistic explanations.) Woodward's [2003: 179] interventionist counterfactual approach to explanation incorporates strong epistemic constraints. These constraints are reflected in a relative disinterest in the fundamental ontology. Importantly, the causal relations are not even taken to be good candidates for fundamental relations. A variable is taken to be a direct (or contributing) cause for another variable only relative to a variable set.

vertex or variable Y means that X *directly causes* Y ... The basic idea is that X is a direct cause of Y if and only if the influence of X on Y is not mediated by any other variables in the system of interest \mathbf{V} in the following sense: there is a possible manipulation of X that would change the value of Y ... when all other variables in \mathbf{V} are held fixed at some set of values in a way that is independent of the change in X ...

It will also be useful to represent causal relationships by means of systems of equations ... in which each endogenous variable Y (i.e., each variable that represents an effect) may be written as a function of all and only those variables that are its direct causes ...

[Woodward, 2003: 42]

Let me use a very simple example to illustrate the framework. Let us say that I want to model a scenario where a boulder falls and starts rolling towards a hiker. The hiker sees the falling boulder and ducks.⁷ To model the scenario we introduce a variable for whether or not the boulder falls, let us call it F , and a variable for whether or not the hiker ducks, let us call it D . For each variable we stipulate that it takes one of two values; 1 if the boulder falls and if the hiker ducks (respectively), and 0 otherwise. In order to represent the relations of causal relevance, we can make use of directed edges. The fall (or non-fall) is directly causally relevant to the ducking or non-ducking of the hiker. A causal diagram captures this simple causal structure.



Figure 1

We also need to represent the endogenous variables—in this case only D —as a function of their direct causes. This is easy to do. The hiker ducks if the boulder falls, but not otherwise. So, D takes the value 1 when F takes the value 1 and the value 0 when F takes the value 0; $D=F$.

Woodward's account is a non-reductive one. In the full account, the notion of a 'possible manipulation of X that would change the value of Y ... when all other variables in \mathbf{V} are held fixed at some set of values in a way that is independent of the change in X ...' relies on a technical notion of an intervention on X with respect to Y . The notion of an intervention is itself defined in causal terms relative to a causal graph.⁸

2.2 Structural equations modelling of grounding scenarios

The cases discussed as examples of grounding explanations are of heterogeneous

⁷ We will expand on this case in the next section.

⁸ See Woodward [2003: 98].

nature but often include examples like the following:⁹

- The fact that P and the fact that Q jointly ground the fact that $P \wedge Q$.
- The fact that $P \wedge Q$ does not ground the fact that P.
- The fact that Socrates exists grounds the fact that the singleton set containing Socrates as its sole member exists.
- The fact that the singleton set containing Socrates as its sole member exists does not ground the fact that Socrates exists.
- The fact that an act is virtuous grounds the fact that the gods love it.
- The fact that the gods love an act does not ground the fact that the act is virtuous.

All of the above seem like cases where there is a *directed* notion of dependence in place. Just as in the causal literature, it seems hard to analyse this directed notion of dependence purely in counterfactual terms. Some counterfactuals seem relevant, but others do not.¹⁰ As Wilson [2016] points out, we need to rule out the counterfactuals that seem to track back along the direction of causation and back along the direction of ground. It seems true that had Socrates not existed, then the singleton containing Socrates would also not have existed: this seems like the right counterfactual for the grounding case. It also seems true that had the singleton containing Socrates not existed, then Socrates would not have existed; this seems like the wrong counterfactual for the grounding case. The hope is that the grounding case could—like the causal case—make use of a non-reductive but informative account to clarify which counterfactuals are relevant.

In particular, notice that in the causal case we start by selecting an appropriate causal model where we specify the endogenous and exogenous variables. The formal notion of intervention is defined relative to such a model. It is, therefore, tempting to take the same strategy in trying to capture the relevant counterfactuals in the grounding scenario, and to capture the interventionist counterfactuals *in terms of* directed graphs and structural equations. This is the strategy of Schaffer [2016] and Wilson [2016]. Wilson [2016: 3] characterises an intervention by reference to a directed graph as “... a ‘clean’ alteration of the value of a particular variable that does not affect the values of upstream causal variables”.

To see the framework in action, let us return to the case of the fact that P together with the fact that Q grounding the fact that $P \wedge Q$.¹¹ We introduce a variable for

⁹ I should note that there is no consensus on what the proper relata of ground are. Here I am following the convention making the relata facts, but nothing that I will go on to say hinges on this choice. I will drop the reference to facts when it makes sentences less cumbersome to do so.

¹⁰ In all of the cases above, the antecedent would be a counterfactual and not a counterpossible. However, counterpossible antecedents will be needed for some examples in the grounding literature. For example, it may be taken to be the case that the fact that 2 exists grounds the fact that {2} exists, but not vice versa.

¹¹ This case is adapted, with some conceptual differences, from Schaffer [2016: 79].

whether or not P obtains; let us call it P . Similarly, we introduce a variable for whether or not Q obtains; let us call it Q . Finally we introduce a variable for whether or not $P \wedge Q$ obtains; let us call it C . As in the previous case, each variable takes one of $\{0, 1\}$ depending on whether or not the fact in question obtains.

In order to represent the relations of grounding relevance, we can make use of directed edges. Whether or not P obtains is directly grounding relevant to whether or not $P \wedge Q$ obtains. Similarly, whether or not Q obtains is directly grounding relevant to whether or not $P \wedge Q$ obtains. Just as in the causal case, a directed graph does not tell us how the endogenous variables depend on their direct grounds (or causes). To specify this we make use of structural equations. We only need one structural equation in order to specify how $P \wedge Q$ depends on P and Q ; $C = \text{Min}(P, Q)$. That is, the value of C is equal to the lowest value taken by either of P or Q (or both).




Figure 2

According to the grounding model, some interventions on P , with Q fixed, entail a “downstream” alteration of C . In particular, if Q takes the value 1, an intervention that changes the value of P from 1 to 0 will have the downstream consequence of changing C from 1 to 0. However, P is not downstream from C , so no conclusions about changes in P based on interventions on C are supported by the grounding model.

So far this is just a formal characterization. We can still ask what the intervention in question represents. This brings with it the worry that interventions on variables in the grounding case are not “clean”. In particular, the variables involved (such as whether $P \wedge Q$ holds and whether P holds) are not *independently fixable* (IF). That is, they will clearly violate the criterion specified in Woodward [2015]).

(IF): a set of variables \mathbf{V} satisfies independent fixability of values if and only if for each value it is possible for a variable to take individually, it is possible (that is, “possible” in terms of their assumed definitional, logical, mathematical, mereological or supervenience relations) to set the variable to that value via an intervention, concurrently with each of the other variables in \mathbf{V} also being set to any of its individually possible values by independent interventions. [Woodward, 2015: 316]

Woodward does not endorse IF as a general requirement on interventionist accounts. Rather, he takes it to indicate a difference between causal and non-causal dependence (with only the former taken to satisfy IF).¹² Wilson [2016] is clear that we will have to consider not only counterfactual scenarios, but sometimes also counterpossible ones, when we move to the grounding case. We may for example need to intervene

¹² Woodward’s (2015) interest is in models with mixed causal and non-causal dependences. Here he suggests respecting definitional constraints so that we should not “keep fixed” values of variables related in a definitional way to a variable under intervention.

on the existence of the singleton set without intervening on the existence of its member or intervene to change the value of $P \vee Q$ from true to false without changing the value of P from true to false. This will clearly violate IF. Wilson [2016] diagnoses the need for non-trivial counterpossibles as one of the reasons for grounding skepticism.¹³

While I do not think that the difficulties in understanding counterpossibles are trivial, I want to focus on a different problem in this paper.¹⁴ Namely, even if we allow that there are non-trivial counterpossibles, there are stark differences in how we can understand the *aptness* of causal versus grounding models. That is, there are important differences in how we can judge whether or not a model is any good as a representation of a scenario. I turn to this question in the next section.

3. Aptness of structural equations models

Once we have a causal model or a grounding model on the table the obvious question to ask is whether the model is an adequate representation of the causal or grounding scenario that we are modeling. For short, is the model an apt one? In this section I move on to consider what we can say about how we go about judging whether a model is an apt one. I will first show how the question of aptness arises and discuss how it is addressed in the context of causal structural equations models. In section 3.2 I will argue that the question of models' aptness is of similar importance in the grounding case, but that we do not have the same resources to address the question.

3.1 Aptness of causal structural equations models

Within the literature on causal modelling it is acknowledged that it is of crucial importance to select appropriate variables.¹⁵ We have already seen one suggested restriction on variable choice for causal models; when representing relations of causal dependence, the variables used should be independently fixable. Here, I want to focus on another aptness restriction: the appropriate grain of representation of the model. A model can fail to be apt by including too few variables or by including too many.

Let me extend the example of the falling boulder from section 2 to illustrate the idea that it is possible to fail to construct an apt model by including *too many* variables on a path. In the new scenario, we add the information that the hiker survives the encounter with the falling boulder. The example is described by Hitchcock [2001: 276] and attributed to an early draft of Hall [2004].

“Boulder”: a boulder is dislodged, and begins rolling ominously toward Hiker. Before it reaches him, Hiker sees the boulder and ducks. The boulder sails harmlessly over his head with nary a centimeter to spare. Hiker survives his

¹³Schaffer (2016: 72) also takes this to open the door to grounding scepticism (although he does not endorse it).

¹⁴I do not want to suggest that these are the only difficulties with extending the interventionist framework to grounding cases. See Koslicki (2016) for some very different concerns from the ones that I will raise here.

¹⁵For an extended discussion about variable choice see, for example, Hitchcock [2012] and Woodward [2016].

ordeal.

We can model the scenario by following Woodward's [2003: 79-80] discussion of the case.

As before, we introduce a variable for whether or not the boulder falls, let us call it F , and a variable for whether or not the hiker ducks, let us call it D . The variables F and D behave as in section 2.1. We also introduce a variable for whether or not the hiker survives, let us call it S . S takes the value 1 if the hiker survives and 0 otherwise. In order to represent the relations of causal relevance we again make use of directed edges. Here, the fall (or non-fall) of the boulder is directly causally relevant to the survival or non-survival of the hiker. The fall (or non-fall) is also directly causally relevant to the ducking or non-ducking of the hiker. Finally, the ducking (or non-ducking) of the hiker is directly causally relevant to the survival or non-survival of the hiker. A redrawing of Woodward's [2003: 79] figure 2.7.2 captures this.




Figure 3

We also need to represent the endogenous variables (S , D) as a function of their direct causes. As before $D=F$. We also need to express S as a function of its direct causes. This is easy to do. The hiker survives if either the boulder does not fall or the hiker ducks. So, the value of S is the highest of the values taken by either D or $1-F$ (or both); $S = \text{Max}(D, 1 - F)$.¹⁶

Let us now return to the question of this section. Is the model an adequate representation of the scenario in question? In particular, have we included the right number of variables? Let us return to our intuitive judgements about the causal relations in the scenario. This case is often used to illustrate that at least some causal notions are not transitive. The boulder falling is a cause of the ducking, and the ducking is a cause of the survival of the hiker. Yet, the boulder falling is not a cause of the survival. The notion of cause that is relevant to our judgement is the notion of an actual cause (or token causation). Woodward's [2003: 79–81] discussion of the case reveals that the result that our model delivers about whether or not the fall of the boulder is an actual cause of the survival of the hiker hinges crucially on whether or not we judge it to be appropriate to include more variables along the F - S path in figure 3.

Formally, for a case that is not one of symmetric overdetermination, we have the following criterion for whether $X=x$ is an actual cause of $Y=y$.¹⁷

(AC1) The actual value of $X = x$ and the actual value of $Y = y$.

(AC2) There is at least one route R from X to Y for which an intervention on X will change the value of Y , given that other direct causes Z_i of Y that are not on this route have been fixed at their actual values. (It is assumed that all

¹⁶ These are slight notational variants on Woodward's [2003: 79] equations 2.7.4 and 2.7.5.

¹⁷ We have to be more careful in cases of overdetermination. See footnote 18.

direct causes of Y that are not on any route from X to Y remain at their actual values under the intervention on X .)

Then $X = x$ is an actual cause of $Y = y$ if and only if both conditions (AC1) and (AC2) are satisfied. [Woodward, 2003: 77]

To assess whether $D=1$ is an actual cause of $S=1$, we evaluate the causal influence along each path from the ducking to the survival, keeping the off-path variables representing direct causes of S fixed to their actual values. If we keep the value of F (the falling of the boulder) fixed to its actual value (the only variable not on the path from D to S that is a direct cause of S) and we intervene to change whether the hiker ducks or not (the value of D), then we change whether or not the hiker survives (the value of S). So the ducking of the hiker is an actual cause of the hiker's survival. This is as we suspected.

However, the falling of the boulder is not an actual cause of the survival of the hiker. Following Woodward's discussion, we have two paths to consider. Let us look at the direct F - S path first. If we keep the ducking (D) fixed at its actual value, then changing whether or not the boulder falls (the value of F) does not alter whether or not the hiker survives (the value of S). We also have a second path to consider: the path from F to S that goes via D . Here there are no variables to keep fixed along the direct F to S path, so we cannot evaluate the influence of F on S via the F - D - S path by keeping any such variables fixed. In neither case do we find that the falling of the boulder makes a difference to the survival of the hiker. So far, things look good. Our use of structural equation modeling has recovered the judgment that the notion of actual cause fails to be transitive in the way that our intuitive judgement about the case leads us to expect.

Notice, however, that our solution depends on there not being a variable included on the F to S path representing a point on the trajectory where the boulder is too close to the hiker's head for the hiker to duck (as Woodward points out in his discussion of the case).

[T]his treatment of the boulder example depends crucially on the absence of any intermediate variable on the direct route from F to S . This raises the obvious question of why it wouldn't be equally or more correct to include such a variable in our representation of the example. [Woodward, 2003: 80]

If we introduce such a variable, then we need to keep this variable fixed to evaluate the influence along the F - D - S route. We are now considering a scenario where the boulder fails to fall but appears at an intermediate point in its trajectory where it is too late for the hiker to duck. If there was such a variable, then when evaluating the influence of F on S via the F - D - S path, we find that had the boulder not fallen the hiker would *not* have survived. What makes it inappropriate to include such a variable?

In our discussion of the falling boulder example..., we rejected the idea that it was appropriate, given the causal structure of this example, to consider the possibility that the boulder both failed to fall and yet (somehow) appeared a few meters from the hiker's head. It was not that this was (in itself) a logical or causal or nomological impossibility, but rather that, to take this possibility seriously, we needed to consider an example with a rather different causal

structure from the one we originally set out to analyze, one in which some independent mechanism or process, other than falling, is responsible for the appearance of the boulder in close proximity to the hiker. At least in ordinary contexts, the possibility that the boulder both fails to fall *and* appears near the hiker's head *and* doesn't get there as a result of following a trajectory from some independent source but instead, say, simply materializes near the hiker's head is not one that we are prepared to take seriously.
[Woodward, 2003: 86]

The judgement that there is no serious possibility of the boulder failing to fall and yet simply materialising near the hiker's head is one that we have good evidence for. Although it is not ruled out by the metaphysical nature of causation or by the very concept of causation that a boulder could fail to fall and simply materialise on a trajectory close to the hiker's head, it is not how we take boulders to behave. The possibility considered is incompatible with our best theoretical understanding of the causal behaviour of objects like boulders. This is not to say that we take it to be incompatible with the *concept* of causation to have a boulder behave such as to simply appear close to the hiker's head. Rather our *a posteriori theory* of what the causal mechanisms and processes responsible for boulder behaviour are like rules out such a possibility. Like all of our theories about the physical world, such a theory is fallible and not typically a conceptual truth. For example, we may have a *theory* that restricts causal influences to propagate at subluminal speeds without taking it to be any part of the *concept* of causation that causal influences are so restricted.

The theory of the causal mechanisms involved is a *local* one. We can have a theory of the causal mechanisms behind boulder behaviour that rules out a boulder simply appearing close to the hiker's head without getting there by travelling in a continuous trajectory. Yet, we can also allow that an apt model for the behaviour of subatomic particles should not include a restriction where a particular kind of subatomic particle can only be present within a given area through the earlier presence of a particle of that kind at some nearby area. There is no conflict between these two judgements of aptness. They both stem from an understanding of what the causal mechanisms in question are—typically, and around here—like.

The considerations that we invoke in order to conclude that a model that includes an extra variable on the *F* to *S* path is not a good representation of the causal situation are knowable only a posteriori, they apply only to a particular type of system, and they do not amount to a conceptual ban on the possibilities that are ruled out. In the next section I will argue that this type of consideration is not available in the grounding case.

Before moving on to the next section I want to push a bit deeper into the reasons that Woodward gives for ruling out the possibility of the boulder failing to fall but simply materialising close to the head of the hiker. Woodward takes this particular case to be one where the possibility is ruled out based on objective facts about how the world operates.

As I have already intimated, I think that it is true that in some cases an investigator's ... interests and purposes (and not just how the world is) influence the possibilities that are taken seriously... On the other hand, as the examples described above illustrate, at least some of the considerations that go

into such decisions are based on facts about how the world operates that seem perfectly objective. For example, there is nothing arbitrary or subjective about the claim that boulders don't materialize out of thin air... [Woodward, 2003: 89]

The causal notions in the interventionist framework are not, however, generally purely a matter of just what the world is like.

... [C]ausal judgements reflect both objective patterns of counterfactual dependence and which possibilities are taken seriously; they convey or summarize information about patterns of counterfactual dependence among those possibilities we are willing to take seriously. In other words, to the extent that subjectivity or interest relativity enters into causal judgments, it enters because it influences our judgements about which possibilities are to be taken seriously. [Woodward, 2003: 90]

To make the non-actual (but causally conceptually possible) scenario of the boulder materialising out of thin air ruled out on grounds of how the world operates, we need to appeal to some feature of the world that constrains not just what is *actually* true about the world; after all, we are ruling out a non-actual possibility. We need to turn to some feature of the world with modal force to achieve this. To introduce the notions of causal mechanism or causal production (or, for that matter, worldly dependence) is attractive in the boulder case since those notions are typically taken to be at least candidates for objective, interest independent, features of the world that do have modal force. A purely worldly notion of cause is also not, however, the notion of cause that is generally represented by causal graphs and structural equations in the interventionist account.

So far, I have tried to convince you that adjudicating whether or not a causal model includes the right variables (both in number and kind) to represent the causal structure of interest can involve considerations that, first, take us outside the realm of a purely conceptual account of causation and to a theory of features of particular causal processes or mechanisms. Crucially, we can have a posteriori evidence for or against a suggested theory of a particular causal process or a causal mechanism. Second, the theories of causal mechanisms or processes can be local. They do not have to be taken to apply to all types of causal processes. Third, in order to take the constraint on aptness to stem from objective consideration from way the world operates we need to introduce a notion of cause (or other notion with modal force) that is not the one that is the immediate target of analysis in interventionist causal models.

Let us see if we can generalise this discussion to help us answer the question of aptness in the case of grounding models.

3.2 Aptness of grounding structural equations models

The question of aptness is as important in the grounding case as in the causal case. In particular, the selection of which variables to include affects the conclusions about ground that can be drawn from the models. As the example of the boulder from the previous section illustrates, whether a model in the causal case delivers the intuitively correct verdict about actual causes depends on not including too many variables on a path. In this section I will show that including too many variables in a grounding

model similarly leads to intuitively mistaken verdicts about the grounding relations in the scenario that is being modelled.

In the causal case this challenge was addressed by appealing to the aptness (or lack of aptness) of the model. At the end of this section I will argue that the kinds of considerations available in the causal case are not available in the grounding case.

Let me now turn to the example that will show that including too many variables is as problematic in the grounding case as in the causal case. Consider the claim below.

Disjunction from Conjunction (DfC): $P \wedge Q$ and R are both potential grounds for $Q \vee R$.

It is uncontroversial among grounding theorists that when R holds, R grounds $Q \vee R$ (it is similarly uncontroversial that Q can ground $Q \vee R$). It is less clear whether $P \wedge Q$ can ground $Q \vee R$. In favour of taking $P \wedge Q$ to be directly grounding relevant, we can note that it is easy to imagine scenarios where it being a fact that $P \wedge Q$ looks like a potential explanation of it being a fact that $Q \vee R$. For example, imagine a scenario where we start from information about the conjunctive fact that $P \wedge Q$; a tempting answer to give as to why $Q \vee R$ holds is that $P \wedge Q$ does. In particular, the extended interventionist motivation for attributing a relation of direct grounding relevance seems to hold: there are some interventions on the variable representing $P \wedge Q$ such that it is associated with changes in the variable representing $Q \vee R$ (in particular intervening in order to make the case that $P \wedge Q$ holds will make it the case that $Q \vee R$ holds, and some interventions to make it the case that $P \wedge Q$ does not hold will make it the case that $Q \vee R$ does not hold). Moreover, we cannot break the grounding structure into two steps; we cannot conceive of the scenario as one where $P \wedge Q$ grounds Q which grounds $Q \vee R$. Any standard account of ground would deny that $P \wedge Q$ is grounding relevant to Q . If we want to capture a grounding scenario where our “initial” conditions (fundamental conditions relative to the scenario modelled) are conjunctive facts then we have to either fail to capture any grounding relevance between $P \wedge Q$ and $Q \vee R$ or include it directly.

Whichever way we go on the truth or falsity of *DfC*, the grounding model below should strike us as problematic.




Figure 4

Here C is a variable that takes the value 1 if $P \wedge Q$ holds and 0 otherwise. D is a variable that takes the value 1 if $Q \vee R$ holds and 0 otherwise. P is a variable that takes the value 1 if P holds and 0 otherwise (and mutatis mutandis for Q and R). We can use the following structural equations to try to represent these relationships. The value of D is the highest value taken by any of R, C, Q ; $D = \text{Max}(R, C, Q)$. The value

of C is the lowest value taken by either P or Q (or both); $C = \text{Min}(P, Q)$

Let it be the case that P , Q , $P \wedge Q$ and $Q \vee R$ all obtain. Let it also be the case that R does not. In this scenario, Q should turn out to ground $Q \vee R$ but R should not. We can extend the terminology of the last section to capture this. In particular, $Q=1$ should turn out to be an actual ground of $D=1$ but $R=0$ should not turn out to be an actual ground of $D=1$. We have a few options available to us for how we go about evaluating the relations of actual ground in this case (corresponding to various alternatives for how to evaluate actual causation). Let us try a fairly straightforward and standard proposal by generalizing Woodward's criterion for cases that do not involve symmetric overdetermination. This is appropriate since, intuitively, in the scenario considered $Q \vee R$ is not actually overdetermined; Q holds but R does not. We will try to isolate the grounding influence along a single path. To do so we will keep any off-path variables representing direct grounds of D fixed at their actual values. If we follow this reasoning, then Q is not an actual ground of $Q \vee R$. Why? There are two paths from Q to D to consider: the direct path and the path via C . Let us consider the direct Q to D path first. There are two variables to keep fixed at their actual values: C and R . The actual value of C was stipulated to be 1. So changing the value of Q from 1 to 0 will not change the value of D . Let us consider the Q to C to D path. Now, we cannot fix the values of intermediate variables along the direct Q to D path (since there are no such variables). So we cannot evaluate the influence along the Q to C to D path. In analogy to the boulder case, $Q=1$ turns out to not be an actual ground of $D=1$. Here, however, this is bad news! Q should have turned out to be an actual ground for $Q \vee R$ in this scenario. The model has clearly gone wrong somewhere. The crucial question for grounding theorists is to identify where it has gone wrong.¹⁸

Intuitively it is clear what has happened. While it is at least not clear that we would want to deny that the fact that $P \wedge Q$ could be grounding relevant to the fact that $Q \vee R$, it is a mistake to represent the fact that $P \wedge Q$ as an intermediate step between Q and $Q \vee R$ in the grounding structure in question. Just as in the discussion of including an intermediate variable on the F to S path in the boulder case, the mistake lies in including an intermediate variable where there ought to be none. In the boulder case we diagnosed the mistake by appealing to considerations from what we take the causal processes and causal mechanisms of boulder motion to be like. Can we do something similar in the grounding case?¹⁹ I think that the answer is that we cannot.

¹⁸ In making this evaluation I have relied on Woodward's [2003: 77] criteria for actual causation. One possible objection is that Woodward already notes that these criteria will have to be amended in order to handle cases of symmetric overdetermination well. However, this will not help the situation. On more involved definitions of actual causation adapted for the grounding case we can get Q to be an actual ground for $Q \vee R$. However, we do so at the cost of allowing $P \wedge Q$ to count as a *separate* actual ground for $Q \vee R$. That is, we make the situation look like a case of actual overdetermination. This is not an improvement. When Q holds and R does not hold, $Q \vee R$ is not actually grounding overdetermined. See Weslake [forthcoming] for a summary of different proposals and for a new suggestion.

¹⁹ The criteria for aptness of grounding models that Schaffer [2016: 74-75] proposes (by adapting to the grounding case the criteria suggested by Blanchard and Schaffer [forthcoming] for the causal case) do not solve the problem. The problem is not that we have *left out* some fact that ought to have been included (we do not have too few variables).

In the causal case we relied on a posteriori evidence for the claim that the causal mechanisms and processes involved in boulder motion operate to produce contiguous motion. In our experience, boulders (and boulder like objects) do not just materialise out of thin air. Importantly, this is merely a theory about the causal behaviour of boulders. It does not rise to the level of a conceptual claim about causation.

We have several reasons to worry about whether something similar could apply to the grounding case. First, can we make use of questions about the relevant grounding mechanisms and processes when evaluating the aptness of a grounding model? I take it to be much easier to see how we have explanatory dependence in the examples of grounding in section 2 than to see how we have a relation of grounding mechanism or some specific grounding process.²⁰ This would already seem to indicate a difference from the causal case; to get something close to the causal notion of process we would have to take notions of metaphysical building or metaphysical making non-metaphorically and to not be mere alternative ways of talking about metaphysical explanation and dependence.

Second, in the causal case we are relying on local theories of the causal behaviour of a specific type of system. We do not need to judge it to be conceptually or metaphysically impossible that boulders would materialise out of thin air in order to rule out a causal model as failing to be apt if it reflects such a possibility. In the grounding case this middle ground option—between including all possible scenarios and only including only one—is not available. What we would need is some information about the way that the actual world operates that would make it inappropriate to consider the possibility represented by including the variable representing the value of $P \wedge Q$, but that would not rule out the possibility that in other possible (or impossible) worlds the inclusion of such a variable would be appropriate. The relations that we are interested in when it comes to grounding cases do not, in general, seem to allow for this. We do not have a theory of disjunction in the actual world that—with good reason—takes disjunction to have features at the actual world that are different from those of disjunction in other possible worlds (although we can, of course consider different types of disjunction). A local theory of disjunction is not forthcoming.

In the causal case, it is the appeal to such local considerations that allow us to put objective constraints on which models are apt. Without them we are left to select the possibilities that are to be taken seriously merely based on criteria such as the interest of the modeller and other non-objective factors. In the grounding case we need to find a substitute for local theories of causal mechanisms or processes that could play the role of selecting the relevant possibilities (or impossibilities) that ought to be captured by the grounding model.

The most obvious solution to the problem—take all possibilities seriously in the grounding case—will not work. After all, we have already seen that the grounding models in question rely on taking some *counterpossibles* seriously. We cannot restrict ourselves to merely taking all possibilities (not even all logical possibilities)

²⁰ This is in line with Wilson (forthcoming). However, Schaffer [2016: 54] takes ground to have most of the features that we may associate with causal relations.

seriously. However, merely including impossibilities does not help. If we extend the proposal to include counterpossibles—take all possibilities and impossibilities seriously in the grounding case—we destroy the ability to rule out the graph in figure 4 as failing to be apt. There is no scenario (possible or impossible) that the graph in figure 4 could mistakenly have represented as a scenario that ought to be taken seriously. They should all be taken seriously on the view under consideration.

So far I have hoped to convince you that when it comes to considerations of aptness the analogy between causal models and grounding models break down. In the next section I will briefly suggest a way forward for grounding models.

4. Conclusion

Grounding is often taken to have very close ties to metaphysical explanation. The existence of a type of non-causal explanation has been cited as the driving reason to postulate a relation of ground. Here, I have simply granted that there are such explanations and focused on the question of whether we have reason to take ground to be akin to causation in its role in explanation. I think that the answer to this question is no.

In the previous section I argued that we do not have an obvious way of extending to the grounding case the resources that we have for adjudicating whether or not a causal model is apt in the causal case. The difficulty in the grounding case of finding reasons for ruling out a possibility (or counterpossibility) as not appropriate also suggests a way forward. We cannot simply include all possibilities or all possibilities and counterpossibilities. Nor can we appeal to a posteriori evidence for local theories of grounding processes. We can, however, appeal to general logical and metaphysical theoretical principles to try to provide an account of which possibilities or counterpossibilities that are relevant.

The focus on exceptionless privileged regularities puts the account more in line with the deductive-nomological account than with causal accounts of explanation.²¹ Part of what is distinctive about causal explanation and causal modelling is that we can use it even when we lack information about exceptionless laws. In the causal case structural equations do not need to be exceptionless, and our judgements about the aptness of models can rely on empirical local theories of particular causal processes. In the previous section I argued that the same does not hold for the case of grounding models.

If the explanations in the grounding case are better understood by analogy to explanations that invoke general laws or general principles than by analogy to causal explanations, then the questions that we raise about ground look rather different. It is now a pressing matter to understand these general principles. They are what we need to articulate in order to substantiate a grounding claim.

Although my argument here has been different from Wilson's [2014: 544-545], I take my argument to support the claim that bare ". . . Grounding . . . claims leave open questions that must be answered to gain even basic illumination about or allow even

²¹ For an explicit articulation of ground by analogy to law-based explanation see Wilsch [2015, 2016] and Jansson [2016].

basic assessment of claims of metaphysical dependence...”. In particular, we cannot avoid answering questions about the aptness of grounding models. Moreover, the answers look like they will—unlike in the causal case—have to come from general principles.

References

- Audi, P. (2012), ‘A Clarification and Defense of the Notion of Grounding’, in F. Correia and B. Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (Cambridge University Press), 101-121
- Blanchard, T., and Schaffer, J. (forthcoming), ‘Cause without Default’, in H. Beebe, C. Hitchcock, and H. Price (eds.), *Making a Difference: Essays in Honour of Peter Menzies* (Oxford University Press)
- Bliss, R., and Trogon, K. (2014), ‘Metaphysical Grounding’, in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition), <http://plato.stanford.edu/archives/win2014/entries/grounding/>
- Dasgupta, S. (2014), ‘On the Plurality of Grounds’, *Philosophers’ Imprint*, 14/20:1-28
- Fine, K. (2012), ‘Guide to Ground’, in F. Correia and B. Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (Cambridge University Press), 37-80
- Hall, N. (2004), ‘Two Concepts of Causation’, in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals* (MIT Press), 225-276
- Hitchcock, C. (2001), ‘The Intransitivity of Causation Revealed in Equations and Graphs’, *The Journal of Philosophy*, 98/6:273–299
- Hitchcock, C. (2012), ‘Events and times: a case study in means-ends metaphysics’, *Philosophical Studies*, 160/1:79–96
- Illari, P. (2013), ‘Mechanistic Explanation: Integrating the Ontic and Epistemic’, *Erkenntnis*, 78/2 supplement:237–255
- Jansson, L. (2016), ‘Explanatory Asymmetries, Ground, and Ontological Dependence’, *Erkenntnis*, online first, doi:10.1007/s10670-016-9802-1
- Jenkins, C.S.I. (2013), ‘Explanation and Fundamentality’, in M. Hoeltje, B. Schnieder and A. Steinberg (eds.), *Varieties of Dependence* (Munich: Philosophia Verlag), 211-242
- Kim, J. (1974), ‘Noncausal Connections’, *Noûs*, 8/1: 41-52
- Koslicki, K. (2016), ‘Where grounding and causation part ways: comments on Schaffer’, *Philosophical Studies*, 173/1:101–112
- Kovacs, D. M. (2016), ‘Grounding and the Argument from Explanatoriness’, *Philosophical Studies*, online first, doi: 10.1007/s11098-016-0818-9

- Rodriguez-Pereyra, G. (2005), 'Why Truthmakers' in H. Beebe and J. Dodd (eds.), *Truthmakers: The Contemporary Debate* (Oxford University Press), 17-31
- Schaffer, J. (2016), 'Grounding in the image of causation', *Philosophical Studies*, 173/1: 49–100
- Weslake, B. (forthcoming), 'A Partial Theory of Actual Causation', *British Journal for the Philosophy of Science*
- Wilsch, T (2015), 'The Nomological Account of Ground', *Philosophical Studies*, 172/12: 3293–3312
- Wilsch, T. (2016), 'The deductive-nomological account of metaphysical explanation', *Australasian Journal of Philosophy*, 94/1:1–23
- Wilson, A. (2016), 'Grounding Entails Counterpossible Non-Triviality', *Philosophy and Phenomenological Research*, online first, doi: 10.1111/phpr.12305
- Wilson, A. (forthcoming), 'Metaphysical Causation', *Noûs*
- Wilson, J. (2014), 'No Work for a Theory of Grounding', *Inquiry: An Interdisciplinary Journal of Philosophy*, 57/5–6:535–579
- Woodward, J. (2003), *Making Things Happen: A Theory of Causal Explanation* (Oxford University Press)
- Woodward, J. (2015), 'Interventionism and Causal Exclusion', *Philosophy and Phenomenological Research*, XCI/2:303-347
- Woodward, J. (2016), 'The Problem of Variable Choice', *Synthese*, 193/4:1047–1072