Running Causation Aground

H.K. Andersen

Simon Fraser University, Department of Philosophy

Abstract: The reduction of grounding to causation, or each to a more general relation of which they are species, has sometimes been justified by the impressive inferential capacity of structural equation modelling, causal Bayes nets, and interventionist causal modelling. Many criticisms of this assimilation focus on how causation is inadequate for grounding. Here, I examine the other direction: how treating grounding in the image of causation makes the resulting view worse for causation. The distinctive features of causal modelling that make this connection appealing are distorted beyond use by forcing them to fit onto grounding. The very inferential strength that makes causation attractive is only possible because of a narrow construal of what counts as a causal relation; as soon as that broadens, the inferential capacity markedly diminishes. Making causation suitable for application to grounding spoils what was appealing about causation for this task in the first place. However, grounding need not appeal to causation: causal modelling does not have exclusive claim to structural equation modeling or other formal techniques of modelling structure. I offer a case in favour of a different kind of metaphysical frugality, which tend towards narrow, more restrictive construals of relations like causation or grounding, because then each relation behaves more homogenously. This more homogenous behavior delivers stronger inferential power per relation even though there may be more relations to which one is committed.

1. Introduction

If one is of a certain frugal bent, then committing to fewer distinctive metaphysical relations is always better than committing to more. The spareness here is a distant relation of Occam's Razor, an avoidance of unnecessary inflation of the apparatus required for metaphysical analysis. In this spirit, there is an attraction to finding a way to combine causation and grounding so that there is only one general kind of relation, which is called grounding when it holds between levels of fundamentality, and causation when it holds across moments of time. The reduction of grounding to causation, causation to grounding, or each to a more general relation of which they are both species cuts down on apparent philosophical baggage. Some of these projects have focused on the advantages of causation has in terms of modeling, involving inferential power that has greatly advanced in recent decades and which could be redeployed for grounding to gain a similar advantage in formal modelling. The motivation here is more than just frugality with respect to metaphysical commitments; it includes the impressive inferential capacity of the cluster of structural equation modelling, causal Bayes nets, and interventionist causal modelling.

A number of criticisms have already been levied against the assimilation of grounding and causation together, especially Bernstein (2016), Koslicki (2016), Raven (2015) (for a more general criticism of ground as supplanting other relations, including but not limited to causation, see Wilson 2014). Many of these concerns highlight the different behavior of causal and grounding relations, and how causation may be inadequate for the job required of grounding. Here, I examine the other direction: how treating grounding in the image of causation makes the resulting view of causation worse off for causation.

Bernstein (2016) makes a helpful distinction between two claims that could be made here: Unity is the claim that causation and grounding are the same relation, or species of one relation-genus; Illumination is the claim that grounding and causation can be illuminated by considering them together or in comparison. Wang (2020) lays out the arguments for and against Illumination, including reasons to be cautious about treating grounding in the image of causation insofar as causation may be insufficient for the requirements of an adequate account of grounding. She notes, though, that she does not take anyone to defend the stronger claim of Unity.

I will be addressing this stronger claim of Unity. I take Wilson (2018) to be espousing something stronger than Illumination, and will use his paper as an exemplar of this stronger claim because of the motivation he provides for Unity. My target in this paper is the stronger claim that grounding and causation are in some important sense two variations of the same relation (this includes but is not limited to locutions like, two species in the same genus; or grounding as a species of which causation is the genus), where the connection is either motivated by or based on the characteristic features of modelling found in interventionism, using variables and arrows to depict relationships, interventions and their outcomes to evaluate claims about relatedness, and structural equation models to represent these relationships numerically.

I will argue that the distinctive features of causal modelling that make this connection appealing in the first place are distorted beyond use by forcing them to fit onto grounding, undermining the motivation for such connection even if it is eventually made workable. The very inferential strength that makes causation attractive is only possible because of a narrow construal of what counts as a causal relation; as soon as what counts as a causal relation broadens, the inferential capacity markedly diminishes. Making causation suitable for application to grounding spoils what was appealing about causation for this task in the first place. The upside, though, is that grounding need not appeal to or somehow involve causation if the appeal is based on the use of structural equation modeling as useful for analysing grounding. Causal modelling does not have exclusive claim to structural equation modeling, nor to other formal techniques of modelling structure. Grounding can be modeled using a different version of these tools, involving different background conditions or assumptions. Causation need not be involved.

One consequence of this is to undermine the general motivation for bringing causation and grounding together as a way of reducing the number of metaphysical relations to some absolutely smallest number. There could be fewer, but then, we can do less with each one, in terms of the inferences that are licensed by identifying instances of it. I offer a case in favour of a different kind of metaphysical frugality, one that identifies those relations that pull their own weight most efficiently in terms of providing stronger inferences when the relation in question is identified. What are the most high-powered relations, where we get the most bang for the buck, as it were? This understanding of metaphysical frugality will tend towards narrow, more restrictive construals of relations like causation or grounding, because then each relation behaves more homogenously. This more homogenous behavior is what supports the inferential power that is the advantage of this understanding of metaphysical frugality with respect to relations. There may be a larger number of resulting relations to which we are committed than on the reductive approach to metaphysical frugality, but each such commitment brings a pay-off in terms of analysis power. Causation has already earned its place, on this more-relations-narrowly-construed understanding of metaphysical frugality.

I begin with an exegesis of Wilson's version of the claim that there is one genus of which nomological and metaphysical (i.e., grounding) causation are species. Section 3 explicates the Causal Markov Condition and its role in supporting the inferences between probabilistic relationships among variables and causal structure, and how it relates to intervention and modularity. Section 4 demonstrates how metaphysical grounding violates the Causal Markov Condition, and how the avenues of response to this compromise what makes causal modelling techniques effective. Section 5 considers narrower versus broader construals of causation, and makes the case for narrower construals of relations to yield higher-powered formal treatments. Section 6 concludes.

2. Grounding as Metaphysical Causation

Wilson (2018) presents the strongest version of the claim that grounding and causation are species of a single genus, or variations of the same relation. I am focusing on Wilson's piece for several reasons. The first is that he makes the strongest case for this connection (not merely claiming the strongest connection) in arguing for grounding as a kind of metaphysical causation. Wilson locates his own proposal with respect to others like Schaffer, Bennett, Fine, and Koslicki, where even other authors who argue for a closer connection between causation and grounding, like Bennett and Schaffer, nevertheless still treat them as connected yet distinct relations (see especially Wilson 2018, p. 724). Wilson wants to go further and claim that there just is one relation, and it is the genus of which we can then distinguish further species like metaphysical causation and nomological causation. This is clearer in terms of what the goals are for which we might be motivated to treat grounding as metaphysical causation. Wilson's proposal offers a specific and detailed proposal where he applies interventionism to cases of grounding to show examples of metaphysical grounding. This is helpful for identifying points of agreement, including on the relevance of interventionism for making claims about grounding and causation, and points of divergence, such as the requirements on causal modelling such as the Causal Markov Condition.

The second reason to focus on Wilson's proposal because we share an orientation toward the meta-metaphysical issue of how one can legitimately make such a case for or against assimilating two relations. He notes that "the having of a theoretical benefit can count in favour of a principle of fundamental metaphysics." (p. 724) Here, I take us to be in agreement in how one should make the selection: in comparing two proposals for how causation and grounding are, or are not, connected, identifying a theoretical benefit for one proposal over the other as more useful is a legitimate reason to adopt that proposal. I will focus on one class of theoretical benefits, loosely collected together as inferential power. Contemporary causation and causal modelling have impressive inferential power: given specific information about a system, it provides justification for a range of inferences to strong conclusions: not merely that there is a causal relation between two variables, for example, but the precise degree of strength, etc. In this sense, what Wilson finds appealing about causal modeling, as we'll see shortly, is also what I find appealing about it, and the disagreement then turns on the more specific issue of how to preserve that in connection to grounding.

This orientation toward effective ways to approach these metaphysical questions also shares a great deal with Hall's (2023) pragmatist approach to metaphysical analysis. Causation has proven its worth in practice; a pragmatist following Hall can respectfully decline to start with questions like, 'but is it *really* there?', or 'but is it really fundamental?', by instead looking at how the relevant concept operates in use and showing that it meets a high bar for usefulness. Techniques of causal modeling can produce genuine empirical success when used for novel predictions about causal structure in data sets; grounding has not yet done anything of the sort. When I speak of the inferential power of causal modelling, this is the understanding of 'use' I have in mind. Contemporary causation is noteworthy precisely because it emerged from a long tradition of views on causation with such detailed, refined, and well-justified inferential practices for modeling actual systems in the world, and has established a wealth of empirical cases where it has been not merely successful but rather astonishingly successful. This theoretical benefit counts in its favour metaphysically, for both Wilson and myself.

It is not common in contemporary metaphysics to use empirical success in methodological application to then reverse-justify the metaphysical commitments required to show that methodology to be justified. It has a faint bootstrapping air of circularity. Yet it is a time-honoured approach on a longer philosophical trajectory going back at least to Newton and Leibniz, in the vein of natural philosophy. The efficacy of a methodology can demonstrate at least prima facie reasons to think that it is justified overall, and then that overall justification of the metaphysical commitments can be made more specific by looking at what, precisely, is involved in the justification of the specific techniques that led to the success.

Wilson (2018) puts forward the view that causation is a genus category, of which there are at least two different species. One of these species is metaphysical causation, also known as grounding. The other species is what he calls nomological causation. Nomological causation is intended to correspond more closely to what is ordinarily thought of as causation, something that is tied to the laws of nature and which has, presumably, a lower degree of necessity than metaphysical causation but a higher degree than mere high probability. He uses nomological causation as a stand-in for probabilistic causation to highlight the similarity to grounding.

Wilson's claim is that G = MC; or, that grounding just is metaphysical causation. "According to G = MC, metaphysical causation and nomological causation are different species of the same genus." (p. 724) His motivation is twofold. The first involves simplicity, where he argues that this counts as ideological parsimony, having just one relation with two versions, rather than two relations. The second motivation involves unification of grounding and causal explanations, where unification is itself explanatory. Wilson also takes this unification of grounding and causal explanations to provide theoretical support for views about explanation where all explanations are causal, by removing some apparent counterexamples of grounding explanations as non-causal explanations.

Wilson proceeds, in the second part of the paper, to gives examples of specific causal systems with parallel grounding examples to illustrate his proposal. He draws on interventionism, and offers examples of how to convert grounding explanations into the form familiar from interventionist causal modeling. The variables are relabeled while the graphs remain structurally the same, within each example. He uses these graphs to demonstrate how both grounding and causation perform similarly, especially for inference and for behavior of counterfactuals.

He argues for G = MC by also noting how "grounding and nomological causation stand in the same general relations to laws, necessity, and inference." (p. 747) He offers grounding as specially connected to fundamentality, analogously to how nomological causation is connected to temporality: each orders similarly, with grounding doing so for levels of fundamentality and nomological causation doing so with respect to time ordering (p. 731). He takes metaphysical causation to be synchronic and not mediated by laws of nature, while nomological causation is diachronic and mediated by laws of nature (p. 730). He takes the differences in mediating role played by laws of nature to be a defining difference between the species metaphysical and nomological causation, under G = MC.

3. Causal Markov Condition: the engine under the hood

This section lays out one of the background assumptions that makes contemporary causal modelling so powerful, the Causal Markov condition (CMC). This condition connects to other key notions in causal modelling such as modularity and the technical definition of intervention, which means that it sits at the center of an especially powerful part of the inferential engine to move back and forth between probabilistic relations among data to directed acyclic graphs of causal structure. The CMC is required, as an assumption, for the inferences that have drawn such attention to causation in recent decades. This section will also briefly examine some ways in which the CMC can be violated, and the consequences this has for making inferences about causal structure.

There are variations on the exact wording, but one standard formulation of the Causal Markov Condition (CMC) is:

> Causal Markov Condition: Let G be a causal graph with vertex set V and P be a probability distribution over the vertices in V generated by the causal structure represented by G. G and P satisfy the Causal Markov Condition if and only if for every W in V, W is independent of V\(Descendants(W) ∪ Parents(W)) given Parents(W). (Spirtes, Glymour, and Scheines 2012, p. 29)

Another version of this condition can be found in Pearl (2009, p. 30). Woodward (2005) offers this formulation:

> CM says that, conditional on its parents or direct causes, every variable is independent of every other variable except its effects:
>
> (CM) For all Y distinct from X, if X does not cause Y, then *Pr(X|Parents(X))= Pr(X|Parents(X)*Y)*
>
> … As Hausmann and Woodward (1999) argue, insofar as there is any systematic connection between causation and conditional independence relations in acyclic graphs, it appears to be captured by CM. (p. 64).

It is a refinement of the screening-off condition, originally introduced by Reichenbach (1956). There is a lot packed into this condition, and it is a foundation stone for interventionist and causal Bayes net modeling, so it is worth getting clear on what meeting this condition requires.

One way to re-describe the CMC is that it picks out the probabilistic dependencies that are due to directed arrow between nodes in a graph. By conditionalizing on the parents of some variable X, all other, more distal causes of X (called ancestors) are screened off from X, which just means they are probabilistically independent of X given the parents of X. Then, either X causes some further variable Y, or Y neither increases nor decreases the probability of X given the parents of X. In the graph in Figure 1, for example, a simple causal chain is given, with one common cause. The CMC says that by conditioning on the parent of X, X will be probabilistically independent of every other variable in the system except Y.

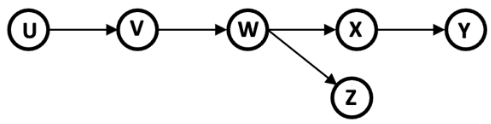

Figure 1

The role that the CMC plays is in connecting the graph of the causal structure with the probability distribution over the variable vertices (also called nodes). There are two directions of inference here. One is from a graph to a set of probabilistic dependencies and independencies. Using Figure 1 and the CMC, we can predict that P(X)≠P(X|Z). But once we condition on W, they will be independent: P(X|W)= P(X|W,Z). We can also predict that conditionalizing on W will screen X off from its ancestors, U and V: P(X)≠P(X|U, V) and P(X|W)= P(X|W,U,V). So, even though U and V have a causal path to X, and wiggling the values of U and V will wiggle the values of X, they are rendered independent by conditionalizing on the parent of X, which screens off their influence on X.

The second direction of inference is from the set of probabilistic in/dependencies, plus the outcomes of intervention, to the correct causal structure in the form of a directed acyclic graph (DAG). If conditioning on W fails to screen off some variable Z, then the CMC justifies the inference that Z is an effect of X. This point is extremely important in the next section of this paper: if it turns out that Z is not an effect of X (see Figure 1), and yet conditionalizing on the parent of X does not screen off Z, then the CMC has been violated. This could happen in several different ways. It could be that there is some mysterious and systematic correlation between X and Z, such that there is no causal connection between them but they remain correlated anyway, perhaps due to some extremely improbable coincidence. Assuming that CMC holds is assuming that there is no such cosmic trickery or coincidence in the system in question.[1] It could also mean that the variables overlap in the extension of their instances, or, they double-count: an instance of variable Z is also an instance of variable X. This will result in a persistent correlation of X and Z that is not causal yet cannot be screened off, violating the CMC.

This connects to the definition of an ideal intervention. Not everything that can be done thereby counts as an intervention. Crucially for the purposes of my argument here, the inference from probability distributions over a system of variables to the directed acyclic graph for that system cannot return a decisive answer about causal relations among variables given conditional and unconditional in/dependencies in the data drawn from the system, without also relying on the outcomes of interventions. This is key to contemporary causal modeling: the 'do' operator, such as introduced by Pearl in 2000, is required to differentiate variables that are probabilistically dependent because, for example, they are both effects of a common cause. By intervening on a target variable, that variable is made independent of its ordinary causes. Changing the value of that variable through intervention then reveals the effects of the target variable, rather than the common effects of the parent(s) of the target variable.

This also connects to the structural equations that represent the system. The definition of the CMC stipulates that variables X and Y must be distinct. There is more to the idea of distinctness than can be covered in this paper, but it is related to, or a distant descendant of, the general requirement that cause and effect be in some important way distinct that has come down from Aristotle in various forms over millenia. In this context, the requirement of distinctness is that X and Y must be independently intervenable-on. It must be at least in principle possible to intervene on Y, for example, without having to also change X. This requirement is often put in terms of modularity: arrows in a graph must be modular, in the sense that each arrow can be broken by an intervention on the variable at the head of the arrow without thereby changing any of the other arrows in the graph. If X and Z double-count instances, then they fail to be modular.

An ideal intervention is defined as fixing the value of the intervened-on variable to a specified value through means that are exogenous to the original graph. Informally, the

[1] This is at the root of a back and forth between Cartwright (2002, 2006) and Hausmann and Woodward (2004).

intervention sets the value of the target variable by doing something other than setting the value of the parents of the target variable in the graph. By intervening, we know what the value of that variable is, *and* that the arrow(s) from the parent variable(s) into the target variable is broken. Figure 2 adds an intervention to the graph from Figure 1. This intervention thus renders the intervened-on variable independent of its own parent(s). Even though W is a cause of X in the original graph, the intervention on X breaks the arrow from W to X, such that the P(X) = P(X|W) after intervention.

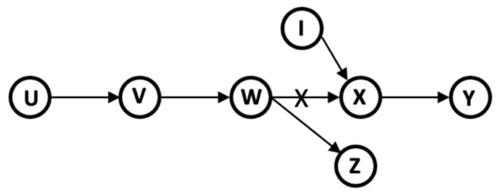

Figure 2

This is a key feature for interventions to serve their role in identifying the downstream effects of X. If we ‘wiggle’ the value of X using an intervention, then no variables in the graph change values *except* the effects of X. If we were to set the value of X using the existing causal structure, so, by ensuring that W took the value that would cause X, then wiggling the value of X, which would require wiggling the value of W, would also wiggle the value of Z. The intervention screens off Z and X from one another by adding a separate causal arrow into X that does not directly affect Y. Modularity of causal relations is required for this. If Z and X violated distinctness, by doublecounting, for example, then intervening on X would fail to be ideal, because it would also affect the value of Z.

If the Causal Markov Condition is violated, then the inferences that take us from probabilistic in/dependencies to graphical structure, or from graphical structure to expected probabilistic relationships among variables, will not be justified. The CMC, and the closely connected features like modularity, distinctness, and ideal intervention, are the engine under the hood driving the efficacy of causal modeling. These powerful inference techniques are not epistemically free: they are justified by, and thus require, assumptions such as CMC.

4. Wilson's Metaphysical Causation Requires Violation of the Causal Markov Condition

This brings us to the redeployment of interventionist modeling for grounding in the form of metaphysical causation. Put briefly: the variables for grounding will often violate the Causal

Markov Condition. It will be impossible to ensure that metaphysical causation meets the conditions required for an intervention, namely that one variable can be set to a specified value while the values of the other variables in the system are held fixed. And thus, interventionism gives us no traction on metaphysical grounding.

This sets up a dilemma: we could say that interventionism only holds for nomological grounding, since metaphysical causation violates conditions like CMC, but then interventionist modeling and structural equation models drawn from causation would fail to hold of metaphysical grounding; or we could make interventionism work for both nomological and metaphysical causation by eliminating requirements like CMC and the possibility of well-defined interventions. Given Wilson's usage of interventionist models and structural equation models, I take him to reject option 1. My argument is that this second horn might be suitable for G = MC, but leaves causation unequipped to account for what it originally did. It makes it less suitable for modelling causation. A key way to fail CMC, as seen in the previous section, is if some variable X remains probabilistically dependent on some other variable in the system that is not its effect, even after conditionalizing on its parents, including when a target variable remains dependent on its parents in the graph even after intervention. I'll illustrate how several grounding examples in Wilson (2018) fail CMC, though there is not space to go through each of Wilson's examples.

A first example involves Wilson's response to concerns about failures of asymmetry in grounding (pp. 727-728). His example is from Thompson (2016), of the weight, volume, and density of a cube. From any two of these quantities, the third can be derived, but there is no clear order or precedence in terms of which must ground which. They are interdependent. For any of these three variables, however, if they are put into a causal system as the potential effect (either of nomological or of metaphysical grounding), it is impossible to have any well-defined intervention on the effect that holds fixed the other two cause variables. The third variables just is a recounting of the other two variables: there is no possible way to screen it off from either of the other two. As a system of variables, for any kind of causation, this example violates the Causal Markov Condition. As such, it might be possible to write a set of equations that look like structural equations for the system, but it will not *behave* like the structural equation models of interventionism. There will be systematic correlations between these variables that do not respond to intervention in a way that is consistent with the CMC.

In general, when doing causal modeling, one aims to avoid these kinds of systems: the kinds of correlations they display are clearly a result of how the terms are defined, not the result of a causal relation between them, insofar as causation is taken to be any kind of genuine structure in the world. There are other resources for modeling these kinds of systems; causal modeling is ill-suited because definitional connections between the variables result in double-counting. This example highlights how a class of grounding examples involving these kinds of relationships among the variables will always violate CMC and thus be ill-fitted for interventionist causal modeling, unless requirements like CMC are removed.

Most of the examples come later, where Wilson works out examples of structural equation modeling in the interventionist approach for classic examples in causation literature, such as pre-emption, double pre-emption, late pre-emption, etc. Wilson helpfully sets out what he takes one of the features of structural causal modeling, namely, that the equations express counterfactual dependencies among the variables including what values would be if other variables were set to specific values. It is the counterfactuals that are doing the explanatory work (p. 740). Wilson's first main example, Simple: Window, uses binary variables to model a rock breaking a window, where the cause is throwing the rock and the effect is the window smashing. He then notes how "models with this structure also describe straightforward cases of grounding" (p. 741). For the parallel metaphysical causation case, the cause variable is whether Socrates exists, and the effect variable is whether Singleton Socrates exists. However, this Socrates example violates the CMC. There is no possible intervention that sets the value of Singleton Socrates in a way that severs the incoming 'causal' arrow from Socrates. If we add a variable that is the causal-ground for Socrates, then Singleton Socrates will not be screened off from changes in the grounds for Socrates, by conditioning on Socrates. The behavior of the grounding example under intervention and conditionalization is different than the behavior of the causation example, even with the same superficial graph structure. The causal example, and not the grounding example, meets the CMC. The arrows in the two graphs don't behave the same way under intervention.

This arises in the other examples as well. The constitution example in the early pre-emption cases is another where it is not possible to screen off ancestors of an effect by conditioning on the parents. The variables P and R will remain systematically correlated, even though they are not causally connected in the right way. P is "whether there is a person with exactly ten fingers here" and R is "Whether there is a person with exactly nine fingers here". C, "Whether my particles are arranged me-wise here" is a parent variable for both P and R, and P and R are both parents of E, "Whether there is a person here." The values for P and R are binary, and will *always* be anti-correlated. The probabilistic correlation between P and R thus does not reflect any real arrow in the graph, even when those arrows are taken to be metaphysical rather than nomological causation. And it cannot be eliminated by any conditionalization among variables. It violates CMC.

It is important to note that I am not claiming these examples don't provide an adequate analysis of grounding. It could be that this suffices for grounding. My point is that these examples, even taken as good illustrations of grounding, violate the basic background tenets of interventionist causal modeling. They might be adequate for grounding, but they fail these more hidden tests for causation; so grounding could not be a species of the genus *causation as found in interventionism*. My challenge to G = MC is that making grounding a species of causation means that *causation* could not continue requiring conditions like the CMC, because grounding violates this condition. The problem is in incorporating grounding into causation by removing what made causation so useful and effective for modeling.

This illustrates the dilemma briefly described earlier. On the first horn, we could say that interventionism only holds for nomological grounding, since metaphysical causation violates conditions like CMC, but then interventionist modeling and structural equation models drawn from causation would simply fail to hold of metaphysical grounding. The motivation for G = MC would be gone. On the second horn, we could make interventionism work for both nomological and metaphysical causation by eliminating these requirements like CMC. Given Wilson's usage of interventionist models and structural equation models in the examples, I take him to reject the first horn. My argument is that this second horn might be suitable for G = MC, but leaves causation unequipped to account for what it originally did - causation of the non-grounding sort.

Consider Wilson's response to the first dilemma horn in more detail. One way around violation of CMC would be to say that conditions such as CMC hold only of nomological causation, but not metaphysical causation; examples of metaphysical causation would thereby not violate a condition that is expected to hold only of nomological causation. Wilson considers this in response to an especially strong formulation called Independent Manipulability (Weslake 2011). Weslake's requirement of independent manipulability is not representative of interventionism specifically, or other variants of contemporary causal modelling, as it is stronger than the CMC or modularity or ideal interventions as described in the earlier section. Wilson's response is that Independent Manipulability can just fail to hold of some kinds of causation, while holding of others.

There are several problems with this response, though, when applied to CMC failures. The first is that it would diminish or even sever the explanatory link that Wilson (and many others; he is one example) takes to be a central motivation of this project. "Our grip on grounding is supposed to come in two other main ways: through examples, and through the connection with explanation" (p. 728). Causal explanation is powerful because there is more to it than simply labeling an explanation as causal. Calling an explanation causal is not a 'dormitive virtue' kind of explanation, where one simply summarizes that it is explanatory by calling it causal. The effectiveness of causal explanation involves the inferences licensed by citing a genuine causal connection in the explanation. And, those inferences do not appear for free - they depend on conditions like CMC. Eliminating the expectation that CMC holds of causation, and specifying that it only holds for one particular subtype of causation, attenuates the unification of the resulting explanations. What is currently called causal explanation, and called nomological causation by Wilson, would then have a higher status as an explanation than the metaphysical causation explanations, because the former and not the latter meet the CMC. By biting the bullet and saying that only some kinds of causation should meet the CMC, and metaphysical causation need not, nomological causation and metaphysical causation then become such different species that it is no longer clear that they do in fact share a genus other than that of being explanations.

A second problem with this response to the first dilemma horn is that it undermines Wilson's approach of using the examples to show why G = MC is appealing. The similarity between metaphysical and nomological causation using structural equation models will be at best superficial. The reasons to use structural equation models for causation is that these can be used for making various inferences and predictions, to coordinate prediction between probabilistic relationships among variables in data and in the causal graph for the system of variables, and for encapsulating modularity of the causal structure in the world into the representation in equation form. Modularity in equations tied to the CMC: when modularity fails, the screening-off of a variable from its ancestors conditional on its parents will fail. But this is exactly what happens in metaphysical causation structural equations. When A grounds B, and B grounds C, C is not screened off from any influence of A when B is conditionalized on. These examples illustrate how inter-connected the key terms in interventionism are, such as intervention, modularity, independence, and CMC.[2] But the appearance of similarity in structural equations belies the underlying differences where only some of those structural equation models (those for what would otherwise be called causation) can support stronger inferences while others (for grounding) cannot support such inferences.

There is a further possible escape from this horn of the dilemma, where the examples don't CMC, by treating the variables as mixed populations where there is a latent cause, not measured in the system, that is the actual common cause of the variables in the system (for example, see Spirtes et al. 2012, pp. 32-35). This wouldn't serve Wilson's purpose, however: it would require of grounding that any system of equations has at least one further latent variable that is not included. Adding more variables, to try and 'catch' the latent variable and add it into the regular system, would never succeed, because there would still be the persistent, leftover correlations that don't respond to interventions in the right way. Grounding as metaphysical causation would require that a latent-in-principle variable is involved in *any* graphical or equation model of grounding. I am leaving this out as a possible way out of the problem of violating CMC, because this seems to be contrary to grounding as involving fundamentality to allow that there must always be some latent variable accounting for the observed dependencies.

This leaves the second horn of the dilemma as the only remaining option. In order to use the modelling apparatus from interventionism, there are some basic axioms or conditions that must be met, such as Causal Markov Condition and Faithfulness Condition, and interventions that must be at least well defined even if not practically feasible. If those conditions are not met, then one simply can't use any of the inferences for which contemporary causal modelling is recognized, for which interventionism is quite literally named. In order to treat causation as comprising both metaphysical and nomological causation (even leaving probabilistic causation aside), the Causal Markov Condition, and conditions like well-defined interventions that break

---

[2] This is related to points made by Ney (2016), though my diagnosis of the issue differs from hers.

causal arrows, must be given up. Even if this suffices to make G = MC an adequate account of grounding, it renders the resulting version weaker for its original purposes, modelling causation.

I am urging us to reject the assimilation of grounding to causation if it requires us to accept that second horn of the dilemma, that of weakening causation by eliminating conditions for metaphysical causation that grounding does not meet. This can be a more appealing option for Wilson's own project that it might seem. I propose instead that there is no monopoly by causation on structural equations, nor on graphical models generically used. Interventionism does involve substantive reliance on those techniques but they are not proprietary to causation. Interventionism also involves reliance on some distinctive conditions or assumptions that are not apt for grounding. Instead of trying to use interventionism, with ill-fitting assumptions for grounding, it is possible to use structural equation or graphical models with *different* assumptions or conditions, ones that reflect how grounding behaves and the inferences it supports. This involves rejecting G = MC but keeps the motivation intact for using formal modeling techniques, just ones with different conditions that those for causation.

5. Narrow versus broad construals of causation

The less we try to expand causation, the more homogeneity there is among the instances we find of it, and the stronger we can make assumptions about it to support more powerful inferences. As the previous section argued, expanding causation to include metaphysical causation means giving up conditions like the Causal Markov Condition; and as section 3 showed, the Causal Markov Condition is central to the extraordinary success of causal modeling in recent years. Putting grounding under the umbrella of causation stretches that umbrella so far it no longer covers what it used to. This section explores the metaphysical advantages of keeping a certain kind of narrow construal of causation. Wilson's nomological causation is too broad in one regard, where the expansion of causation to include metaphysical causation waters down what was explanatorily distinctive about causal explanation. In another regard, though, it is too narrow to capture causation as it is found in interventionism, because it leaves out probabilistic causation.

I've argued elsewhere (Andersen 2018) that we should take a narrow construal of causation, counting fewer examples as ones of causation rather than some other relation, in order to preserve its power in making inferences between probabilistic relationships among variables and causal structure. Those arguments were directed against the very broad construal of what counts as a causal explanation offered by Lange (2013, 2016), where any explanation that situates the explanadum in the network of the world's causal relations is thereby causal. Lange's motivations were not to illuminate anything about causal explanation: he wanted to show how, even given the broadest and most generous construal of causal explanation, there were still explanations that fell outside that range in the sciences, specifically ones involving certain mathematical relationships. As such, there was no reason related to causation itself for Lange to give such a broad construal; it served as a generous way of setting up an opponent's viewpoint.

The reasons I gave there for holding a narrow view also apply here, though. If we want causal explanation to be genuinely substantive, and not watered down into a mere synonym of explanation generally, we need criteria by which to differentiate causal explanation from a plethora of other types of explanation. By drawing on the success of causal explanation to attempt to illuminate other types of explanation, we end up diminishing the success of causal explanation. By trying to make causal explanation serve both for causal explanations narrowly conceived, and also for grounding explanations, what those explanations have in common will of necessity be weaker. It might add explanatory benefit to grounding but it does so at a steep cost to causal explanations, which would now imply far less about the phenomena when successfully identified. Broadening causal explanation to include grounding means that less is conveyed in such explanations.

There is another issue that arises with Wilson's proposal of making causation the genus of which nomological causation is a species, namely, that nomological causation is too narrow for contemporary causation. Wilson distinguishes metaphysical from nomological causation because only the latter, and not the former, involves mediation by a law of nature (p. 740). Yet it is key to contemporary discussions of causation, including the interventionist approach on which he draws, that causation *not* be required to be nomological. It must allow for, and in the overwhelming majority of cases simply is, probabilistic rather than deterministic.[3] Nomological

[3] Wilson's nomological causation seems to involve deterministic laws, so I follow that here. There could be probabilistic or statistical laws; such cases would attenuate rather than strengthen the connection between grounding and nomological causation.

causal explanation thus leaves out a massive amount of existing causal explanation, anything where the connections in the graph are probabilistic rather than binary. Even when laws are involved, they play a highly mitigated role, and in interventionism, it is extremely unusual for a law to make any explicit appearance in systems with probabilistic causal connections. One often has to deliberately set up a system to only involve deterministic relationships of 0 or 1; in science, the overwhelming number of systems, however they are counted, involve probabilities between 0 and 1. Nomological causation is thus changing the topic: interventionism is not a theory of *nomological* causation. It is a theory for causation that includes the full range, probabilities between 0 and 1 inclusive.

There is a helpful way to situate Wilson's proposal against a long historical trajectory involving laws, determinism, explanation, and causation. Russell famously rejected causation, but he rejected a version of it that involves necessitation, what we would now call deterministic laws; or, as aptly, nomological causation. Russell's arguments are directed toward a much older understanding of causation more akin to the Principle of Sufficient Reason; it bears no real resemblance to contemporary accounts of causation. Hempel and Oppenheim (1948) also treated causal explanation as equivalent to deterministic lawful relationships; what they describe matches Wilson's nomological causation closely. After that, though, from Reichenbach's (1956) treatment of conjunctive forks and screening-off, discussions of causation began to distinguish between lawful relationships, which involved some degree of modality stronger than mere probabilistic connection, and causation, which was distinguished by including probabilities. Suppes (1970) decisively established probabilistic changes as characteristic or revealing of causal relationships, furthered in work such as Rosen (1978) and Salmon (1980). It was the advent of the 'do' operator (Pearl, 2000 edition) and formal definitions of interventions that added to the probabilistic relationships the final required step needed to different causal structure from correlations. The better those accounts of causation and causal inference got at accommodating genuinely probabilistic causation, the less well-suited causation became for treating the nomological, where outcomes are either necessitated or rendered impossible. It is an open philosophical question whether lawful relationships among variables in a physics system (to take a toy example, such as force, mass, and acceleration in F=ma), are best understood as causal, in that form, or nomological *rather than* causal. But if they are part of causation, they are at best a peripheral example, and they are problematic precisely because they do not admit of well-defined interventions (see also Ney 2016; Woodward 2005, pp. 129-131).

This means that changing nomological causation in order to accommodate genuinely probabilistic causal relationships is going to pull grounding and causation even further apart from one another. This could be extremely revealing: whatever it is that has to be given up with respect to the attempted analogy or structural similarity between nomological causation and grounding just is part of the distinctive character of either grounding, or causation. Having to put the probabilistic treatment back into causation, expanding the overly narrow construal of causation as nomological causation, means that there is an opportunity to identify more precisely where there is a rupture or divide between what is required for an adequate account of causation to serve as the basis for the justified inferences discussed above, versus what would be required for such an account of grounding that could serve as the basis for formal treatment.

Recall that part of Wilson's motivation is the prospect of unifying explanation (2018, p. 724). By expanding the range of causal explanation, there is some unification of apparently different kinds of explanation as all related to causal explanation. But at the same time, what counts as causal explanation has now been watered down so much as to have little if anything in common, other than being explanatory. This renders causal explanation equivalent to explanation simpliciter. While that does unify, it does so at a great cost: there is no longer anything substantive to having a *causal* explanation, rather than some other kind of explanation. To say it is causal explanation says nothing more than just saying it is an explanation, an avenue that is already available to us. The 'causal' part becomes superfluous, contributing nothing over and above mere explanation, and unable to adequately convey what remains distinctive about the class of explanations that were formerly known as causal under the narrower construal. We can instead choose to keep the narrower construal of causation, so that there is real substance in saying that an explanation is specifically a *causal* explanation, rather than some other kind of explanation, leaving it for an account of explanation simpliciter to articulate in what the unity consists. And then it can be broadened to include probabilistic causal relationships as well as deterministic ones, since we already have very good reason to think these behave in a sufficiently similar way (as evidenced by the effectiveness of contemporary causal modelling in doing just this).

I conclude this exploratory section by noting a useful parallel in this discussion about CMC with a discussion about another key assumption in causal modeling, Causal Faithfulness (CF). CF can fail when there are systems that maintain homeostatic equilibrium: thermostats are an artificial example (Zhang and Spirtes 2008). Biological systems are rife with examples of CF failure - we should expect evolution to have produced highly effective CF-violating causal structures for maintaining important homeostatic systems in organisms (Andersen 2013). What does this mean, for using causal modeling for systems like this? It largely means that the strongest versions of the inferences between probabilistic in/dependencies and causal structure cannot be used. The conditions for their justification are not met and they will be unreliable. Instead, there are weaker assumptions that can be used instead, to work around the specific failures of CF (e.g., Ramsay, Zhang, and Spirtes 2012; Weinberger 2018). These weaker versions are less powerful but are the strongest available under those circumstances, and will vary depending on the precise circumstances of the CF violation(s). Finding that CF does not hold of some system means losing modeling power. Weaker assumptions justify weaker inferences.

Bringing this back to the case of the CMC, it is worse. It is possible to work around a failure of CF, as long that failure is not rampant and it only takes certain triangular forms in the graph. The CMC is more foundational: if the CMC does not hold, then there are no inferences between probabilistic relationships among the variables and causal structure. There are no interventions in the system that will be revealing; modularity fails. If the CMC does not hold, then we have lost out on the inferential power that makes causal modeling so attractive in the first place. The more that gets included under the umbrella of causation, the weaker the assumptions are that will hold of everything under that umbrella, and the fewer inferences will be available. G = MC might achieve some clarity for grounding but does so at great cost for causation.

6. Conclusion: an optimistic note for proliferation rather than reduction

I'll conclude by considering the meta-metaphysical question of what is at stake in introducing or retaining a new metaphysical relation and what reasons can count for or against. What are the conditions that must be met for such a new relation to get to join the metaphysical club? Or for an existing member to stay, if a shinier, newer member might take its place? There is jostling, as it were, where the introduction of a new relation may kick out some older relation. Wilson invokes considerations of use and aptness in philosophical analysis as a reason to endorse or reject particular accounts of grounding. He argues that G = MC provides parsimony: there are fewer relations posited as primitive, which is, on one way of thinking, preferable. If one *can* reduce from 2 relations to 1, one *ought* to do so, if fewer relations are always better than more; it keeps the metaphysical club properly exclusive. The urge to treat grounding in the image of causation seems to be motivated by appreciation for the frugality of desert landscapes. Wilson (2018) has served as a representative in this paper, but it is a widespread inclination. The overarching metaphysical picture is one where the ideal is some single, unifying, broadest relation or category, from within which further distinctions foliate into species unified by that one genus.

I'd like to endorse Wilson's focus on use and aptness as reasons to accept metaphysical proposals like this, and then urge a different view of how to think about frugality with respect to metaphysical relations, how many of them we ought to aim at, and what is required to introduce a new one or dismiss an old one. By putting inferential efficacy as a central desideratum for a relation, each given relation becomes more inferentially powerful when there are more of them, each construed more narrowly. This is another way of being metaphysically frugal: making each relation earn their keep inferentially, rather than reductively. By having more relations, stronger conditions can hold for each, justifying stronger inferences.

If one is serious about grounding as metaphysical causation, then this connection should be doing some work, not reshuffling labels. Given the impressive advancements in causal modeling, in several different formal treatments, it is part of the motivation, not merely a corollary, in putting grounding and causation together that grounding be modeled using the same kinds of formal treatments as causation. But as we have seen, that will not work: what it would take to modify those formal tools to make them suitable for application to grounding also compromises them for causation. It would involve giving up much of the inferential power of causal modelling for causation to make space under the umbrella for grounding as well.

However, this should not be understood as leaving grounding in the cold in terms of taking advantage of formal tools. Causation has no particular claim to exclusivity with respect to these formal tools. Many of these were borrowed from other fields' modeling practices, such as structural equation modeling in economics. Formal search methods and inference algorithms are not the sole purview of causation. What it is that differentiates ground from causation, rather than what they have in common, is precisely where we are likeliest to find the kinds of assumptions required to make formal treatments of grounding work.[4]

References:

Andersen, H. 2023. "Causal Modeling and the Efficacy of Action," in Brent and Miracchi Titus, eds. (2022, 164-181).

- 2018. “Complements, Not Competitors: Causal and Mathematical Explanations,” British Journal for the Philosophy of Science, 69 (2): 485-508.

- 2013. "When to expect violations of causal faithfulness and why it matters," Philosophy of Science, 80 (5): 672-683.

Andersen, H. and Mitchell, S, eds., 2023. The Pragmatist Challenge: Pragmatist Metaphysics for Philosophy of Science, Oxford, UK: Oxford University Press.

Bernstein, S. 2016. “Grounding Is Not Causation,” Philosophical Perspectives 30: 21–38.

Brent, M. and Miracchi Titus, eds 2023. Mental Action and the Conscious Mind, New York, NY: Routledge.

Campbell, J. 2010. "Independence of variables in mental causation," Philosophical Issues, 20: 64-79.

Cartwright, N. 2002. "Against Modularity, the Causal Markov Condition, and Any Link Between the Two: Comments on Hausman and Woodward," British Journal for the Philosophy of Science, 53 (3): 411-453.

- 2006. "From metaphysics to method: Comments on manipulability and the causal Markov condition," The British Journal for the Philosophy of Science 57 (1): 197-218.

Hall, E. 2023. “Respectful Deflationism”, in Andersen and Mitchell, eds (2023, 133-159).

[4] Many thanks to Jenn Wang and to Al Wilson and his research group for discussions on this, and to Christopher Daly and Fraser McBride for helpful comments. I am grateful for the opportunity to live and work on the traditional and unceded territory of the Squamish, Musqueam, Tsleil Waututh, and Kwikwetlam First Nations.

Hausman, D. M., & Woodward, J. 2004. Modularity and the Causal Markov Condition: A Restatement. British Journal for the Philosophy of Science, 55 (1): 147-161.

Hempel, C.G., & Oppenheim, P. 1948. Studies in the Logic of Explanation. Philosophy of science, 15 (2): 135-175.

Jago, M. ed., 2016. Reality Making, Oxford, UK: Oxford University Press.

Koslicki, K. 2016. "Where Grounding and Causation Part Ways: A Reply to Jonathan Schaffer," Philosophical Studies 173 (1): 101–12.

Lange, M. (2013). "What makes a scientific explanation distinctively mathematical?" The British Journal for the Philosophy of Science 64 (3): 485-511.

- (2016). Because Without Cause: Non-Casual Explanations in Science and Mathematics, Oxford, UK: Oxford University Press.

Ramsey, J., Zhang, J., & Spirtes, P. L. 2012. "Adjacency-faithfulness and conservative causal inference," arXiv preprint arXiv:1206.6843.

Raven, M. 2015. "Ground," Philosophy Compass 10 (5): 322–33.

- ed., 2020. Routledge Handbook of Metaphysical Grounding, New York, NY: Routledge.

Reichenbach, H. 1956. The direction of time. University of California Press.

Rosen, D. 1978. "In Defense of a Probabilistic Theory of Causality," Philosophy of Science 45 (4): 604-613.

Salmon, W. (1980). "Probabilistic Causality," Pacific Philosophical Quarterly 60 (1-2): 50-74.

Schaffer, J. 2016. "Grounding in the Image of Causation," Philosophical Studies 173 (1): 49–100.

Suppes, P. (1970) A probabilistic theory of causality. Amsterdam: North-Holland.

Thompson, N. 2016. "Metaphysical interdependence," in M. Jago, ed. (2016, 38-56).

Wang, J. 2020. "Cause," in M. Raven, ed. (2020, 300-311).

Weinberger, N. 2018. "Faithfulness, coordination and causal coincidences," Erkenntnis, 83 (2): 113-133.

Weslake, B. 2011. "Exclusion excluded," International Studies in the Philosophy of Science.

Wilson, A. 2018. "Metaphysical Causation," Noûs 52 (4): 723–51.

Wilson, J. 2014. "No Work for a Theory of Grounding," Inquiry 57 (5–6): 535–79.

Zhang, J., & Spirtes, P. 2008. "Detection of unfaithfulness and robust causal inference," Minds and Machines, 18 (2): 239-271.