

Complements, Not Competitors: Causal and Mathematical Explanations

Holly Andersen

ABSTRACT

A finer-grained delineation of a given explanandum reveals a nexus of closely related causal and non-causal explanations, complementing one another in ways that yield further explanatory traction on the phenomenon in question. By taking a narrower construal of what counts as a causal explanation, a new class of distinctively mathematical explanations pops into focus; Lange's ([2013]) characterization of distinctively mathematical explanations can be extended to cover these. This new class of distinctively mathematical explanations is illustrated with the Lotka–Volterra equations. There are at least two distinct ways those equations might hold of a system, one of which yields straightforwardly causal explanations, and another that yields explanations that are distinctively mathematical in terms of nomological strength. In the first case, one first picks out a system or class of systems, and finds that the equations hold in a causal–explanatory way. In the second case, one starts with the equations and explanations that must apply to any system of which the equations hold, and only then turns to the world to see of what, if any, systems it does in fact hold. Using this new way in which a model might hold of a system, I highlight four specific avenues by which causal and non-causal explanations can complement one another.

1. *Introduction*

2. *Delineating the Boundaries of Causal Explanation*

2.1. *Why construe causal explanation narrowly? The land of explanation versus grain-focusing*

2.2. *Reasons to narrow the scope of causal explanation*

3. *Broadening the Scope of Mathematical Explanation*

4. *Lotka–Volterra: Same Model, Different Explanation Types*

4.1. *General biocide in the Lotka–Volterra model*

4.2. *Two ways a model can hold, yielding causal versus mathematical explanations*

5. *Four Complementary Relationships between Mathematical and Causal Explanation*
 - 5.1. *Slight reformulations of explananda*
 - 5.2. *Causal distortion of idealized mathematical models*
 - 5.3. *Partial explanations requiring supplementation*
 - 5.4. *Explanatory dimensionality*
 6. *Conclusion*
-

1 Introduction

Causal explanation has been the focus of intense work in the last two decades, which makes it useful to consider the nature of the boundary between causal and non-causal explanations. I will consider the question of what we see when we do a fine-grained examination of that boundary. What we find is that many explananda are more like clusters of phenomena: the original phenomenon for which an explanation was sought breaks into a variety of more precisely delineated phenomena that when taken together with the right interrelationships between them, constitute the explanatory cluster that was the original phenomenon. As a consequence of this, the more clearly we specify an explanandum in question, the more we reveal closely related causal and non-causal explanations for slight variations in the formulation. Small changes are just enough to rock us back and forth across the distinction between causal and non-causal explanation, which is helpful in understanding the boundaries between causal and non-causal explanation while highlighting the importance of precisely specified explananda.

The tendency to conceive of scientific explanation solely or primarily in terms of causal explanation results in part from the ubiquity of causal structure in the world. At the same time, there are many relationships that can serve an explanatory role that are not causal. Mathematical relationships, part-whole relationships, and even relationships in a taxonomy can explain by situating the explanandum with respect to the explanans in terms of a relationship that is not straightforwardly causal. Explanations might seem to be causal, or mostly causal, if we are insufficiently 'focused' in terms of specifying which explanandum is at stake. But once we clarify exactly what is being explained, and by what it is being explained, then a host of non-causal relationships pop into focus as key to explanation, situated closely all around causal explanation(s). Slight reformulations in the explanandum will change the relevant explanation(s) from causal to non-causal and back again. When considering scientific explanation, we are overwhelmingly often in the territory where causal and non-causal explanations fall closely together.

Regardless of the specific kind of non-causal scientific explanation being considered, there will be a multitude of causal explanations just next door, so that small reformulations of what precisely is being explained will shift the explanans from non-causal to causal.

I'll be using the term 'non-causal' as a catch-all term for any kind of explanation that isn't causal. This negative characterization of non-causal in terms of what it is not results in a wildly heterogeneous category. This sheds light on both causal explanation and non-causal explanation in terms of where they leave off. I will consider a particular kind of non-causal explanation—distinctively mathematical explanations—as a way of understanding an important proper subset of causal versus non-causal explanations.

This article offers several new points about causal explanations, how explanations can complement, and a distinction between two ways in which a model might hold of a system such that it changes the character of the resulting explanations from causal to mathematical. I do so in the service of also making a broader point about causal versus non-causal explanation, where my goal is not to offer a novel account of that distinction. The materials to characterize the boundaries where causal explanations leave off and non-causal explanations begin are already around, in my view, and simply need to be dusted off and redeployed. With this aim, I make use of some of the classic features of explanation, ones which have appeared over the last fifty years, even while specific accounts of the structure of explanation have changed. The points made here will be applicable to any particular account of explanation, causal or not—this includes new mechanistic explanation, interventionist causal explanation, distinctively mathematical scientific explanations, and more.

The approach offered in this article for causal versus non-causal explanations will have consequences for further questions such as the explanatory role that mathematics and mathematically formulated laws may play in explanation (consider especially Pincock [2012], [2015]; Saatsi [2011]), the role of models in explanation (Sterrett [2002]; Bokulich [2008]), and alternative forms of explanation that are not straightforwardly causal (Batterman [2002], [2010]; Batterman and Rice [2014]; Reutlinger [2014]). My discussion here should be useful to a broad swath of these other, more specifically focused, discussions about particular ways in which distinct explanatory techniques are to be understood or deployed. For instance, Skow ([2014]) has argued that existing examples of non-causal explanations are unconvincing, in that they can be accommodated within causal explanation. On the view offered in this article, the unconvincing examples indicated by Skow look more like insufficiently well-clarified explananda, such that further clarification would reveal both that he has identified genuinely causal explanations, and that the original

examples could be reformulated in slightly different ways to retain their non-causal character.

A main take-away message will be the way in which philosophical fruit is borne of an increased focus on precision and clarity in stipulating what precisely is getting explained, what precisely is doing the explaining, and what precisely the explanatory relation between explanans and explanandum is taken to be. Simply fine-tuning our focus in characterizing explananda will have an enormous benefit in allowing a rich array of explanatory techniques and relations to pop into view.

2 Delineating the Boundaries of Causal Explanation

There are three main elements in any explanation. The details about what can fill these roles will vary depending on other commitments, such as an ontic versus propositional construal of explanation. But all accounts agree on these general features. The explanation also includes an explanans, and some relationship between the explanans and explanandum. The explanation explains by situating the explanandum in terms of the explanans via that relationship. An explanandum is a target for explanation: anything that can be identified for which an explanation is possible. While this might seem unhelpfully vague, it reflects the wild heterogeneity of targets for explanation that we actually find in the sciences and beyond. Why we identify particular targets as worth explaining is an important question I set aside for this article, and instead start from the point when a target for explanation has already been identified. An explanation need not be total or even particularly satisfying to count as explanatory. Weak, partial, even barely helpful explanations are still explanatory of something, even if they leave much else unexplained.

2.1 Why construe causal explanation narrowly? The land of explanation versus grain-focusing

Compare a narrow versus broad construal of causal explanation:

Narrow: Any explanation that provides parts of the explanandum's causal history, including contextually relevant features that might not be in its direct causal past, as explanans. Or, the connection between explanans and explanandum is a causal relationship(s)

Broad: Any explanation that explains by virtue of situating an explanandum in the network of causal relations in the world. Or, any of the explanans, explanandum, or connection between them involves a causal relationship(s) or relata

The key difference is that on the narrow construal, only those explanations where the relationship connecting explanans and explanandum is causal will thereby be causal explanations. There might be causal relationships that are themselves the explanandum, or that are the explanans, but where the relationship connecting explanans and explanandum does not involve tracing out a causal history, and thus are thereby non-causal explanations. On the broad conception, any explanation that has a causal element anywhere—a causal relationship or causal relata as the explanans, explanandum, or the connection between them—will count as a causal explanation. On the narrow construal, for instance, there might be non-causal explanations of causal structures, whereas on the broad construal, any such explanation would have to count as causal.

The narrower our construal of causal explanation, the more useful it is for identifying and analysing particular instances of causal explanation that we find in the sciences. It allows us to say something significant by calling a particular explanation 'causal'. The more broadly we construe causal explanation, on the other hand, the less significance is involved in labelling a particular explanation as a causal explanation. At the extreme end of a broad construal, causal explanation merely means nothing other than explanation. This prevents us from being able to say anything about the distinctive character of causal explanation, and runs a host of different explanatory techniques together into a muddy wash.

On the broad view, one might be tempted to think of the boundary that divides causal and non-causal explanations like provinces in the land of explanation. The provincial view construes an entire domain as to be explained by one or other kind of explanation, be it causal or non-causal. The boundaries between them take you entirely out of one kind of explanation and entirely into another; the border strictly divides them from one another. One might even think that they occasionally have border skirmishes, fighting over control of a given explanation as belonging in one province rather than another. Some explanations might be hard to locate, in terms of finding the right land in which they live.

This is a view I will argue we should entirely reject in favour of a different metaphor that goes along with the narrow construal. Consider an old-fashioned darkroom enlarger used for printing black-and-white photographs from film. Once the image to be printed looks sharp enough to the naked eye, one can still improve the picture quality by using a grain focuser. Using it to make the individual grains of silver visible, one usually sees a gently variegated smear of greys, black, and whites, continuously shading into one another. Slight additional adjustments with the knob, all within the range of what looked sharp to the naked eye, result in a sudden pop into focus of the

individual grains of silver on the film.¹ With the grain enlarger, it becomes apparent that the image really is composed of distinct dots or grains, large and small, often with irregular edges, all nestled up closely and together composing the image.

This metaphor results in a very different picture of how causal and non-causal explanations fit together. Explanations might seem to be causal, or mostly causal, if we squint or if we are insufficiently ‘focused’ in terms of specifying which explanandum is at stake. But, once we clarify exactly what is being explained and by what it is being explained, then a host of non-causal relationships pop into focus, situated around the causal explanation(s). Slight reformulations in the explanandum will change the relevant explanation(s) from causal to non-causal and back again. The tendency to conceive of scientific explanation solely or primarily in terms of causal explanation results in part from the ubiquity of causal structure in the world. Regardless of the specific kind of non-causal scientific explanation being considered, there will be a multitude of causal explanations just next door, in that small reformulations of what precisely is being explained will shift the explanans from non-causal to causal.

The metaphor of grain-focusing thus means using specificity in delineating given explananda. Shifting the way in which an explanandum is formulated, even subtly, is overwhelmingly likely to change the required explanans enough to skew the result in the direction of causation. There are always some causal explanations of something in the vicinity for almost every explanandum, but they may not be causal explanations of the original explanandum.

2.2 Reasons to narrow the scope of causal explanation

A useful counterpoint is Lange’s ([2013]) construal of distinctively mathematical explanations. He takes up the question of what makes some scientific explanations distinctively mathematical rather than causal. The explanations he discusses are not simply mathematical explanations, where one bit of mathematics explains another, nor are they simply explanations where mathematical representations are involved. They are scientific in that they are about explananda in the world, but involve mathematical relationships in the explanation in a way that goes beyond mere mathematical representation.

Lange’s ([2013], p. 493) characterization of distinctively mathematical explanations starts with what he takes causal explanation to be: ‘I will adopt a broad conception of what makes an explanation “causal”: it explains by virtue of describing the contextually relevant features of the result’s causal history or,

¹ Note that the issue of physical size is not what drives this metaphor. It is not that zooming in to a microscopic size reveals better explanations; rather, the magnification is analogous to conceptual precision in formulating the explananda and explanans.

more broadly, of the world's network of causal relations'. He cites (Lange [2013], p. 487) the following passage of Salmon as evidence of Salmon holding that all explanations are causal, and adopts a broad reading of Salmon's description of explanation: 'To give scientific explanations is to show how events and statistical regularities fit into the causal structure of the world'.

This construal of causal explanation is, as he puts it, broad. This makes sense in light of Lange's task of distinguishing distinctively mathematical explanations from causal explanations. A more generous definition of causal makes his argumentative task more difficult and the resulting conclusion more compelling. However, I'll urge adoption of the narrow rather than broad construal for general use: the original construal of causal explanation is so broad that it dilutes the usefulness of the category of causal explanation by dramatically reducing the commonality between the different kinds of explanations that all thus qualify as causal.²

Why should the narrow rather than broad construal of causal explanation be adopted? There are several reasons: The first is that we can assent to Salmon's description of explanation as situating the explanandum in the network of causal relations in the world without thereby committing to all such situatings being causal situatings. For non-causal situatings of explananda in the network of causal relationships, the explanatory work is not done through causal relationships—it is not by tracing along the pathways of the causal network that the explanandum gets explained. Consider 'new mechanism' explanations (Andersen [2014a], [2014b]). A mechanism can explain a phenomenon when that phenomenon is the product of the organized causal chain at the termination conditions of the mechanism: this would be a causal explanation on the narrow and broad construals both. A mechanism can also explain when it gives rise to, or constitutes, the phenomenon to be explained. This clearly situates the explanandum in the network of causal relations, since a mechanism is just a special kind of recurrent, organized causal chain in that network. On this broad construal, this would also be a causal explanation, just like the first explanation. But on the narrow construal, it would be a constitutive, rather than causal, explanation. It is non-causal even though there are clearly causal relationships and relata involved in the explanation; it is the special causal structure of the mechanism that is the explanans. But the connection between explanandum and explanans is

² Skow ([2014]) offers a characterization of causal explanation that involves a broad versus narrow distinction, but his concern is more about the totality of explanation: while he rejects an overly narrow view of causal explanation, he takes it to be a rejection of the idea that a causal explanation must be complete, total, or otherwise sufficient to count as explanatory. In this regard, both Skow and I agree that a causal explanation can be both causal and explanatory without explaining the totality of an explanans' past causal history. It is not clear whether he would endorse the construal of narrowness I offer here.

one of constitution, not causation. The broad construal is unable to mark this difference.

Another reason to adopt the narrow construal is that it allows for a more accurate portrayal of the diversity of explanatory practices, especially in the sciences. By allowing any explanation involving causation in any way to count as causal, we smear together a host of distinct but closely related explananda and connecting relationships. Maintaining our explanatory resources in terms of distinct types of connections that can be explanatorily deployed is necessary to accurately describe and distinguish the wide variety of species of explanations found in the sciences.

A third reason involves a stronger way in which two kinds of explanations might be related. While more precise specifications of explananda may often enough yield two closely related but distinct explananda, that is not always or necessarily the case: sometimes there may be more than one kind of explanation available for a single well-specified explanandum. An explanation in which the explanandum is constituted by the explanans is another example of a non-causal situating in the causal network. In contrast, that same explanandum could be given an alternative causal explanation in terms of the causally upstream portions of the network, rather than in terms of what constitutes it. These are two distinct explanatory perspectives on the same explanandum, which can be extraordinarily illuminating. Losing focus on these distinct ways of situating something in the causal network means losing the ability to understand the full range of existing explanations and how they fit together.

Finally, adopting the narrow construal encourages good practice in terms of formulating explananda in a sufficiently precise way. As we'll see in the next section, rocking back and forth between slightly different formulations of the phenomenon to be explained results in shifting between causal and non-causal explanations for the closely related but non-identical explananda. When considering scientific explanation, we are often in the territory where causal and non-causal explanations fall closely together. If we accept Salmon's view that scientific explanations involve situating explananda in the network of causal relations in the world, we are always in the immediate vicinity of causal relationships. This means that it is deceptively easy to phrase explananda in vague ways that make it appear causal, by latching onto whatever causal relationships are in the vicinity. If we switch what precisely is getting explained, we thereby switch what, precisely, can explain it. Fudging the explanandum will almost always result in apparently causal explanations. More precise formulation of explananda, on the other hand, will bring into focus the non-causal relationships nestled nearby, and will highlight the details in those explananda that result in shifts between causal and non-causal explanations.

3 Broadening the Scope of Mathematical Explanation

Lange identifies distinctively mathematical explanations by the extra modal force they have compared to the necessity associated with causal laws. He illustrates this with the example of a mother trying to divide twenty-three strawberries evenly among three children without cutting any berries. There is no way to do this, and the explanation—that twenty-three is not evenly divisible by three—puts a constraint on any causal laws. There are no causal laws, actual or possible, such that they would result in a causal process by which twenty-three strawberries could be evenly divided. The modal force of the explanation, therefore, must come from something other than a causal relation—in this case, from the mathematical relation.

The class of explanations Lange thus identifies is that whose explanations that apply anywhere, in any possible world. There is no causal process, under any circumstances anywhere, such that it could involve an even division of twenty-three by three. No physical laws or facts are presupposed or utilized. These distinctively mathematical explanations are indeed quite interesting in terms of the explanatory relationship connecting explanandum and explanans, which results in a degree of necessity higher than any causal laws could achieve. It is striking to find explanations of physical events that involve purely mathematical relationships as key to the explanation (this is in contrast to the role mathematics plays when causal relationships are represented mathematically, or even when arguably non-causal physical laws in mathematical form are involved).

In this section, I show how some of the extra ‘space’ in the territory of explanation opened up by a narrow construal of causal explanation, in combination with an expanded version of Lange’s criterion, reveals the existence of an additional set of distinctively mathematical explanations. These explanations have a weaker degree of necessity than the very strongest degree of necessity associated with claims such as ‘twenty-three is not evenly divisible by three’, but they are nevertheless modally stronger than ordinary causal relationships. They constrain possible causal relationships, but only for properly picked-out subsets of the causal network and with sets of conditions in place that might, for instance, include constraints on what the physical laws must be for the explanations to hold. The way in which these subsets are picked out, however, means that once we have found them, we can give explanations about their causal structure that are distinctively mathematical, rather than simply causal.

Once again, distinctively mathematical explanations cannot be identified simply by having a mathematical form. There are many straightforwardly causal relationships that can be represented mathematically without thereby rendering them mathematical rather than causal. The key feature ought to

involve mathematical relationships as the, or part of the, connection between explanandum and explanans:

Distinctively mathematical explanations are such that, given some set of conditions A, then the connection between explanandum and explanans is mathematical in character, such that the resulting explanation involves a modal force stronger than that of causal generalizations.

This just is Lange's characterization, preceded by the stipulation that some set of conditions must hold; these conditions concern the nature of the relationship between explanans and explanandum. Instead of asserting that the mathematical modal character holds (call that B), it asserts a conditional: If A, then B. For the examples Lange identifies, the conditions hold vacuously, or are trivially fulfilled—no conditions must be met for twenty-three to not be evenly divisible by three. What about non-trivial conditions? One important way, perhaps one of the most central ways, to cash out what those conditions are that must be met for a distinctively mathematical explanation to result, concerns the notion of 'holding of'. When a mathematical equation or set of equations holds of the world in the right way, it is the mathematical relationship(s) doing the explanatory work, even though they are representing causal relationship(s).

What does it mean to hold of the world in different ways? Since I'll be using a model-based example in the next section, I will consider this question in terms of what it means for a model to hold, or fail to hold, of a particular system in the world. This is a rich topic for discussion, on which this article will only briefly touch. In general, it means that the model is applicable to the system, such that elements within the model provide a sufficiently veridical representation of parts of the system thus represented. It might also correctly describe at least some of the causal structure of the system (although, to be clear, it need not). The model holds of the system just in case the system is within the domain of systems to which the model can even be applied, and that application yields at least some degree of fit between the model as a representational device and the actual system in the world. Holding of can be a matter of degree: a model can hold to a greater or lesser extent, or it can hold of one system more than it holds of another, even while holding of both above some level. This is vague and broad, but requires as few philosophical assumptions as possible.

To say that a model holds of a particular system, then, is a situating per Salmon. Since the world is densely packed with causal relationships, this is a situating of the model in the causal network of the world. It lines up the model and a proper subset of the world, conceptually super-imposing the former over the latter to show how it fits there. To say that a model holds of some part of the network is not to make a causal claim. 'Holding of' has the wrong relata to

be causal, since one relatum is a chunk of the world, and one relatum is a representational device.

There are different ways in which a model might hold, or fail to hold, of any particular system. Some of the ways in which a model holds are not themselves explanatory—it is merely a representational relationship, with no added explanatory force of its own. But there are also times when it is explanatory simply to say that the model holds. Note that these are different ways in which it might hold, not merely different degrees to which it might hold. Some of the ways of holding of a system yield a causal version of the model. Explanations using the model that holds in this way will be causal explanations. But there are ways a model might hold of a system such that explanations using the model are distinctively mathematical, modally stronger than merely causal ones. This is a kind of situating of a model in the network of causal relations in the world that gives explanatory leverage on the system as an explanandum, but in a way that is not causal, even though it is because of the causal structure of the system that the model holds of it. This will be clearer with an example.

4 Lotka–Volterra: Same Model, Different Explanation Types

My discussion so far has been rather abstract. A simplified version of the Lotka–Volterra model serves as a useful illustration of how the same model can be deployed in causal and in non-causal explanations, of how to find distinctively mathematical explanations that only hold when given conditions are met, and of how it can be explanatory as opposed to merely representational to say that the model holds of a system. This example also demonstrates the general advantage of the narrow construal of causal explanation.

4.1 General biocide in the Lotka–Volterra model

The equations of the simplified version of the Lotka–Volterra model represent a predator and prey population over time as coupled harmonic oscillators:

$$\begin{aligned}dV/dt &= rV - (aV)P, \\dP/dt &= b(aV)P - mP,\end{aligned}$$

where V represents number of prey, P number of predators, r prey growth rate, and m predator death rate.

On the broad construal of causal explanation, the Lotka–Volterra (LV, henceforth) equations provide causal explanations. Consider a recent description of the LV mathematical model:

Predation is of great interest to ecologists because it often represents a *force* that *keeps* populations below their environments' *carrying capacities*. [Theoretical ecologists] construct models to study the factors that control the maximum *population size* as well as the phase, amplitude, and *frequency of oscillations* in the populations. (Weisberg [2013], p. 10; emphasis added)

Emphasised phrases and terms in this quotation involve thick causal terminology, such as 'keeps' and 'control', as well as terminology that is not causal per se, but is closely connected to or implicitly represents causal factors, such as the factors that control the relevant oscillation parameters. The LV equations represent many kinds of causal relationships. Individual prey and predator organisms are born, eat, sometimes get eaten, reproduce, and die. Some of these are explicitly represented, like the predator death rate, while some influence a variable but are not directly represented, such as prey eating. Populations sizes grow or diminish over time; they causally interact in a variety of ways, such as when a scarcity of prey making hunting harder. The size and rate at which each population grows is causally affected by the size and rate of which the other grows, or crashes. The predator–prey systems represented in the model are rich in causal structure.³

The LV model is a useful example because it is often utilized as a toy model to illustrate a simplified version of population dynamics, in full awareness that few if any actual populations precisely mirror these dynamics. Now consider what Weisberg and Reisman ([2008]) call the 'Volterra principle': any general biocide, something that indiscriminately kills both prey and predators, will result in an increase in the ratio of prey to predators. This follows purely from a consideration of the mathematics in terms of solutions to the equations for equilibrium values (in other words, not immediately upon introduction of the general biocide, but once the oscillations have stabilized with the new causal factor of the biocide), and for the stability of those equilibrium solutions to

³ There are extremely interesting questions about getting from a welter of causal relationships to the spare mathematical lines in the model, though, such that the model might not be directly representing anything causal. There are different ways of summing across causal histories, some of which may yield higher-level causal histories, and some of which may yield some other kind of relationship that is about causal relationships and is constituted by them, but is not thereby itself causal. This connects to debates about drift as a causal force, or fitness as causal versus mathematical. Another question raised by the LV model worth exploring is how, on the face of it, a causal–mechanical account and an interventionist account of causation seem to yield different answers as to whether a summed set of causal trajectories is itself causal. These questions are left to the side for this paper.

small perturbations in parameter values (Weisberg and Reisman [2008]). Consider:

p = average predator population size/average prey population size,

$$p = rb/m.$$

Recall, r is prey reproduction rate and m is predator death rate. In the presence of general biocide, r goes down (that is, prey are not reproducing as quickly) and m goes up (that is, the predator death rate increases). Thus, a general biocide must decrease p . Changes in r and m just are changes to average population sizes, so that the ratio of predators to prey increases or decreases. The effect of the general biocide on both populations need not be equal or balanced in any particular way. By definition, a general biocide will be such that it affects those two factors, regardless of the quantitative distribution of that effect.

Some sample values can be plugged into the equations to illustrate this point. Figure 1 shows a baseline for the two populations; R and m are the two parameters that will change. Now compare to Figure 2, which is the graph for a light general biocide. This involves changing the relevant parameters slightly, so the predator death rate is higher and the prey birth rate is lower. Note how the peaks for the prey population have gotten higher, while the peaks for the predator population stay very close to where they were before. This illustrates an increase in the ratio of prey to predators. The frequency of oscillations has also changed, but that does not affect the ratio. The effect is more pronounced with a massive biocide, in Figure 3. The prey population peaks are strikingly higher than in Figure 1 and the predator population peaks have dropped significantly.

We have something here that looks distinctively mathematical. The modal force of the ‘must’ by which a general biocide must increase the relative proportion of prey to predators is stronger than any causal relationship. It will hold of any possible causal relationship that instantiates those equations. But the key question is, if this is a distinctively mathematical explanation, of what is it an explanation?

The examples we’ve just seen in the figures involve toy values—they were made up to provide a baseline and to illustrate how changing the predator death rate and prey birth rate changed the relationship between the two average populations. There is nothing in the world that is being modelled by this toy example using these values. Yet there is something undeniably explanatory about comparing those three graphs in terms of understanding the changes to them resulting from a general biocide. If there are any systems in the world of which this toy model holds, then we already have in hand a powerful explanation that constrains the possible causal mechanisms in such a system: they would be unable to violate this principle.

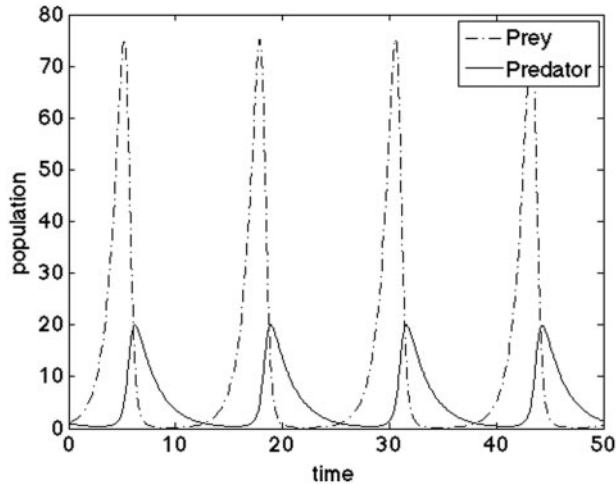


Figure 1. Without biocide; $r = 1$, $a = 0.2$, $c = 0.2$, $m = 0.5$.

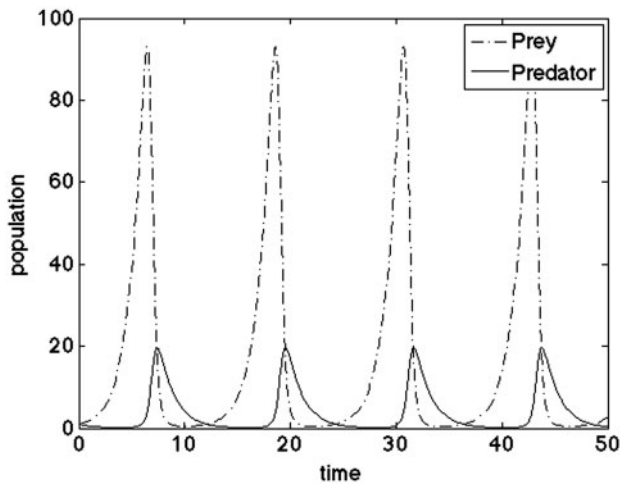


Figure 2. With biocide; $r = 0.8$, $a = 0.2$, $c = 0.2$, $m = 0.7$.

4.2 Two ways a model can hold, yielding causal versus mathematical explanations

This leads to the contrast I want to draw between two different ways in which the LV model can hold of a system. The first, which is causal on both the narrow and broad construal, involves starting with a particular system in the world, such a given population of moose and wolves, modeling it as accurately as we can, and seeing what comes out of it. By starting




Figure 3. Massive biocide; $r = 0.5$, $a = 0.2$, $c = 0.2$, $m = 1.5$.

with a particular system, modelling it, and finding that the LV equations hold of the system, we end up with a causal model of the system, and a merely representational relationship between the model and the system. There is no reason to think in this case that the resulting model would apply elsewhere, or to assume similar outcomes for another system with broadly similar mechanisms. If we find out that the causal mechanisms of this system are such that under some circumstances, the model fails to hold, we would have to develop a different model for this system, not choose a different system.

It could turn out that this model does not apply, or that we could change the dynamics of the system such that it no longer applies. This means that the explanation for why a general biocide in that population results in an increase in the proportion of prey to predators does not have the modal force of a mathematical explanation. It doesn't have to hold of the causal relationships when the model holds of the system in this way. They could be found, by the failure of the Volterra principle in this system, to violate the LV equations instead. Insofar as it does hold, it holds with the contingency associated with causal explanations, not with the stronger necessity of mathematical explanations. It is a situating in the network of causal relations that starts with a specific chunk of the network and then tries to find a model that holds of it well enough.

In contrast, if we start with the toy model in which we know with mathematical certainty that the Volterra principle holds, we can then go looking for some part of the world of which it might hold. This turns out to license a host

of additional explanatory resources, including ones that have the modal force of distinctively mathematical explanations. If the model holds of some chunk of the network of causal relations, that of which it holds must conform to the Volterra principle. We know, before ever finding such a system, that the causal structure of the system must yield that result. This is because we are picking out the system(s) in question because they conform to the LV equations.

This is an entirely different use of the model than when we start with a system and build a model for it. In this second approach, the model is already in hand and we go looking for parts of the world to which it fits—that is, the group of systems found to conform to the model must have specific features to which the systems fit. In the first case, in which we start with a system, it could turn out that the model developed to represent it fails. This is a representational failure, and means that conclusions drawn from the model needn't apply to the system itself. But when we start with the model and then pick out only those systems to which it applies, we don't get representational failures because, by definition, only those systems to which it actually applies were picked out. It failing to apply doesn't mean the conclusions drawn from the model don't fit the systems it models, it means that the system that we have picked out is not the right kind of system for that model.

In the first case, the chunk of the network of causal relations being considered is held fixed and if the equations do not apply, we need different equations. In the second, the equations are held fixed and if one part of the network of causal relations doesn't fit those equations, we move on to another part of the network. By contrasting these two cases, we can see the modified definition of the distinctively mathematical explanation at play. If the conditions are met—namely, if this model holds of a system in the world—then such a system must conform to the Volterra principle.

That the model holds already provides explanatory traction on that part of the world, in the way that had thus far been represented but not yet explained. This is a way of situating the model in the network of causal relations by selectively picking out pieces of that network such that the pieces of the network thus picked are governed by the mathematical relationships of the model in a way that is modally stronger than merely identifying instances of a particular causal relation. The model holds of that set of systems differently than it holds of a system that was found to follow the equations.

On the broad construal, both of these modelling scenarios involve causal explanation, because they are about a causal system. On the narrow construal, the first is a causal explanation involving situating the LV model in the network of causal relations. One could find that a particular set of causal relations in the system constitute a general biocide and that the system itself thus conforms to the Volterra principle.

But the second way of using the model is not providing a narrow causal explanation when we say that the Volterra principle holds, even though it is also a way of situating the model in the network of causal relations, and even though it is explanatory of the system in the world thus identified to say that the LV equations hold of it. The Volterra principle must hold, with a mathematical and not causal degree of necessity, in any system we identify to which this model applies.

When we start from a mathematical relationship within a model that is being considered separately from particular systems, it is explanatory of these systems to which it applies that the model holds of them. It is putting all such systems together into a special type, namely, that of the systems of which this model holds. It situates the model in the network of causal relations in the world such that we can then make further claims about any system of that type, since that type is defined as any system of which this particular model holds.

Construing causal explanation narrowly allows us to make this useful distinction between ways of building or deploying a model. It highlights the intriguing and distinctive character of the explanation that general bio-cide results in an increase in the proportion of prey to predators: it is not merely that there are causal structures described by these equations, but that any causal structures described by these equations must conform to this principle, no matter how differently they are implemented in terms of mechanistic detail. There are substantive conditions that must be met for the distinctively mathematical explanation to hold. It does not apply under any circumstances, which is why it requires amending Lange's original characterization; but once those conditions are met, it carries a distinctive modal necessity with it. We discover something about the world by finding that this model holds of a particular system, and this allows us to recognize two ways of making such a discovery: by starting with the system and finding that the model holds, or starting with the model and finding a part of the world of which it holds.

This also illustrates the usefulness of the grain-focusing metaphor for discussions about the role of models in explanation. Does the LV model provide mathematical or causal explanations? This is not yet a sufficiently well-defined question—it can be used for either, depending on the specific explanandum in question. The model itself is neither intrinsically causal nor non-causal. Applied in one way, to one part of the causal network, it yields mathematical explanations; applied a different way, to a different part of the network, it yields causal explanations.

5 Four Complementary Relationships between Mathematical and Causal Explanation

There are (at least) four different ways in which causal and distinctively mathematical explanations can complement one another in terms of filling out a richer explanatory picture of a target phenomenon.⁴ Some of these points also apply directly to, or can be extended to, other forms of non-causal explanation. I'll focus on the complementary roles for distinctively mathematical and causal explanations here, in keeping with the example in the previous section.

There is a weaker and a stronger version of the claim that these explanations complement rather than compete, and in what follows I argue for both. The weaker claim is that a more precise specification of a broad but fuzzy explanandum results in closely related but non-identically formulated explananda, each of which will thus involve a slightly different explanans. The resulting explanans fill out a richer picture of the original fuzzily formulated explananda as comprising these more specific explananda. The stronger claim is that there may be times when a single well-specified explanandum itself is amenable to both mathematical and causal explanation. This is an extremely interesting situation, where two distinct explanatory perspectives can be taken on a single explanandum.

5.1 Slight reformulations of explananda

The first complementary role for distinctively mathematical and causal explanation is simply that where slight reformulations of the explanandum pivot between a causal and mathematical explanation. In the LV example above, one can consider two different explananda: one is the behaviour of this system, to which the LV equations can be applied; another is any system to which the LV equations can be applied. Using the LV model, similar explanations can be given of behaviour in the specific system being modelled and in the class of systems to which the model applies. But they are not identical explanations, since they will involve slightly different explanans according to the slight differences in explananda. A particular population of moose and wolves on a single island in a wilderness area might be one of which the LV equations holds such that the changes over the past five years in their populations can be causally explained in terms of the model. Or their population changes have been tied together in the way that they have over the past five years can be explained in terms of showing that this island is an instance of the LV equations. Finding the precise boundaries across which slight reformulations of

⁴ These are not intended to be exhaustive or exclusive. My goal is to pick out some central avenues for complementary roles with sufficient richness of detail to enable identification of such instances.

explananda shift between causal and non-causal explanations gives us further insight into the overall phenomenon of which they are reformulations.

5.2 Causal distortion of idealized mathematical models

A second complementary role can be that of providing a causal explanation of why a particular non-causal explanation doesn't hold, or only holds to a low degree. For instance, we might find a particular set of predator–prey populations that initially seem like good candidates for modelling using the LV equations. However, it can turn out that this particular set of populations are not very well modelled by these equations, because there is some further causal factor(s) not included in the equations but which have a relevant effect on population dynamics. The way(s) in which the mathematics does not do an adequate job representing the actual system can itself be explained in terms of those additional causal factors and the way(s) in which they affect the solutions to the original equations. In the moose–wolf case, a particular island might have otherwise been well modelled by these equations, except for the presence of hunters that regularly cull a specific number of moose, which does not appear in the equations and is not tied to the population of moose. This can be especially enlightening if the LV equations can partially model the system: if they get some explanatory traction on what is going on, the way in which they fail to explain more than they do is itself a target for explanation in terms of additional causal relations. The equations might be good enough up to a limit in terms of numerical accuracy, or they might be good but only under a constrained set of conditions, such that causal details provide information about the system conditions in which they will break down, and so on.

5.3 Partial explanations requiring supplementation

A third complementary role involves distinctively mathematical explanations that provide a genuine explanation of some phenomenon, but only a partial one. This is different than failing to apply; a mathematical explanation might explain part of the explanandum, but require supplementation to adequately account for what happens. Often a number of different elements must be taken together to constitute a full explanation, and some of those elements may be distinctively mathematical, right alongside straightforwardly causal ones. For instance, a classic example of mathematical explanation is that of the prime year life cycle of the cicada (Baker [2005], [2009]; Saatsi [forthcoming]). It is extremely difficult for predators of cicadas to time their own reproductive cycles to that of their prey when the cicadas reproduce in prime year life cycles. The explanation involves pointing to the way in which primes are

divisible only by themselves and one; odd numbers like nine could still let a predator with a three-year reproductive cycle capitalize on the bounty of cicadas every three generations, whereas a thirteen-year cicada cycle does not allow that.

This is certainly some kind of explanation of why cicadas have a prime year life cycle, rather than even or non-prime odd year life cycles. But it is not a complete explanation. It has to be supplemented with a lot of other explanations in order to finish the explanatory task (Lehmann-Ziebarth *et al.* [2005]). This includes information such as the fact that cicadas, and not other insect species, have a safe place for larvae to wait for that many years without getting eaten and with sufficient nourishment that nymphs don't compete with one another. Bees could not implement a prime year life cycle like this, even if it would be equally advantageous for them to avoid predation. Other insects might have also benefited from such a life cycle and the right conditions for larvae to wait it out, but did not have the requisite genetic possibilities to actually evolve such a life cycle. As such, part of the total explanation of why cicadas have a prime year life cycle involves distinctive mathematical features of prime numbers, but a more complete explanation of that phenomenon must involve causal explanations as well.

5.4 Explanatory dimensionality

Finally, a fourth complementary role for distinctively mathematical and causal explanations comes closest to what one might have thought would have been alternative explanations of the very same explanandum. Rather than competing with one another, however, these explanations provide a powerful and underappreciated dimensionality to their explanandum. This fourth complementary role involves causal explanations that illustrate, instantiate, or fall under a distinctively mathematical explanation. A distinctively mathematical explanation provides limits on the space of possible causal structures that could be involved in a given explanandum: only those within the bounds of the mathematically possible could be realized, and any within the bounds delineated mathematically would suffice. Every causal system of the relevant sort must conform, in the sense of being locatable somewhere in that space of possibility. Such a mathematical explanation explains not merely particular instances of that kind of system, but also situates different instances in the same space of possible structure. Adding a causal explanation involves picking out some proper subset within the mathematically delineated bounds. It provides further explanation of some particular causal system to see how it fits in that space of possibility, how distinct but related systems fill out or fail to fill out that space of possibility, and to see how those systems evolve through the space of possible causal structures over time.

The distinctively mathematical explanations provide the underlying topography over which the actual causal systems are laid and across which they travel through time. Picking out one of the systems to which the distinctively mathematical explanation applies, it is explanatorily illuminating to see the way in which the system's causal structure, including mechanistic detail, conforms to or implements the spare equations of the model. The same phenomenon—the behaviour of those predator and prey populations in the system—is grouped as a single instance of two distinct types, such that distinct explanations are provided by situating the system as an instance of each type of explanation. In one case, it is picked out as an instance of the type of system to which the equations apply, such that its behaviour can be held up to that of other systems picked out by the same criteria, to see how it is similar and in what regards it differs. But it can also be picked out in terms of the detailed causal structure of that system, which may differ relevantly from the other systems, and which may group this system into a different type in terms of causal detail. It is the same token, but considered as instances of different types.

Picking out a part of the network of causal relations and zooming in to get more details about its causal structure, versus picking out a part of the network of causal relations as one instance of a type that is found elsewhere in the network of causal relations: this is a complementary way for two kinds of explanations to illuminate the same explanandum. Taking the causal and mathematical explanatory perspectives on the very same explanandum yields a powerful dimensionality that neither kind of explanation could provide alone.

This is not an exhaustive or exclusive list of the ways in which causal and distinctively mathematical explanations can complement one another. And again, 'non-causal explanation' is a heterogeneous category such that other non-causal explanations, besides the distinctively mathematical, can complement causal explanations in yet further ways. However, it should be clear that distinctively mathematical and causal explanations are not in explanatory competition for explananda; rather, when taken together, they fill out a richer view of their explananda.

6 Conclusion

Distinguishing causal explanation more sharply and precisely from other forms of explanation by restricting what counts as causal has the, perhaps unexpected, consequence of strengthening causal explanation. Such clarification avoids lumping together other kinds of explanation, such as constitutive or mathematical, with causal; it means more to provide a causal explanation, rather than reducing causal explanation to a synonym for explanation in

general. On a broad construal of causal explanation, arguably anything with empirical content will turn out to be causal, and labelling an explanation as 'causal' consequently comes to mean little or nothing about that explanation. On a narrow construal of causal explanation, though, adding the label 'causal' to an explanation adds substantive information about what is doing the explanatory work. The result is that causal explanations are a more homogenous category.

The difference between a causal and non-causal explanation often turns on precise ways of formulating the explanandum in question: one way will require an explanans that is causally upstream of the explanandum, while a slight shift in formulation will change the explanandum sufficiently that a nearby non-causal explanans pops into focus. Imprecise formulation of explananda reduces the explanatory resources at our disposal by flattening genuinely different explanatory relationships into one generic category.

Distinctively mathematical explanations are those where the modal force of the explanation is stronger than any causal explanation could provide—it constrains all possible causal structures. Drawing on Lange ([2013]), I've argued that we should amend the characterization of distinctively causal explanations. Some mathematical explanations hold anywhere, always, under all conditions, such as twenty-three not being evenly divisible by three. But some apply once certain conditions are met, such as that a given model (such as the LV model) holds of the system. Once those conditions are met, some explanations have the modal necessity associated with distinctively mathematical explanations, stronger than that of causal explanations. 'Holding of' is thus a relation that can, under some circumstances, pack explanatory punch separately from the causal relations involved.

By combining a narrow construal of causal explanation with a broader construal of distinctively mathematical explanation, we can see how there are (at least) two distinct ways in which a model might hold of a system. Considering the toy version of the LV equations highlights different ways of situating that same model against the background framework of causal relations. Each distinct way of situating the model will result in different explanatory resources that can be used for the system of which the model holds.

Acknowledgements

This work was supported in part by a grant from SSHRC of Canada. Much thanks to the audience at the PSA 2014 symposium, and to co-symposiasts Alexander Reutlinger, Marc Lange, Lawrence Shapiro, and Laura Ruetsche. This work has benefited from feedback from and discussion with Alexander Reutlinger, Chris Pincock, Alisa Bokulich, and three anonymous referees.

Thanks to Nic Fillion for creating of the biocide graphs. I am grateful for the opportunity to live and work on unceded Coast Salish territory.

*Department of Philosophy,
Simon Fraser University, Burnaby,
British Columbia, Canada
holly_andersen@sfu.ca*

References

- Andersen, H. [2014a]: ‘A Field Guide to Mechanisms: Part I’, *Philosophy Compass*, **9**, pp. 274–83.
- Andersen, H. [2014b]: ‘A Field Guide to Mechanisms: Part II’, *Philosophy Compass*, **9**, pp. 284–93.
- Baker, A. [2005]: ‘Are There Genuine Mathematical Explanations of Physical Phenomena?’, *Mind*, **114**, pp. 223–38.
- Baker, A. [2009]: ‘Mathematical Explanation in Science’, *British Journal for the Philosophy of Science*, **60**, pp. 611–33.
- Batterman, R. [2002]: *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, New York: Oxford University Press.
- Batterman, R. [2010]: ‘On the Explanatory Role of Mathematics in Empirical Science’, *British Journal for Philosophy of Science*, **61**, pp. 1–25.
- Batterman, R. and Rice, C. [2014]: ‘Minimal Model Explanation’, *Philosophy of Science*, **81**, pp. 349–76.
- Bokulich, A. [2008]: ‘Can Classical Structures Explain Quantum Phenomena?’, *British Journal for the Philosophy of Science*, **59**, pp. 217–35.
- Lange, M. [2013]: ‘What Makes an Explanation Distinctively Mathematical?’, *British Journal of the Philosophy of Science*, **64**, pp. 485–511.
- Lehmann-Ziebarth, N., Heideman, P. P., Shapiro, R. A., Stoddart, S. L., Hsiao, C. C. L., Stephenson, G. R., Milewski, P. A. and Ives, A. R. [2005]: ‘Evolution of Periodicity in Periodical Cicadas’, *Ecology*, **86**, pp. 3200–11.
- Pincock, C. [2012]: *Mathematical and Scientific Representation*, New York: Oxford University Press.
- Pincock, C. [2015]: ‘Abstract Explanations in Science’, *British Journal for the Philosophy of Science*, **66**, pp. 857–82.
- Reutlinger, A. [2014]: ‘Why Is There Universal Macro-behavior? Renormalization Group Explanation as Non-causal Explanation’, *Philosophy of Science*, **81**, pp. 1157–70.
- Saatsi, J. [2011]: ‘The Enhanced Indispensability Argument: Representational versus Explanatory Role of Mathematics in Science’, *British Journal for the Philosophy of Science*, **62**, pp. 143–54.
- Saatsi, J. [forthcoming]: ‘On the “Indispensable Explanatory Role” of Mathematics’, *Mind*, doi:10.1093/mind/fzv175.
- Skow, B. [2014]: ‘Are There Non-Causal Explanations (of Particular Events)?’, *British Journal for the Philosophy of Science*, **65**, pp. 445–67.

- Sterrett, S. [2002]: 'Physical Models And Fundamental Laws: Using One Piece of the World to Tell about Another', *Mind and Society*, **5**, pp. 51–66.
- Weisberg, M. [2013]: *Simulation and Similarity: Using Models to Understand the World*, New York: Oxford University Press.
- Weisberg, M. and Reisman, K. [2008]: 'The Robust Volterra Principle', *Philosophy of Science*, **75**, pp. 106–31.