

Proportionality, Determinate Intervention Effects, and High-Level Causation

Abstract: Stephen Yablo’s notion of proportionality, despite controversies surrounding it, has played a significant role in philosophical discussions of mental causation and of high-level causation more generally. In particular, it is invoked in James Woodward’s interventionist account of high-level causation and explanation, and is implicit in a novel approach to constructing variables for causal modeling in the machine learning literature, known as causal feature learning (CFL). In this article, we articulate an account of proportionality inspired by both Yablo’s account of proportionality and the CFL account of variable construction. The resulting account has at least three merits. First, it illuminates an important feature of the notion of proportionality, when it is adapted to a probabilistic and interventionist framework. The feature is that at the center of the notion of proportionality lies the concept of “determinate intervention effects.” Second, it makes manifest a virtue of (common types of) high-level causal/explanatory statements over low-level ones, when relevant intervention effects are determinate. Third, it overcomes a limitation of the CFL framework and thereby also addresses a challenge to interventionist accounts of high-level causation.

1. Introduction

To explain an effect of interest, different causal stories, some being more fine-grained than others, can usually be told. Consider Stephen Yablo (1992)'s example. Suppose a pigeon is sensitive to red—no matter what shade of red it is—and is thus far more likely to peck at red food than food with other colors, say, green. Now this pigeon is provided with some red food, where the shade of redness happens to be scarlet, together with some green food. Unsurprisingly, the pigeon pecks at the red (and scarlet) food. An explanation of the behavior can be told: the food's *being red* caused the pigeon to peck. However, it seems a finer-grained story can also be told: the food's *being scarlet* caused the pigeon to peck. Call the first a *higher-level causal statement* and the second a *lower-level causal statement*.¹

The question is which causal statement is better with respect to explaining the fact that the pigeon pecked at the food, given the assumption that the pigeon is only sensitive

¹ High-/low-level refers to the grain size of a cause (or a causal statement), namely, how fine- or coarse-grained a cause is. By contrast, macro-/micro-level has something to do with constitution, with the macro-level representing a system's overall property (or state) and the micro-level representing the working (or mechanism) underlying that system. Though different, however, they share two key features essential to our discussion of proportionality. First, in both cases, the low-/micro-level determines (or necessitates) the high-/macro-level but not vice versa. Second, a high-/macro-level property, event or state of affair can usually (though not necessarily) be realized by multiple low-/micro-level properties, events, or states of affair. Although proportionality is originally linked to the high-/low-level discourse, our discussions will also apply to the macro-/micro-level discourse. For these reasons, higher-/lower-level and macro-/micro-level will be used interchangeably throughout the essay.

to red—no matter what shade of red it is. Some philosophers claim that the higher-level one is better because it satisfies the requirement of *proportionality* (Yablo 1992; 1997; Woodward 2010; 2021a; Griffiths et al. 2015), while others—who doubt the plausibility of the preference—hold that the higher-level is no better than its lower-level rivals (e.g., Sober 1999; Shapiro and Sober 2012; Franklin-Hall 2016). Following Woodward (2010; 2021), we see proportionality primarily as a *criterion* (among others such as invariance, stability, etc.) concerning the choice of variables for causal analysis.² That is, in the context of modelling causal relations for the purpose of making predictions or giving causal explanations,³ proportionality requires that we choose a candidate cause-variable (or variables) that best or better fit the given effect of interest. However, we must bear in mind from the outset that there might be other criteria of causal explanation pulling in different directions, such that proportionality may sometimes be properly compromised for the sake of other desiderata.⁴

² Notice that this understanding differs from the metaphysical conception that takes proportionality to play a defining/constitutive role in characterizing causation. Yablo (1992; 1997) himself takes a metaphysical stance towards proportionality, for he distinguishes causation from causal relevance and causal sufficiency, and thinks that causation requires a cause to be proportional to its effect. For debates over this conception, see McGrath (1998), List and Menzies (2009), Dowe (2010), Zhong (2014; 2022), and McDonnell (2017).

³ Causal relations and causal explanations are of different kinds. We take causal relations to be something in the real world, which are represented by causal hypotheses/explanations. That said, we sometimes use these terms interchangeably for convenience, with or without the indication that a causal explanation/hypothesis is based on a corresponding causal relation.

⁴ A different way to understand this is that proportionality is a *ceteris paribus* desideratum, requiring that *other things being roughly equal*, we choose a proportional cause-variable relative to the given

We side with those who regard proportionality as a virtue.⁵ One of our purposes in this essay is to demonstrate and highlight an important reason for why it is such a virtue. Towards this goal, we will articulate an account of proportionality that generalizes Yablo’s account of proportionality to a probabilistic and interventionist setting, which is inspired by a novel approach to constructing variables for causal modeling in the machine learning literature, known as causal feature learning (CFL, Chalupka et al. 2014, 2016a, 2016b, 2017). Our account has at least three merits. First, it reveals an interesting connection between the notion of proportionality and the notion of determinate intervention effects in spite of ambiguous interventions (Spirtes and Scheines 2004). The notion of “determinate intervention effects” is actually a central component of the notion of proportionality when it is adapted to a probabilistic and interventionist framework. Second, our account has a simple consequence: when relevant intervention effects are determinate, a high-level causal/explanatory statement is strictly more *informative* than a low-level one, in that the former *entails* the latter but not vice versa. Since a proportional cause, in our account, is simply one at a highest level with determinate intervention effects, one vindication for proportionality readily follows: *other things being equal and given the effect to be explained*, we prefer proportional causal/explanatory statements because they are the most informative.

effect-variable.

⁵ We agree with many (e.g., Bontly 2005; McLaughlin 2007; Woodward 2010) that proportionality can be viewed as a *pragmatic virtue* of (or constraint on) causal explanation/statement, but “pragmatic” here is better understood as relating to “means/ends justification”, rather than relating to features associated with the “pragmatics” of language use, or to our computational/epistemic limitations (this is emphasized by Woodward (2021a, 362–65)); that is, we think proportionality is a virtue because a proportional causal explanation conduces better to some goals of explanation, e.g., informativeness, as we shall argue below.

Third, our account also amounts to a generalization of the CFL account of variable construction and relaxes its restriction on comparisons of variables that are coarser and finer partitions of *a common space*. As a result, we can easily address a challenge to Woodward’s (2010) earlier account of proportionality posed by Franklin-Hall (2016), in a simpler fashion than Woodward’s own response in his most recent account (2021a, 2021b). The challenge is that the consideration of proportionality cannot favor a high-level binary variable over a low-level binary variable, such as *red/non-red* versus *scarlet/cyan* in Yablo’s example. In our account, *red/non-red* rather than *scarlet/cyan* is unequivocally picked out as the proportional cause-variable for the effect-variable *pecking/non-pecking*. And again, we have a clear reason for this preference: the relevant causal statement in terms of the former is strictly more informative than that in terms of the latter.⁶

Two preliminary remarks are in order. First, like many participants in the debate, we embrace Woodward’s (2003) interventionist account of causation (and hope to improve his interventionist account of proportionality). According to this account, one variable *X* is a cause of or has a causal influence on another variable *Y* when there exists an intervention on *X* with respect to *Y* that can change *Y*’s value or probability distribution. The concept of an intervention can be roughly understood in the following way: “An intervention on *X* with respect to *Y* changes the value of *X* in such a way that if any change occurs in *Y*, it occurs only as a result of the change in the value of *X* and

⁶ Similar responses are given by Blanchard (2020) and in Woodward’s more recent work (Woodward 2021a, 2021b). As we explain in Section 5, our way of making the point is more general and rigorous, and stems from an account that is closer to the original notion of proportionality.

not from some other source” (Woodward 2003, 14).⁷ Second, in this essay we are concerned with “vertical” comparisons between candidate causal statements but not “horizontal” ones. They are vertical because the candidate causes in the alternative causal statements stand in logical, metaphysical (e.g., supervenience or realization), coarse-graining, determinable/determinate, among many others, relationships, so that one cause’s state can entail or necessitate the other cause’s state, e.g., an object’s being scarlet entails that it is red. This differs from horizontal comparisons, where the candidate causes in the alternative causal statements are *distinct* (Lewis 1986), meaning that one cause’s state does not logically or metaphysically constrain the other cause’s state, e.g., an object’s being scarlet and its being heavy.⁸

The rest of this essay will proceed as follows. Section 2 briefly reviews Yablo’s notion of proportionality, defined in terms of a notion of “screening off”. Section 3 introduces the CFL framework developed in the machine learning literature, analyzes its connections to Yablo’s account of proportionality, and develops an account that generalizes both. Based on this generalized account, we show in Section 4 that proportional causal/explanatory statements are the most informative among those that feature determinate intervention effects, and argue that this supplies a simple and clear rationale for sometimes favoring high-level causal/explanatory statements over low-

⁷ For a more comprehensive discussion of Woodward’s notion of intervention, see Woodward (2003, Chapter 3).

⁸ As pointed out by Potochnik (2010), in real scientific settings the comparisons between alternative causal statements can run both vertically and horizontally. However, for simplicity this essay will concentrate on the vertical cases.

level ones. Section 5 discusses some important related work and further clarifies a few relevant issues.

2. Yablo's notion of proportionality

Yablo (1992)'s pigeon, Sophie, has been a poster bird for the literature on proportionality. As mentioned previously, Sophie is only sensitive to red, namely, she is insensitive to colors other than red and appears to be unable to tell apart different shades of red. Hence, she is far more likely to peck at red food. So, when she is presented with some red food, which happens to be *scarlet*, and consequently pecks at the food, two stories can be told: that the food's *being red* caused her to peck, or that the food's *being scarlet* caused her to peck. Obviously, the first one is more coarse-grained than the second, for the property invoked in the first stands to that invoked in the second as a determinable stands to a determinate.

Given the assumption that she is only sensitive to red, Yablo (1992, 1997) holds that the more coarse-grained causal statement is *true* while the more fine-grained one is *false*. His reason is that the cause in the more coarse-grained causal statement is *proportional* to the effect, whereas the putative cause in the more fine-grained one is not. Using a notion of *screening off*, he defines *proportionality* as follows:

Given a pair of determinable and determinate C and C^* and a property E ,

“(1) *C screens off C** from *E* iff, had *C* occurred without *C**, *E* would still

have occurred.

(2) *C* is *required by E* iff none of its determinables screens it off, and *C* is *enough for E* iff

it screens off all of its determinates.

(3) *C* is *proportional to E* iff it is both required by and enough for *E*.” (adapted from Yablo

(1997, 266–67)).

With this definition, we can understand why Yablo thinks the more coarse-grained causal statement is true while the more fine-grained causal statement is false in the pigeon example. This is simply because the property $C=RED$, but not the property $C*=SCARLET$, is proportional to the property $E=PECKING$. Let us check the conditions. First, had *C* occurred without *C** (meaning that *RED* is present whereas *SCARLET* is absent), *E* would still have occurred (meaning that *PECKING* is present). Namely, even if the food had been red without its being scarlet, e.g., if the food had been crimson, Sophie would still have pecked. This goes in tune with the assumption that Sophie is only sensitive to red no matter what kind of red it is. Therefore, *C* screens off *C** from *E*, which means that *C** is not required by *E*.

Second, *C* is both required by *E* and enough for *E*. Consider the *required by* condition. Suppose we obtain a very coarse-grained property: $X=BRIGHT\ COLORS$. Now we ask if *X* can screen off *C* from *E*: had *X* occurred without *C*, would *E* still have occurred? The answer is negative, for Sophie is not sensitive to all bright colors and it might be the case that had the food been green (which is one way to realize the

antecedent “had X occurred without C ”), E would not have occurred. The same goes for all the other properties more coarse-grained than C , by the setup of the example. So, the *required by* condition is satisfied. Consider the *enough for* condition. Suppose we obtain a very fine-grained property: $Y=CRIMSON$. Can C screen off Y ? Namely, had C occurred without Y would E still have occurred? The answer is affirmative, for our assumption is just that Sophie is sensitive to red no matter what kind of red it is. By the same token, C can screen off all the other properties more fine-grained than C , and therefore the *enough for* condition is also satisfied. Hence, C is proportional to E .

Unlike Yablo, we do not regard proportionality as a defining condition for causation or a necessary condition for an acceptable causal explanation (cf. Woodward 2010; Griffiths et al. 2015). However, we do take it as a virtue (among other virtues that may sometimes pull in different directions) and one of our purposes is to present a simple and compelling reason for answering why it is a virtue. Our result, however, will be presented in a more general framework that accommodates probabilistic causation and is compatible with causal relata other than properties. The basic setup of this more general framework is presented by Chalupka et al. (2014, 2016a, 2016b, 2017), where they aim to construct the “right” macro-variables for causal modelling and inference and reach essentially the same idea as proportionality in their construction. It is instructive for our purpose to make explicit the parallel between Chalupka et al.’s proposal and Yablo’s, to which we now turn.

3. A generalized account of proportionality

Drawing on the theory of causal Bayesian networks (Kiiiveri et al. 1984; Spirtes et al. 2000; Pearl 2009) and computational mechanics (Shalizi 2001; Shalizi and Crutchfield 2001), Chalupka et al. (2014) develop their framework in the context of computer vision research, where a motivating task is to figure out what a macro visual cause (say, a red traffic light in a digital image) is for a given behavior (say, stopping of a self-driving car) from more micro-level data (say, pixel values). They dub the task *visual causal feature learning*, but the framework is by no means confined to computer vision, so we refer to it simply as *causal feature learning* (CFL). For the present purpose, the relevant part of their work is their conception of the “right” cause-variable given an effect-variable, where the candidate cause-variables are various partitions of an underlying state space.

Here is an illustration of the basic idea. In Yablo’s pigeon example, suppose the effect-variable of interest Y is the pigeon’s response to food (*pecking/non-pecking*), and imagine that the state space we are considering on the cause side is the hue space, in which each hue is represented in degrees ranging from 0 to 360. Each candidate cause-variable is a partition of this space, and the most fine-grained one takes each point in the space, i.e., a hue value, as a possible value. The question is which partition of the hue space is the right cause-variable for Y . In Chalupka et al.’s setup, it is assumed that for each point x in the state space, there is a well-defined *intervention effect* on Y , in the form of a probability distribution $p(Y | do(x))$, where $do(x)$ is Pearl’s (2009) celebrated notation that denotes an intervention to force the state x . Then the right partition is to

group together all and only states that have the same intervention effect on Y . For example, in Yablo's example, each hue value corresponding to redness has the same intervention effect (i.e., pecking with probability 1), and each other hue value has the same intervention effect (i.e., pecking with probability 0), so the right partition is to divide the space into two cells, corresponding to red and non-red, respectively.

More generally, the CFL account of the best cause-variable (CFL-BCV) can be summarized as follows:⁹

CFL-BCV: Given an effect-variable Y and a space \mathcal{X} , such that $p(Y=y | do(x))$ is defined for every $x \in \mathcal{X}$, the right or best cause-variable relative to Y is the partition of \mathcal{X} induced by the following equivalence relation between states in \mathcal{X} :

$$x_1 \sim x_2 \Leftrightarrow p(Y=y | do(x_1)) = p(Y=y | do(x_2)), \text{ for every value } y \text{ of } Y.$$

That is, in the variable or partition picked out by CFL-BCV, each value or cell is composed of, so to speak, micro-level states that have the same intervention effect on the given effect-variable. Therefore, to manipulate the variable to take a value has the same effect on Y *regardless of which micro-level state is realized by the intervention*. In other words, although the intervention is *ambiguous* with respect to the micro-states (because the intervention could force any of the micro-states to realize its target value of the target variable), it has an *unambiguous* or *determinate* effect on Y , in the sense

⁹ The CFL framework is also concerned with identifying the "right" effect variable from a fundamental space, which we leave aside as it is not relevant to our focus in this essay.

introduced and discussed by Spirtes and Scheines (2004). More precisely, in the CFL framework, a candidate cause-variable taking a value such as $C=c$ picks out a subset of the state space. For each point in the state space, it is assumed that there is a well-defined interventional probability distribution for Y , so in general $p(Y | do(C=c))$ is also understood to be a set: $\{p(Y | do(x)) | x \in C=c\}$. The intervention effect of $C=c$ on Y is said to be determinate just in case this set of interventional distributions is a singleton: the same distribution for Y results no matter which micro-state is realized by the intervention.

The notion of having a determinate intervention effect (despite the intervention being ambiguous) is actually the central notion here, because CFL-BCV can be equivalently formulated as stating that the best cause-variable is the coarsest partition such that each cell has a determinate intervention effect on the given effect-variable. We can hence think of the requirement of having determinate intervention effects as the first criterion in the CFL-BCV account. Among variables that meet this first criterion, we then have the second criterion which favors coarser partitions over finer partitions.

Yablo's account of proportionality, on the other hand, also picks out a "coarsest" property that is enough for the effect, because a proportional cause is by definition both enough for and required by the effect, and being required by the effect entails that no determinable of the property is enough for the effect. Therefore, we can also think of Yablo's account as requiring in the first place that a candidate cause should be enough for the effect, and among candidate causes that are enough for the effect, a coarser property (a determinable) is preferred to a finer one (a determinate).

So there is at least a formal analogy between the two accounts. More importantly, there is also a substantive connection because the concept of “having a determinate intervention effect” is a natural generalization of Yablo’s concept of “enough for” to the probabilistic and interventionist context. Before we proceed to establish this connection, a remark on causal relata is in order. In Yablo’s account the causal relata are properties,¹⁰ and a property’s occurrence amounts to that property’s presence or instantiation in an individual, whereas its non-occurrence amounts to its absence or non-instantiation. We can then represent a property with a binary variable, with one value denoting its occurrence and the other denoting its non-occurrence. Using variables is hence compatible with and more general than working with properties, as they can also be used to represent events, event types, states of affairs, or other candidate causal relata one might propose. Moreover, we can go beyond binary variables if needed, as is routinely done in causal modeling in the sciences.

Recall that Yablo’s concept of “enough for” is defined in terms of a concept of “screening off”, which is defined for properties that stand in a determinable/determinate relation. We thus define a corresponding relation for variables taking values.

Definition 1 (Value fine-graining/coarse-graining): Given variable-value pairs (C, c) and (C', c') , where c is a possible value of C and c' is a possible value of C' , (C', c') is

¹⁰ One anonymous referee says Yablo’s causal relata are, strictly speaking, events rather than properties. Yet, even this is true, it does not concern us here for, as we shall see below, we use variables to represent causal relata which is neutral with all these options.

said to be a fine-graining of (C, c) , and (C, c) a coarse-graining of (C', c') , if C' taking the value of c' necessitates C taking the value of c .

The notion of necessitation will be taken as a primitive in this essay. In the CFL setup with a state space, a value of a candidate cause-variable is simply a subset of the state space, and the coarse-graining/fine-graining relation can be simply identified with a superset/subset relation, which means that the necessitation in play is a sort of logical necessitation. Definition 1 is more general and can be combined with other notions of necessitation in different contexts, though our focus here is on the CFL framework. For convenience, we will also refer to a fine-graining as a *refinement*.

Recall that Yablo defines “ C screening off C^* from E ” by the following counterfactual: had C occurred without C^* , E would still have occurred (as it would have if C^* had). How to generalize this definition to a probabilistic context? As far as we can see, a most natural generalization is this: had C occurred without C^* , E would still have occurred *with the same probability* as it would have if C^* had. Cast in the interventionist language and notations, where the counterpart to the occurrence of a property is a variable taking a certain value and the counterpart to a counterfactual supposition is a hypothetical intervention, a formulation of this generalization is the following:

Definition 2 (Screening off): Let (C, c) be a variable-value pair, and (C', c') be a refinement of (C, c) , (C, c) is said to screen off (C', c') from (E, e) if

$$p(E=e \mid do(C'=c')) = p(E=e \mid do(C=c)).$$

Again, in the CFL framework, a candidate cause-variable taking a value such as $C=c$ picks out a subset of the state space, and $p(E=e \mid do(C=c))$ is also understood to be a set: $\{p(E=e \mid do(x)) \mid x \in C=c\}$. Then how to understand the equality invoked in Definition 2? There are at least two options here, and either will do for our purpose. We can either understand the equality as referring to equality of sets, or understand the equality as applicable only when the sets in question are singletons, i.e., when the interventional probabilities in question are determinate. It does not matter which option we take, because either way, the next definition—which is a straightforward adaptation of Yablo’s notion that C is enough for E if C screens off all of its determinates—will apply only when determinacy obtains.

Definition 3 (Enough for): (C, c) is enough for (E, e) if (C, c) screens off all of its refinements from (E, e) .

It is easy to see that in the CFL framework, (C, c) is enough for (E, e) just in case $p(E=e \mid do(C=c))$ is determinate, i.e., the set $\{p(E=e \mid do(x)) \mid x \in C=c\}$ is a singleton. For if (C, c) is enough for (E, e) , then by Definition 3, (C, c) screens off all of its refinements, including those refinements that correspond precisely to a single point in the state space. But for any single point x in the state space, $p(E=e \mid do(x))$ is supposed to be determinate in the CFL framework, so $p(E=e \mid do(C=c))$ must also be determinate in order to screen

off all of its refinements. Conversely, if $p(E=e \mid do(C=c))$ is determinate, then since in the CFL framework, fine-graining $C=c$ just picks out subsets of $C=c$, $p(E=e \mid do(C'=c'))$ will remain determinate and equal $p(E=e \mid do(C=c))$, for every (C', c') that is a refinement of (C, c) . This means that (C, c) is enough for (E, e) .

Therefore, the generalized concept of “enough for” given by Definition 3 is equivalent to the notion of having a determinate intervention effect, at least in the CFL framework. Hence our claim that the latter, though originally motivated by quite different considerations (Spirtes and Scheines 2004), is actually a natural generalization of Yablo’s concept of “enough for” to the probabilistic and interventionist context.

Definitions 1-3 are formulated in terms of variables taking values, as a generalization of Yablo’s talk of occurrence and non-occurrence of properties. It is straightforward to extend the notions to variables, by quantifying over the possible values. Specifically, we can define the notion of a cause-variable being enough for (or equivalently, having determinate intervention effects on) an effect-variable as follows:

Definition 4: (C, c) is said to be enough for a variable E if (C, c) is enough for (E, e) , for every possible value e of E . And a variable C is said to be enough for a variable E if for every possible value c of C , (C, c) is enough for E .

As we remarked earlier, the CFL-BCV account can be reformulated as picking out a coarsest cause-variable that is enough for the given effect-variable (in the sense of Definition 4), just as Yablo’s account of a proportional cause can be seen as picking out

a coarsest property that is enough for the given effect. Let us now add two definitions to make this reformulation explicit, which also serves to generalize the CFL-BCV account to overcome a significant limitation.

Definition 5 (Variable fine-graining/coarse-graining): A variable C_2 is said to be a fine-graining or refinement of a variable C_1 (and C_1 a coarse-graining of C_2) if every value of C_2 is a fine-graining of some value of C_1 , and every value of C_1 is a coarse-graining of some value of C_2 .

Definition 6 (Proportional cause-variable): C is said to be a proportional cause-variable for E if C is enough for E and no proper coarse-graining of C (i.e., coarse-graining of C that is not identical with C) is enough for E .

We label Definition 6 as defining a proportional cause-variable to highlight the obvious and close affinity to Yablo's notion of proportionality. It is straightforward to verify that in the CFL framework, there is a unique proportional cause-variable according to Definition 6 for a given effect-variable, which is precisely the cause-variable picked out by the original CFL-BCV.

This reformulation of the CFL-BCV account through Definitions 1-6 not only illuminates its connection to Yablo's account of proportionality, but also, thanks especially to Definition 5, overcomes a limitation of the original CFL-BCV account. The limitation is that the original account is restricted to candidate cause-variables that

are partitions of the same space, that is, partitions of a given state space. As we shall see in Section 5 below, this limitation has led some, e.g., Franklin-Hall (2016), to challenge Woodward (2010)'s earlier account of proportionality, which is automatically resolved in ours.

However, it is one thing to provide an account of proportionality that fits our intuitions, but quite another to provide a principled and compelling justification for the superiority of the proportional. In the next section, we show that our account implies a simple and compelling rationale for citing proportional causes or cause-variables in providing explanations.

4. Determinate difference-makers sink down

Definitions 1-3 in the previous section make one thing obvious: refinement of a value always preserves “enough for” or determinacy of intervention effects relative to an effect, thanks to the transitivity of the relation of necessitation. For example, in Yablo's pigeon example, the value *red* has a determinate effect on the effect *pecking*, so is any refinement of *red*, such as the value *scarlet* or the value *crimson*. Similarly, by Definitions 4-5, if a variable is enough for or has determinate intervention effects on an effect-variable, then any refinement of the variable remains so. For example, since *red/non-red* has determinate intervention effects on *pecking/non-pecking*, every refinement of *red/non-red*, including, for example, *scarlet/non-scarlet red/non-red* (which partitions the same space as *red/non-red* does), or *scarlet/cyan* (which partitions

a smaller space than *red/non-red* does), also has determinate intervention effects on *pecking/non-pecking*.

On the other hand, Definition 6 makes it clear that proportionality (of a cause-variable with respect to a given effect-variable) can be viewed as resulting from two desiderata. First, a cause-variable should have determinate effects on the effect-variable. Second, among cause-variables with determinate effects, the more coarse-grained (i.e., higher level), the better. The first amounts to “pulling down” the level while the second “pushing up”, together yielding a neither too high- nor too low-level cause-variable.

Our purpose in this section is to highlight a rationale for the second desideratum: more coarse-grained variables (with determinate effects) are better than more fine-grained variables (with determinate effects). The rationale is simply that causal statements in terms of the former are more *informative* than causal statements in terms of the latter—more informative in the sense that the former logically entail the latter but not vice versa. Similar ideas can be found in Blanchard (2020) and Woodward (2021a, 2021b), but our development of this rationale will be more general and rigorous.

Consider a common kind of statement in the causal modelling literature: variable X is a (determinate) cause of variable Y . Following a standard (albeit simplified) interventionist account, let us *stipulate* that this statement means that X is a determinate difference-maker of Y . In other words, X is a determinate cause of Y just in case (1) X has determinate intervention effects on Y , and (2) X is a difference maker for Y , i.e., there exist two values of X , $x_1 \neq x_2$, such that $p(Y | do(X=x_1)) \neq p(Y | do(X=x_2))$.

We make this stipulation to simplify the statements of the theorems below as well

as the accompanying discussions. We believe that this stipulation captures some common uses of such statements, but it is not our intention to claim that a statement of variable causation is always interpreted along this line. Our goal is to illustrate how using higher-level variables can be more informative than using lower-level variables in those causal/explanatory statements that include in their meanings a requirement of determinacy and a requirement of difference-making, perhaps among others. The stipulation made here can be seen as picking out the weakest interpretation of such statements.

With this stipulative definition of definite variable causation, the following theorem readily follows:

Theorem 1 (Determinate causes sink down): Given an effect-variable Y , if X is a determinate cause of Y , then every refinement of X is a determinate cause of Y .

Proof: The argument is straightforward. Suppose X is a determinate cause of Y . As mentioned at the beginning of this section, Definitions 1-3 together with the transitivity of necessitation guarantee that for every value of X , every refinement of that value has a determinate intervention effect on every value of Y . This, together with Definitions 4-5, entail that every variable that is a refinement of X also has determinate intervention effects on Y . Moreover, since X is a determinate cause of Y , there exist two values of X , $x_1 \neq x_2$, such that $p(Y | do(X=x_1)) \neq p(Y | do(X=x_2))$. Then by Definition 5, for every refinement X' of X , there is a value x_1' that refines x_1 , and there is a value x_2' that refines

x_2 . It follows that $p(Y \mid do(X=x_1)) = p(Y \mid do(X=x_2)) \neq p(Y \mid do(X=x_1)) = p(Y \mid do(X=x_2))$, which means that X remains a difference maker for Y . Q.E.D.

The idea of this theorem is not entirely new: Shapiro and Sober (2007) have already noted part of it; they also pointed out that the converse of this theorem is not true—we will explore these connections in Section 5.2.

Given an effect-variable Y , call a variable that is both a proportional cause-variable for Y and a cause of Y a *proportional cause* of Y . The following corollary of Theorem 1 is obvious. (The part in the parentheses follows from Definition 6, also illustrating that the converse of Theorem 1 is false.)

Corollary 1: Given an effect-variable Y , if X is a proportional cause of Y , then every refinement of X is a determinate cause of Y (and no proper coarse-graining of X is).

This result supplies a straightforward rationale for favoring certain types of causal/explanatory statements invoking proportional cause-variables over those invoking more fine-grained cause-variables. For example, consider a question raised by Franklin-Hall (2016): in Yablo's pigeon example, why is it better to assert that the variable $R=red/non-red$ is a cause of the variable $Y=pecking/non-pecking$ than to assert that $S=scarlet/cyan$ is a cause of Y ? Simply because the former is strictly more informative; it entails the latter but not vice versa. For such statements of (determinate) variable causation, the one with the proportional cause is the most informative true

statement, for all other true statements involving more fine-grained cause-variables are entailed (we will further discuss this question in Section 5.1).

Thus, the statement that R is a (determinate) cause of Y entails that S , as a refinement of R , is a (determinate) cause of Y , and that $C=crimson/amber$, as another refinement of R , is a (determinate) cause of Y , and so on. And this is not restricted to binary variables. For example, instead of fine-graining $R=red/non-red$ into a binary variable $S=scarlet/cyan$, we may also fine-grain R into a three-value variable $S'=scarlet/non-scarlet\ red/non-red$, or a four-value variable $S^*=scarlet/non-scarlet\ red/cyan/lime$, and so on. It follows from our theorem that the statement that R is a (determinate) cause of Y entails the statement that S' is a (determinate) cause of Y and the statement that S^* is a (determinate) cause of Y .

5. Further discussions

5.1. Franklin-Hall's challenge

As we mentioned in Section 3, the original CFL-BCV account has a limitation, namely, it is restricted to candidate cause-variables that are partitions of the same space, i.e., partitions of a given state space. As a consequence, the framework cannot accommodate the contrast between, e.g., binary variables at intuitively different levels, such as the variable $red/non-red$ versus the variable $scarlet/cyan$, because they are not partitions of the same space (the latter partitions a smaller space than the former does). This omission

is philosophically significant, because one of Franklin-Hall's (2016) criticisms of Woodward's (2010) earlier account of the potential superiority of high-level explanations rides precisely on such a contrast. Woodward's (2010) account of proportionality requires that a proportional cause-variable with respect to an effect-variable be one that, for one thing, "explicitly or implicitly conveys accurate information about the conditions under which alternative states of the effect will be realized", and for another, "conveys only such information—that is, the cause is not characterized in such a way that alternative states of it fail to be associated with changes in the effect" (298). In other words, the proportional cause-variable should contain *all and only* the causally relevant information with respect to accounting for the effect-variable.

Franklin-Hall (2016) presents the following argument to challenge Woodward's account. She invites us to consider a variable with only *scarlet* and *cyan* as its possible values. Intuitively this variable *scarlet/cyan* is clearly a lower-level variable compared to *red/non-red* and seems to be less proportional than the latter. Then, she evaluates whether this variable *scarlet/cyan* is proportional to the effect-variable *pecking/non-pecking* in terms of Woodward's notion of proportionality, and argues that *scarlet/cyan* also satisfies Woodward's conditions for proportionality, for *scarlet/cyan* "conveys accurate information about the conditions under which alternative states of the effect will be realized" and "conveys only such information—that is, the cause is not characterized in such a way that alternative states of it fail to be associated with changes in the effect" (Woodward 2010, 298). Therefore, concludes Franklin-Hall, Woodward's

account fails to entail that *scarlet/cyan* is not proportional (or less proportional than *red/non-red*), nor does it provide any other reason to favor the higher-level cause-variable *red/non-red* over *scarlet/cyan*. The original CFL-BCV account is also vulnerable to this criticism, as it only serves to pick out the “best” variable among partitions of the underlying state space, and is silent about *red/non-red* versus *scarlet/cyan*.

Since our account is not confined to variables partitioning the same space, it has a straightforward response to this criticism. Our definitions readily allow variables to partition a sub-space of the underlying state space, and hence are applicable to the contrast between *red/non-red* versus *scarlet/cyan*, for the latter can be seen as partitioning a sub-space of the space partitioned by the former. In particular, Definition 5 will rule that *scarlet/cyan* is a fine-graining of *red/non-red*, because the value *scarlet* is a fine-graining of the value *red* and the value *cyan* is a fine-graining of the value *non-red*, and therefore every value of *scarlet/cyan* is a fine-graining of some value of *red/non-red*, and every value of *red/non-red* is a coarse-graining of some value of *scarlet/cyan*, meeting Definition 5.¹¹ As a result, it is clear by our Definition 6 that in Yablo’s pigeon example, the variable *scarlet/cyan* is not proportional to the effect-variable *pecking/non-pecking*, because there exists a proper coarse-graining of *scarlet/cyan*, i.e., *red/non-red*, which is also enough for *pecking/non-pecking* (namely,

¹¹ Note that Definition 5 requires not only that every value of the more fine-grained variable refine some value of the more coarse-grained variable, but also that every value of the more coarse-grained variable be refined by some value of the more fine-grained variable, so that *scarlet/crimson* does not count as a fine-graining of *red/non-red*, because the value *non-red* is, so to speak, completely lost in *scarlet/crimson*.

red/non-red screens off *scarlet/cyan* from *pecking/non-pecking*). In addition, the general rationale for proportionality established in the previous section also applies: typical causal statements formulated in terms of the variable *red/non-red* entail those formulated in terms of *scarlet/cyan*.

In his critique of Franklin-Hall's arguments, Blanchard (2020) also appealed to the fact that using the variable *red/non-red* to formulate an explanation for an instance of pecking is more informative than using the variable *scarlet/cyan*. This is closely related to our result that the corresponding high-level causal statements entail the low-level ones. More recently, Woodward (2021a, 2021b) proposed a new account of proportionality, partly in response to Franklin-Hall's criticism. His new account highlights the informational aspect of the proportionality constraint and rules that in the pigeon example, the variable *red/non-red* is more proportional or "satisfies proportionality better" than the variable *scarlet/cyan* with respect to the effect variable *pecking/non-pecking* because the causal dependency claim associated with the latter fails to represent some existing dependence relations involving the effect variable that is accurately represented by the dependency claim associated with the former (whereas the latter dependency claim represents every dependency relation accurately represented by the former). Thus, on his most recent view, Woodward is likewise committed to linking the more proportional causal claim to the more informative causal claim. We applaud this development in Woodward's account as well as the similar insight in Blanchard's work. However, as we showed in previous sections, this informational virtue of proportionality naturally falls out of a principled and

straightforward generalization of Yablo's original notion; it does not require a significant and fairly complex reengineering of the concept as Woodward's definitions seem to have attempted (the spirit of which is also adopted by Blanchard). We therefore submit that the account developed in this paper is not only more general and rigorous than Blanchard's and Woodward's, but also reveals a simpler and more elegant way to express their shared insight concerning the informational value of proportionality.¹²

¹² Franklin-Hall's challenge has a second part, which asks that if we prefer red/non-red over scarlet/cyan with respect to pecking/non-pecking, why not prefer an even more abstract variable that encodes other potential difference-makers of pecking/non-pecking. Suppose, for example, that Sophie the pigeon would also peck if she were tickled. Then consider the variable RT, which takes value 1 if Sophie is presented with red food or is tickled, and value 0 otherwise. Shouldn't we favor RT over red/non-red, in the same way that we favor red/non-red over scarlet/cyan? But intuitively, suggests Franklin-Hall, RT is not a better choice. Blanchard's (2020) answer to this part of Franklin-Hall's challenge is interesting, but a key element in his response appears problematic. Blanchard contends that a key difference between red/non-red versus scarlet/cyan, and RT versus red/non-red is that in the former pair of variables, the more general red/non-red covers more possibilities that are conceptually incompatible with scarlet/cyan (because different shades of color are mutually exclusive), whereas in the latter pair, the additional possibilities incorporated by the more general RT are compatible with those encoded by red/non-red (because Sophie can be ticked while being presented red food). But this difference seems to be bogus. RT can be equivalently defined as taking value 1 if Sophie is presented with red food or is ticked *without red food*, and 0 otherwise. Seen this way, the additional possibilities are clearly also incompatible with the original ones. **In our account, the definition of proportionality is relative to a chosen space of possibilities (for framing a causal inference problem).** If the chosen space does include the possibility of being tickled, then we see no problem regarding RT as the more proportional variable than red/non-red. However, in the original pigeon example, the implicit assumption seems to be that all other potential causes of pecking are taken to be off and not considered, and hence a more appropriate choice of possibility space is to confine to the possible change of colors.

5.3. *The significance of determinateness and contrastive causal statements*

As we mentioned in Section 4, Shapiro and Sober (2007) have already touched part of the idea related to Theorem 1. They pointed out that difference-makers always sink down. That is, if some values of a macro-variable make a difference to an effect-variable, then there must exist some values of an underlying micro-variable that make a difference to the effect-variable. From this they concluded that macro-causation entails micro-causation, as they define macro-causation as difference-making without considering the issue of ambiguous interventions. We find Shapiro and Sober's discussions very insightful, but we also think that when it comes to macro-causation, it is important to keep alert to the matter of ambiguous interventions. When we make a difference-making statement such as $p(Y | do(X=x_1)) \neq p(Y | do(X=x_2))$, we need to make sure that the terms are well-defined so that the meaning of the disequality is clear. An obvious and simple option is to take such statements as requiring or presupposing that the relevant intervention effects are determinate. Of course this is not the only option, but it is at least a sensible choice in many contexts.¹³ In any case, the above theorem can be viewed as an extension of Shapiro and Sober's insight: even if we require both determinacy and difference-making in variable causation, macro-causation still entails micro-causation.

They also pointed out that the converse of the theorem is not true. However, we

¹³ For fans of the Lewisian (as opposed to the interventionist) semantics of counterfactuals, this requirement of determinacy is akin to the requirement of unanimity when the set of closest possible worlds is not a singleton.

think their argument for this point falls a little short. Their argument is based on an example which stipulates that an intervention to a macro-state amounts to a certain probability distribution over the underlying micro-states that can realize the macro-state. It is then possible that no *such* interventions on the macro-variable make a difference to the effect-variable whereas some interventions on the micro-variable do, because differences can average out.¹⁴ As we see it, this argument only shows that *some ways* to intervene on the macro-variable do not make a difference, but there are other ways to force a macro-state that deviate from the specified probability distribution over the underlying micro-states and will produce different effects. Therefore, it is not shown in the example that *no* intervention on the macro-variable can make a difference. In our view, a more promising example to show that difference-making does not percolate up would be one in which micro-states that do make a difference are subsumed into the same macro-state: an extreme, trivial example would be a “macro-level constant” that has only one value and subsumes all the micro-states. In the nontrivial case where the macro-level variable has at least two possible values, to show that difference-making

¹⁴ Their example is this: Suppose there is a macro-variable X that has values x_i and x_j , and there is a micro-variable V that has values v_{i1} , v_{i2} and v_j , and v_{i1} and v_{i2} are realizers of x_i and v_j is the realizer of x_j . X is a cause of another macro-variable Y , for an intervention on X changing its value from x_i to x_j can result in a change in Y . Further suppose that x_i is as likely to be in micro-state v_{i1} as it is in micro-state v_{i2} , and that half the micro-realizations of each x_i confer on Y a probability of 0.3 and half confer on it a probability of 0.9. Hence, with this particular setup, it is possible that “there are changes from one micro-state to another that will be associated with a change in the probability of Y , but an intervention that shifts the system from x_i to x_j will, on average, not affect the probability of Y , since $\Pr(Y \mid x_i) = \Pr(Y \mid x_j) = 0.6$ ” (Shapiro and Sober 2007, 257). As we see it, this shows only that *one* way to implement $do(x_i)$ does not make a difference as compared to $do(x_j)$. From this to draw the conclusion that the macro-variable X does not make a difference to Y seems problematic.

does not percolate up necessarily involves making clear what it is for interventions to make a difference when the intervention effects are indeterminate. The issue of ambiguous interventions looms large and cannot be simply ignored.

Our reason for rejecting the converse of Theorem 1 is different. In our account, the converse fails not because difference-making does not percolate up. In fact, if both a variable and a refinement of the variable have determinate intervention effects on the given effect-variable, difference-making does percolate up: the high-level variable will be a difference-maker for the effect-variable as long as the low-level variable is. Instead, the converse fails because determinacy of intervention effects (i.e., “enough for”) is not preserved under coarse-graining. It is obvious that even though a variable has determinate intervention effects on the effect-variable, a coarse-graining can easily break the determinacy. Therefore, high-level determinate causation entails low-level determinate causation, but not vice versa.

In another article, Shapiro and Sober (2012) argue against defences of proportionality. This is a little ironic from our perspective, since their earlier insight points at a clear justification for the preference of proportional causal statements. We hasten to add that statements of variable causation provide just one simple example of a type of causal/explanatory statement for which using a proportional cause-variable is most informative. To give another example, consider the *contrastive* form of causal statements, to which Shapiro and Sober (2012) appeal in one of their arguments against proportionality. The same argument we used to prove the theorem above can also be used to prove a similar theorem for causal/explanatory statements of an explicitly contrastive form: $X=x$ rather than $X=x'$ causes/explains $Y=y$ (rather than $Y=y'$). Suppose we understand such statements as saying that (1) X 's actual value is x and Y 's actual

value is y , (2) (X, x) and (X, x') both have determinate intervention effects on (Y, y) (and on (Y, y')), and (3) $p(Y=y \mid do(X=x)) > p(Y=y \mid do(X=x'))$ (and $p(Y=y' \mid do(X=x)) < p(Y=y' \mid do(X=x'))$). We have the following theorem:

Theorem 2 Suppose variable C 's actual value is c , D 's actual value is d , and E 's actual value is e , (D, d) is a fine-graining of (C, c) , and (D, d') is a fine-graining of (C, c') . Then the statement that $C=c$ rather than $C=c'$ causes/explains $E=e$ (rather than $E=e'$), entails the statement that $D=d$ rather than $D=d'$ causes/explains $E=e$ (rather than $E=e'$).

The argument for this theorem is extremely similar to that for Theorem 1 and will be left to readers. For such statements, therefore, there is again a simple and compelling rationale for citing proportional cause-variables, which can be relatively high-level. In particular, in the deterministic and binary setup illustrated by Yablo's pigeon example, the statement that the food being *red* (rather than *non-red*) causes/explains the pigeon's *pecking* (rather than *non-pecking*) is the most informative among true statements of this type; it entails such a statement featuring *scarlet* versus *cyan*, or *scarlet* versus *non-red*, or *bright red* versus *cyan*, but not vice versa.

5.4. Other informational accounts

Weslake (2010) criticized two kinds of informational account purporting to justify the preference of higher-level or non-fundamental explanations to lower-level or

fundamental explanations. One of them attempts to argue that a higher-level explanation may convey modal information that is not conveyed by a lower-level explanation, regarding what changes in the explanans would or would not have made a difference to the explanandum. The other seeks to show that some taxonomic information regarding how higher-level variables are related to lower-level variables may be omitted by lower-level variables. Our account is obviously unrelated to the second kind, and we agree with Weslake that the taxonomic information in question, though useful for other purposes, is not of explanatory relevance. Our account is also different from the first kind. We do not wish to claim that lower-level causal or explanatory statements fail to convey crucial information about difference-making or what-if-things-had-been-different questions. Theorem 2, for example, is concerned with causal/explanatory statements that make the difference-making information explicit at both levels. What we do claim is that difference-making at a higher level is more informative than difference-making at a lower level, because the former entails the latter, but not vice versa.

6. Concluding remarks

We have articulated an account of proportionality that adapts Yablo's original account to an interventionist framework and generalizes it to accommodate probabilistic causation. In developing the account, we drew an instructive parallel between a theory of causal variable construction in the machine learning literature and Yablo's account

of proportionality, revealing that the notion of “enough for” at the center of Yablo’s account corresponds to the notion of determinate intervention effects (which was originally introduced for very different purposes). Proportionality, we suggest, can be viewed as resulting from two desiderata, the first requiring determinacy of intervention effects and the second favoring the more coarse-grained or high-level over the less among those that meet the first desideratum.

By extending and improving on an insight from Shapiro and Sober (2007), we provided a justification for the second desideratum. When common types of causal/explanatory statements are understood to require or presuppose determinacy of intervention effects, those featuring high-level cause-states or cause-variables are more informative than those featuring low-level cause-states or cause-variables. And this justification is directly applicable to the example employed by Franklin-Hall to criticize Woodward’s defense of high-level explanation via proportionality; in the pigeon example, a causal/explanatory statement featuring *red/non-red* is more informative than that featuring *scarlet/cyan* (and the latter is unequivocally ruled to be not proportional in our account of proportionality.) Although similar responses were suggested recently by Blanchard (2020) and Woodward (2021a, 2021b), our way of developing this response was more general and rigorous, and followed more closely the original spirit of the notion of proportionality.

The first desideratum, on the other hand, rules against values of a variable that are too coarse-grained. For example, in the pigeon example, the variable *scarlet/non-scarlet* is often implicitly pitted against the variable *red/non-red*, but notice that these

variables do not stand in a fine-graining/coarse-graining relation. Although the value *scarlet* is a fine-graining of the value *red*, the value *non-scarlet* is not a fine-graining, but rather a coarse-graining, of the value *non-red*. According to the assumption made in the pigeon example, the value *non-scarlet* is too coarse-grained by the criterion of proportionality, because it has indeterminate intervention effects on (i.e., is not enough for) the effect-variable; for example, both *crimson* and *cyan* can realize *non-scarlet*, but their effects on *pecking* are different: $p(\textit{pecking} \mid \textit{do}(\textit{crimson})) \neq p(\textit{pecking} \mid \textit{do}(\textit{cyan}))$. Therefore, the variable *scarlet/non-scarlet* is not proportional not only because the value *scarlet* is too specific (because it is “screened off” by a coarser value *red*), but more importantly, because the value *non-scarlet* is too general, in that its intervention effect on the given effect-variable is not determinate.

As we remarked previously, in the case of ambiguous interventions, the requirement of determinate intervention effects is a natural and sensible choice to make sense of causal statements whose content concerns intervention effects. It is not unusual to regard determinateness or uniqueness as a condition for being well-defined. Moreover, an analogous requirement is part of Lewis’s (1973) influential definition of the truth condition of a counterfactual when the realization of the antecedent is ambiguous due to the presence of multiple closest antecedent-worlds. However, we do not wish to argue that this treatment is mandatory. As far as we can see, it is possible and potentially fruitful to adopt some framework of imprecise or indeterminate probabilities and define various causal notions in terms of possibly indeterminate intervention effects, so that causal statements can make sense without presupposing

determinacy of intervention effects. To further investigate this possibility, however, is beyond the scope of this essay.

The point we want to stress here is that the notion of proportionality amounts to prescribing that we work with cause-variables whose intervention effects on the given effect-variable are determinate (when they are available¹⁵), and that a high-level variable is preferred to a low-level one if they both have determinate intervention effects. We think our results in Section 4 provide a compelling reason for this preference, but we grant that it remains an open possibility to relax or even reject the requirement of having determinate intervention effects. We hasten to reiterate that this requirement is actually a force to resist going too high-level, so dropping this requirement may further liberate the use of high-level variables.

Our claim that a high-level causal statement may be more informative than a low-level causal statement may sound a little odd to some ears. When it comes to informativeness, the more familiar wisdom seems to be that a low-level causal statement may contain irrelevant information, rather than that it omits useful information. But this is precisely one of those situations in which more is less; adding irrelevant information in the description of a cause can eclipse what is relevant,

¹⁵ In this essay we follow the CFL framework to assume that there is a “fundamental” level or state space where the intervention effects of each state are well-defined or determinate, which then determine whether a higher-level variable has determinate intervention effects on the given effect-variable. With this structural assumption, there always exist some cause-variables that satisfy the requirement of having determinate intervention effects. In general, which level plays this role of a “fundamental” level may involve some convention or decision, depending on the context.

resulting in a logically weaker statement. In a way this is analogous to the fact that a material conditional gets weaker when its antecedent is strengthened. Finally, it is important to note that our claim about informativeness is relative to certain types of causal or explanatory statements. We believe it applies to a range of commonly used statements, but we by no means insist that it is universally applicable. Sometimes a low-level explanation is much more complex than the kind of explanatory statement considered here and may contain useful mechanistic information that cannot be described at the higher level. Obviously, our arguments are not meant to cover such cases.

References

- Blanchard, Thomas. 2020. "Explanatory Abstraction and the Goldilocks Problem: Interventionism Gets Things Just Right." *The British Journal for the Philosophy of Science* 71 (2): 633–63. <https://doi.org/10.1093/bjps/axy030>.
- Bontly, Thomas D. 2005. "Proportionality, Causation, and Exclusion." *Philosophia* 32 (1): 331–48. <https://doi.org/10.1007/BF02641629>.
- Chalupka, Krzysztof, Tobias Bischoff, Pietro Perona, and Frederick Eberhardt. 2016a. "Unsupervised Discovery of El Nino Using Causal Feature Learning on Microlevel Climate Data." *ArXiv Preprint ArXiv:1605.09370*.
- Chalupka, Krzysztof, Frederick Eberhardt, and Pietro Perona. 2017. "Causal Feature Learning: An Overview." *Behaviormetrika* 44 (1): 137–64.
- Chalupka, Krzysztof, Pietro Perona, and Frederick Eberhardt. 2014. "Visual Causal Feature Learning." *ArXiv Preprint ArXiv:1412.2309*.
- . 2016b. "Multi-Level Cause-Effect Systems." *Proceedings of Machine Learning Research* 51: 361–69.
- Dowe, Phil. 2010. "Proportionality and Omissions." *Analysis* 70 (3): 446–51. <https://doi.org/10.1093/analys/anq033>.
- Franklin-Hall, L. R. 2016. "High-Level Explanation and the Interventionist's 'Variables Problem.'" *The British Journal for the Philosophy of Science* 67 (2): 553–77. <https://doi.org/10.1093/bjps/axu040>.
- Griffiths, Paul E., Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunju Kim, and Rob

- Knight. 2015. "Measuring Causal Specificity." *Philosophy of Science* 82 (4): 529–55.
- Kiiveri, Harri, Terry P. Speed, and John B. Carlin. 1984. "Recursive Causal Models." *Journal of the Australian Mathematical Society* 36 (1): 30–52.
- Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell.
- . 1986. "Events." In *Philosophical Papers: Volume II*, 242–70. Oxford: Oxford University Press.
- List, Christian, and Peter Menzies. 2009. "Nonreductive Physicalism and the Limits of the Exclusion Principle." *The Journal of Philosophy* 106 (9): 475–502.
- McDonnell, Neil. 2017. "Causal Exclusion and the Limits of Proportionality." *Philosophical Studies* 174 (6): 1459–74. <https://doi.org/10.1007/s11098-016-0767-3>.
- McGrath, Matthew. 1998. "Proportionality and Mental Causation: A Fit?" *Philosophical Perspectives* 12: 167–76.
- McLaughlin, Brian P. 2007. "Mental Causation and Shoemaker-Realization." *Erkenntnis (1975-)* 67 (2): 149–72.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge: Cambridge University Press.
- Potochnik, Angela. 2010. "Levels of Explanation Reconceived." *Philosophy of Science* 77 (1): 59–72.
- Shalizi, Cosma Rohilla. 2001. "Causal Architecture, Complexity and Self-Organization in the Time Series and Cellular Automata." PhD Thesis, University of

Wisconsin–Madison.

Shalizi, Cosma Rohilla, and James P. Crutchfield. 2001. “Computational Mechanics: Pattern and Prediction, Structure and Simplicity.” *Journal of Statistical Physics* 104 (3–4): 817–79.

Shapiro, Larry, and Elliott Sober. 2007. “Epiphenomenalism – the Do’s and the Don’ts.” In *Thinking About Causes: From Greek Philosophy to Modern Physics*, by Gereon Wolters and Peter Machamer, 235–64. Pittsburgh: University of Pittsburgh Press.

———. 2012. “Against Proportionality.” *Analysis* 72 (1): 89–93.

Sober, Elliott. 1999. “The Multiple Realizability Argument against Reductionism.” *Philosophy of Science* 66 (4): 542–64.

Spirtes, Peter, Clark N. Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. Cambridge: MIT press.

Spirtes, Peter, and Richard Scheines. 2004. “Causal Inference of Ambiguous Manipulations.” *Philosophy of Science* 71 (5): 833–45.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

———. 2010. “Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation.” *Biology & Philosophy* 25 (3): 287–318.

———. 2021a. *Causation with a Human Face*. Oxford: Oxford University Press.

———. 2021b. “Explanatory Autonomy: The Role of Proportionality, Stability, and Conditional Irrelevance.” *Synthese* 198 (1): 237–65.

<https://doi.org/10.1007/s11229-018-01998-6>.

Yablo, Stephen. 1992. "Mental Causation." *The Philosophical Review* 101 (2): 245–80.

<https://doi.org/10.2307/2185535>.

———. 1997. "Wide Causation." *Philosophical Perspectives* 11: 251–81.

Zhong, Lei. 2014. "Sophisticated Exclusion and Sophisticated Causation." *The Journal of Philosophy* 111 (7): 341–60.

———. 2022. "Why Causation is Biconditional but not Proportional," *American Philosophical Quarterly*, 59 (3): 263–271.