Philosophy of Science (2024), 1–8
doi:PSA 2024 conference version

RESEARCH ARTICLE

Proportionality and the Effect-of-Cause Question

Frederick Eberhardt[1], Wei Fang[2] & Jiji Zhang[3]

[authors listed in alphabetical order]
[1]Division of the Humanities and Social Sciences, Caltech, Pasadena CA 91125, USA
[2] Research Center for Philosophy of Science and Technology, Shanxi University, Taiyuan, Shanxi, China
[3] Department of Philosophy, The Chinese University of Hong Kong, Hong Kong
*Corresponding author. Email: jijizhang@cuhk.edu.hk

(Received xx xxx xxxx)

Abstract
One of the most influential accounts of high-level causation appeals to a notion of proportionality, which aims to identify the cause (variable) at an appropriate level given an effect (variable) of interest. Here we ask the dual question: Given a cause (variable) of interest, what is an appropriate level to model the effect? We argue that this question is not yet settled by the proportionality account and develop a natural counterpart to the proportionality account for this question. We also discuss several challenges for the effect-of-cause question that do not have an analog in the cause-of-effect question.

1. Introduction

Providing an account of the causal relations within a system is widely seen as a fundamental goal of research across the natural and social sciences, as it can provide the basis for prediction, intervention and counterfactual analysis. However, with each science developing its own causal model of its phenomena, the question inevitably arises whether there is a "right" or privileged level of causal analysis. In philosophy, this issue has long been prominent in the debates on mental causation and the autonomy of the special sciences [Kim, 1989, 2000, 2005, McGrath, 1998, McLaughlin, 2007, Shapiro, 2010, Franklin-Hall, 2016, McDonnell, 2017, Woodward, 2018, Vaassen, 2022] Now, with large high-resolution data sets becoming available in an increasing number of scientific domains, the philosophical question studied abstractly in the context of mental or supervenient causation acquires practical scientific relevance. For example, in neuroscience, the nervous systems of different animals are now measured at a variety of different scales, from recordings of almost every single neuron using lightsheet microscopy (e.g. in zebrafish [Ahrens et al., 2013]), to functional magnetic resonance imaging (fMRI) of the whole human brain at a resolution of $1mm^3$ voxels every $0.7s$ [Van Essen et al., 2012]. In each case a major goal is to understand how the operations of the nervous system lead to behavior. But even with vast amounts of high-quality data at the neural level, it remains unclear what the appropriate level of spatial (and temporal) aggregation is to explain the relation between brain activity and output behavior.

Such questions dovetail with discussions in the causal inference literature about how to think about causal variables. Most discussions on this topic have explicitly or implicitly focused on the following question: Given an effect of interest (or phenomenon to be explained), at what level of granularity should we consider causes of the effect? In the philosophical literature, Yablo [1992]'s influential theory of proportionality is a sample answer to such a question, which has stimulated various related accounts about the right level of cause given an effect of interest [Woodward, 2010, 2021, Griffiths et al., 2015, Pocheville et al., 2017, Bourrat, 2018].

In this paper we consider the dual question: Given a cause of interest, at what level of granularity should we model effects of the cause? Call this kind of question the effect-of-cause (EOC) question, and the previous one the cause-of-effect (COE) question. The EOC question is apparently much less discussed than the COE question. This disparity is, we suspect, due at least to two reasons. First, causal questions and explanatory questions are often run together. In explanatory terms, it may sound more natural to ask what explains a given explanandum than to ask what a given explanans explains. However, the apparent unnaturalness does not necessarily reflect any defect in the latter explanatory question. It seems clear that at least in causal terms both questions are meaningful and important. For example, in addition to determining the efficacy of treatment, medical trials attempt to demarcate the boundaries of the treatment effect by identifying side-effects. And both in mundane everyday circumstances and in scientific inquiry it is common to ask about the consequences of a particular intervention without restricting the question to consequences for a predefined target.

Second, it may seem plausible to think that EOC and COE are just two ways of asking the same question, and that once the COE question is properly answered, the EOC question is automatically addressed. One of our aims in this paper is to argue that although the two questions are obviously connected, EOC is not simply reducible to COE. In particular, the proportionality account for the COE question does not settle the EOC question, though there is a natural counterpart to the proportionality account for the EOC question. Moreover, we will highlight some challenges for the EOC question that do not arise for the COE question.

The rest of the paper will proceed as follows. In Section 2, we review Yablo's theory of proportionality and explain why the EOC question is not simply answered by requiring that given a cause of interest C, the effect should be at a level of granularity to which C is proportional. Then, in Section 3, we describe a generalization of Yablo's account to address the COE question in a probabilistic and interventionist setting, and propose a dual account for the EOC question. While the account for COE requires the cause to be as general as possible subject to a constraint, the dual account for EOC requires the effect to be as specific as possible given the same constraint. In Section 4 we highlight some extra complications for the EOC question that are not present for the COE question. Some of these complications provide reasons to resist going as specific as the dual account implies.

2. Proportionality and the EOC question

Our point of departure is Yablo [1992, 1997]'s influential account to address a version of the COE question. Consider his oft-cited example: A pigeon has been trained to peck at red objects, regardless of the specific shades of red. The pigeon is then regularly given scarlet food, and she unfailingly pecks at the food. Suppose the pigeon pecking at the food is the effect (type) we are interested in, and consider the COE question: What is the cause of this effect? Many have the intuition that answering "the food being scarlet" is either false, or at least worse or less felicitous than answering "the food being red".

Yablo accounts for this intuition by his theory of proportionality. According to Yablo's theory, "the food being scarlet" is too specific or fine-grained for the given effect, because it is screened off from the effect by a more general or coarse-grained candidate, i.e., "the food being red". Yablo defines screening-off in terms of a counterfactual conditional. Suppose C_2 is a specification or fine-graining of C_1 in the sense that the instantiation of C_2 necessitates the instantiation of C_1.[1] Then C_1 screens off C_2 from E just in case if C_1 were instantiated without C_2, E would still be instantiated. In the above example, if the food were red without being scarlet, the pigeon would still peck at the food. This shows that "being scarlet" is too specific for the effect. By contrast, "being red" is just right for the given effect in the pigeon example. In Yablo's terms, it is both *required by* and *enough for*, and hence *proportional to*, the effect. It is required by the effect in the sense that no more coarse-grained candidate cause screens it off from the effect. It is enough for the effect in the sense that it screens off all more fine-grained candidate causes from the effect.

In Yablo's view, therefore, proportionality is an answer to the COE question. One matter of debate concerns the nature of this answer: Does it pinpoint a semantic condition for statements of causation, or does it highlight a non-semantic criterion for selecting a cause that is most suitable for some explanatory purposes? Yablo maintains that proportionality is a semantic condition for statements of causation. A statement "C causes E" is false if C is not proportional to E. Thus, in the previous example, it is false that the food being scarlet causes the pigeon pecking at the food (for a similar view, see e.g., List and Menzies [2009], Zhong [2014]). Other authors (e.g., Woodward [2010, 2021], Griffiths et al. [2015]) recognize proportionality as a virtue but not as an essential condition in the definition of causation. According to these authors, it is not false to say that the food being scarlet causes the pigeon's pecking, but this statement is in some way inferior to the statement that the food being red causes the pigeon's pecking. In other words, "being red" is a better answer than "being scarlet" to the COE question in the pigeon example, but the comparative advantage is not what truth enjoys over falsehood.

We side with the latter group and have a novel reason for this preference. The reason has to do with the focus of this paper, the EOC question. In the pigeon example, suppose we are interested in "the food being scarlet" as a cause, and ask what its effect is. Should the answer be "none", because it is not proportional to "the pigeon's pecking" (or any more specific kind of pecking, as can be suitably built into the example)? We think not. In our view, it is more reasonable to say that the pigeon's pecking is an effect of the given cause, even though the given cause is not proportional to this effect. After all, we are interested in the EOC question regarding "being scarlet" presumably because we

[1] Yablo (1992, 1997) takes C_1 and C_2 to be properties, and the coarse-graining/fine-graining relation to be the determinable/determinate relation. We prefer to be uncommitted about the causal relata, and will use the more generic terminology(coarse-graining/fine-graining, general/specific, etc.) to refer to different levels of granularity.

know how to bring it about and we care about what we can manipulate through it. The answer "pecking" is clearly relevant to this interest and better than the answer "none".

Therefore, in our view, pecking is an effect of the food being scarlet, even though being scarlet is not the proportional cause of pecking. More important for our present purpose, the above consideration suggests that the EOC question is not automatically settled just because the COE problem is solved. In particular, assuming proportionality is accepted as the answer to the COE question, we cannot in general answer the EOC question by saying that given a cause of interest, the right or best level of granularity for the effect is that to which the given cause is proportional, because the given cause may not be proportional to any effect, as illustrated by our version of the pigeon example. That said, the EOC question is of course closely related to the COE question. One of our purposes in this paper is to develop an answer to the EOC question that is a natural counterpart to the proportionality account of COE, in a probabilistic and interventionist framework that focuses on causal relations between variables.

The main spirit of the account can be illustrated informally in Yablo's framework and using his terminology. Recall that Yablo defines proportionality in terms of a notion of "enough for" and a notion of "required by". A reformulation of this account [Chalupka et al., 2015, Shapiro and Sober, 2012] keeps the notion of "enough for" and makes it a criterion of admissibility: Only a candidate cause that is enough for the given effect is admissible. The notion of "required by", on the other hand, is replaced by a preference for the more general over the more specific. That is, between admissible candidate causes, the higher-level one is preferred to the lower-level one. (Example: both "being red" and "being scarlet" are enough for the pigeon's pecking, and "being red" is preferred to "being scarlet".) The proportional cause is simply the maximum in this preference order, the most general candidate that is enough for the given effect. And this is taken to be the best level of granularity for the cause given an effect.

From this perspective, it is easy to construct a dual account for the EOC question. Again, the notion of "enough for" defines admissibility: Given a cause of interest, only admit effects for which the cause is enough. Then, between admissible candidate effects, the more specific, lower-level one is preferred to the more general, higher-level one. The best level of granularity for the effect given a cause is the most specific candidate for which the given cause is enough (if there is such a candidate). For instance, consider, in the pigeon example, a more general or coarse-grained candidate effect, say, "pigeon touching the food (in any way)" [Zhong, 2014]. As we see the matter, it is also true that the pigeon touching the food is an effect of the food being scarlet. However, by the dual account for EOC, "pecking" is preferred to "touching" in answering the EOC question, because the given cause is enough for both but the former is more specific.

This outline is only a heuristic guide to the account we will discuss next, as we are concerned in this paper with variable selection for causal modelling (rather than with causation between values of variables that represent properties or events). But the main idea will be similar: given a certain admissibility constraint, generality is preferred in answering the COE question and specificity is preferred in answering the EOC question.

3. Proportionality for cause variable selection and a dual account for EOC

Following Chalupka et al. [2015]'s setup, we assume a "fundamental" space $\mathscr{X}$ for cause variables and $\mathscr{Y}$ for effect variables. For simplicity, assume these spaces are finite. For each $x \in \mathscr{X}$, there is an interventional probability measure over $\mathscr{Y}$, denoted as $P(\mathscr{Y}|do(x))$, using the celebrated do-operator [Pearl, 2000]. More generally, for any partition E of $\mathscr{Y}$, we write $P(E|do(x))$ as the interventional probability distribution over the cells of E that is entailed by $P(\mathscr{Y}|do(x))$. For any subset $\mathbf{x} \subseteq \mathscr{X}$, we use $P(\mathscr{Y}|do(\mathbf{x}))$ to denote a set: $\{P(\mathscr{Y}|do(x))|x \in \mathbf{x}\}$. Similarly, $P(E|do(\mathbf{x}))$ is a set of interventional probability distributions over the cells of E. We call $P(E|do(\mathbf{x}))$ *definite* if it is a singleton and we call any partition C of $\mathscr{X}$ a definite cause variable given E if for each of its cells c, $P(E|do(c))$ is definite; otherwise C is an ambiguous cause variable given E.[2] Consider two types of question:

- Cause of Effect (COE) question: Suppose we are given an effect variable E, a partition of $\mathscr{Y}$. What cause variable (partition of $\mathscr{X}$) should we use?
- Effect of Cause (EOC) question: Suppose we are given a cause variable C, a partition of $\mathscr{X}$. What effect variable (partition of $\mathscr{Y}$) should we use?

For the purpose of this paper, we simply assume that the COE question is settled by Chalupka et al. [2015]'s proposal[3]: The "right" cause variable is the coarsest partition of $\mathscr{X}$ whose values have definite effects on E. In the setup we consider here, this partition is unique, in which each cell of the cause partition C is an equivalence class for the equivalence relation:[4]

$$x_1 \sim x_2 \quad \iff \quad P(E|do(x_1)) = P(E|do(x_2))$$

Here is another way of thinking of this account: The property of having a definite effect on E is taken as the first desideratum for an appropriate cause variable. Only variables with definite effects on E are admissible cause variables. Given an effect variable E, the property of having definite effects on E is preserved downwards: if a variable C has definite effects on E, then every refinement of the partition C also has definite effects on E. The second desideratum then dictates that between two cause variables that both have definite effects on E, the more coarse-grained, the better. Since there is a unique coarsest cause variable that has definite effects on E, these two desiderata pick out that variable as the best or right cause variable.

This theory is akin to Yablo's proportionality account. Note that a definite cause variable C for the given effect E screens off any more fine-grained variable R from E in the following sense: $P(E|do(r)) = P(E|do(c \wedge \neg r))$, for every value c of C and every value r of R that is a fine-graining (i.e., a subset) of c. That is, to borrow Yablo's expression (see Section 2), if c were instantiated without r, the probability distribution of E would still remain the same (as if r were instantiated). Thus, "having definite effects" is an adaptation of Yablo's notion of "enough for" to variables in the probabilistic setting. The resulting state space of the cause variables describes candidate interventions whose effect is insensitive to the microlevel instantiation of the intervention. Consequently, the unique coarsest cause variable that has definite effects on E just corresponds to the *proportional* cause variable for E.

[2] Spirtes and Scheines [2004] introduced the notion of ambiguous manipulations that we adopt here and we use the term "definite" to describe their "unambiguous" manipulations.

[3] For alternative approaches, see e.g. Hoel [2017].

[4] By "$P(E|do(x_1)) = P(E|do(x_2))$" we mean that two probability distributions are the same, i.e., $P(e|do(x_1)) = P(e|do(x_2))$, for every value e of E.

The EOC question now reverses the query: What is the effect variable corresponding to a given cause variable C? Given that the cause variable (some partition of $\mathscr{X}$) is fixed, we first note that there may be no effect variable (partition of $\mathscr{Y}$) to which the cause variable is proportional. An extreme kind of example is when C does not have definite effects on any partition of $\mathscr{Y}$ except for the trivial one that merges all of $\mathscr{Y}$ into a single state.[5] In such a case, C can at best be a proportional cause variable to the trivial partition of $\mathscr{Y}$. But unless C itself is also the trivial partition of $\mathscr{X}$, it will not be the coarsest partition to have definite effects on the trivial partition of $\mathscr{Y}$, and hence will not be a proportional cause variable to any effect variable. More generally, the given cause variable may have definite effects on many partitions of $\mathscr{Y}$, but nonetheless is not the proportional cause variable with respect to any of those partitions.[6] This highlights why the EOC question is not automatically addressed by the above answer to the COE question: We cannot simply select an effect partition such that the cause C is proportional to it, since there may not be such an effect partition.

However, the (reformulated) proportionality account for the COE question naturally suggests a dual account for the EOC question. The dual account also consists of two principles, one about admissibility and the other a preference among admissible effect variables. The first principle still takes the absence of ambiguity of intervention effects to be the desideratum. That is, given a cause variable C, we should consider only effect variables on which C has definite effects. Note that the property of being influenced by C in a definite manner is preserved upwards: if C has definite effects on a variable E, then C has definite effects on every coarsening of E, simply because the distribution of the coarsened variable is determined by that of E. (Recall the dual fact in the COE case: The property of having definite effects on a given effect variable is preserved downwards.) This simple observation suggests the following principle of preference: between two effect variables on both of which C has definite effects, the more fine-grained, the better.

This account for EOC and the proportionality account for COE enjoys an elegant duality. Both accounts take the absence of ambiguity of intervention effects as a criterion of admissibility. Under this constraint, the account for COE prefers more general or coarse-grained cause variables, and the account for EOC prefers more specific or fine-grained effect variables. Those who are attracted to the proportionality account for the COE question will probably find this account for the EOC question a natural counterpart to work with. Moreover, like the proportionality account for COE, the dual account provides a plausible explanation of some intuitions about appropriate choices of variables in various domains. For example, the effect of a mental cause tends to be mental or behavioral rather than more fundamentally physical, as a mental cause will typically have ambiguous effects on too fine-grained physical variables. In contrast, the effect of a (more fundamental) physical cause is typically not mental as there are more specific effect variables on which the cause variable has definite effects. That is, this account will naturally tend to retain a similar granularity between cause and effect, while keeping the question of the appropriate granularity of the effect subject to empirical scrutiny.

[5] For the simplest example, suppose $\mathscr{Y}$ has only two atoms y_1 and y_2, and C has ambiguous effects on the distribution over y_1 and y_2. Then y_1 and y_2 must be merged into a single state for C to have definite effects. Such a degenerate situation can also easily obtain when the given cause is "gerrymandered".

[6] For instance, suppose C is the ternary variable in the pigeon example whose possible values are: *scarlet*, *non-scarlet red*, or *non-red*. Then it has definite effects on the binary behavioral variable *pecking/non-pecking* but is not proportional to it, for a coarser variable *red/non-red* also has definite effects.

4. Challenges for the duality of COE and EOC

Despite the appealing duality between COE and EOC, there are important differences between the two sides. We will briefly note three. First, unlike for cause variables, the two desiderata for effect variables (definiteness and preference for specificity) do not in general yield a uniquely best effect variable. For example, it is easy to construct a case such that one (and only one) value of a cause variable C, say, $C=1$, has an ambiguous effect on a variable E with four values $\{e_1, e_2, e_3, e_4\}$, such that $P(E|do(C=1))$ contains two distributions over the four states: $\{(0.2, 0.3, 0.3, 0.2), (0.3, 0.2, 0.2, 0.3)\}$. Then the coarsening of E into $E_A : \{e_{1\vee 2}, e_{3\vee 4}\}$ will make the effects of $C=1$ definite. But so will the coarsening into $E_B : \{e_{1\vee 3}, e_{2\vee 4}\}$. Since E_A and E_B are not comparable in their granularity, the desiderata do not imply a preference between them.[7]

Second, a proposal made in Chalupka et al. [2015] suggests an interesting reason to sometimes go against the preference for specificity as implied by the dual account for EOC. Consider a simple causal model $C \to E$ for a cause variable C with three states $\{c_1, c_2, c_3\}$ and an effect variable E. Suppose that for both partitions $E_X = \{e_1, e_2, e_3, e_4\}$ and $E_Y = \{e_1, e_2, e_{3\vee 4}\}$ definite probabilities $P(E=col|do(C=row))$ can be specified as in the following two transition probability matrices:

$$\begin{array}{cc} & \begin{array}{cccc} e_1 & e_2 & e_3 & e_4 \end{array} \\ \begin{array}{c} c_1 \\ c_2 \\ c_3 \end{array} & \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.1 & 0.6 & 0.3 \\ 0 & 0.7 & 0.2 & 0.1 \end{bmatrix} \end{array} \qquad \begin{array}{c} \begin{array}{ccc} e_1 & e_2 & e_{3\vee 4} \end{array} \\ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0.1 & 0.9 \\ 0 & 0.7 & 0.3 \end{bmatrix} \end{array}$$

The dual account for EOC implies that – given that both partitions E_X and E_Y imply definite causal effects – the preference should be for the more specific E_X. But note that while the cause C allows some control over whether the outcome is e_1 or e_2 or e_3, there is no way for C to make any difference to the chance of the effect being e_3 vs. e_4: For c_1 those outcomes will never happen, and for c_2 and c_3, in each case there is a 2/3 chance that the outcome is e_3 and a 1/3 chance that it is e_4. So, while the partition E_X offers definite causal effects, the distinction it makes between e_3 and e_4 seems to be irrelevant to the causal influence of C. So, one may argue that the appropriate effect of C should be the coarser partition E_Y. Once the coarser value $e_{3\vee 4}$ is obtained, the probability of having the more specific e_i is not affected by interventions on C. In this sense, we may also say that $e_{3\vee 4}$ screens off more specific e_i from C.

This rationale for sometimes preferring a coarser effect variable does not seem to have a counterpart in the COE context. If deemed desirable, we can accommodate this preference in the dual account by adding a principle to *re-coarsen* a maximally specific definite effect through merging values according to the following equivalence relation:

$$e_i \sim e_j \quad \iff \quad \forall c_1, c_2 \in \mathscr{C} \quad \frac{P(e_i|do(c_1))}{P(e_i|do(c_2))} = \frac{P(e_j|do(c_1))}{P(e_j|do(c_2))},$$
$$\text{if } P(e_i|do(c_2)) > 0 \text{ and } P(e_j|do(c_2)) > 0.$$

[7] A related phenomenon is that the given cause variable may be proportional to multiple effect variables. All these effect variables may be on the same footing, as illustrated by the above example. Sometimes, however, these effect variables are comparable in granularity, in which case the dual account for EOC will still imply a preference.

This is a generalization of a principle proposed by Chalupka et al. [2015], which we consider to be too narrow.[8]

Finally, there is another possible reason to coarsen even further: In the above example, if one collapsed states e_2, e_3 and e_4, then one would make the causal relation between the given cause variable and the resulting effect variable deterministic. There may be circumstances where deterministic causal relations are the goal, and coarsening the effect variable moves toward that goal due to the summation of probabilities. This again is a consideration that does not apply to the COE question. One way to accommodate this preference in the dual account, if deemed desirable, is to add it as another criterion of admissibility. That is, only consider effect variables on which the given cause variable has definite and deterministic effects, and among such effect variables, specificity is to be preferred.

5. Conclusion

Our analysis makes explicit the dual nature of the cause-of-effect and effect-of-cause questions: One should maximally coarsen the cause and maximally refine the effect while maintaining definiteness of the relation. However, we show that the focus on specificity of the effect needs to be dampened with a cautionary note regarding a false sense of precision: there are definite causes of effects whose (effect-)specificity is exaggerated: effect distinctions are drawn over which the cause has no control. These subtleties in the characterization of the effect of a given cause seem to have no analog on the cause side.

For scientific practice, our analysis suggests a two-step process in characterizing a high-dimensional effect given a candidate cause. The first is a search for maximally specific partitions of the effect space on which the given cause variable has definite effects. The second step is an optional re-coarsening of the effect, in which definite effects for which the cause makes no difference, are combined, or where probabilistic effects are combined to obtain coarser, but deterministic outcomes. How to implement this process efficiently is an open and difficult research question.

References

Misha B Ahrens, Michael B Orger, Drew N Robson, Jennifer M Li, and Philipp J Keller. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods*, 10(5):413–420, 2013.

Pierrick Bourrat. Variation of information as a measure of one-to-one causal specificity. *European Journal for Philosophy of Science*, 9(1):11, 2018.

Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual causal feature learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 181–190. AUAI Press, 2015.

L. R. Franklin-Hall. High-level explanation and the interventionist's variables problem. *British Journal for the Philosophy of Science*, 67(2):553–577, 2016.

Paul E. Griffiths, Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunju Kim, and Rob Knight. Measuring causal specificity. *Philosophy of Science*, 82(4):529–555, 2015.

[8] Their rule is to merge values according to the following relation:

$$e_i \sim e_j \quad \Longleftrightarrow \quad \forall c \in \mathscr{C} \quad P(e_1|do(c)) = P(e_2|do(c))$$

Erik P Hoel. When the map is better than the territory. *Entropy*, 19(5):188, 2017.

Jaegwon Kim. Mechanism, purpose, and explanatory exclusion. *Philosophical Perspectives*, 3:77–108, 1989.

Jaegwon Kim. *Mind in a Physical World*. MIT Press, 2000.

Jaegwon Kim. *Physicalism, or Something Near Enough*. Princeton University Press, 2005.

Christian List and Peter Menzies. Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy*, 106(9):475–502, 2009.

Neil McDonnell. Causal exclusion and the limits of proportionality. *Philosophical Studies*, 174(6):1459–1474, 2017.

Matthew McGrath. Proportionality and mental causation: A fit? *Noûs*, 32(S12):167–176, 1998.

Brian P. McLaughlin. Mental causation and shoemaker-realization. *Erkenntnis*, 67(2):149–172, 2007.

J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2000.

Arnaud Pocheville, Paul Edmund Griffiths, and Karola C. Stotz. Comparing causes - an information-theoretic approach to specificity, proportionality and stability. In Hannes Leitgeb, Ilkka Niiniluoto, Elliott Sober, and Päivi Seppälä, editors, *Proceedings of the 15th Congress of Logic, Methodology and Philosophy of Science*, pages 260–275. 2017.

Larry Shapiro. Lessons from causal exclusion. *Philosophy and Phenomenological Research*, 81(3):594–604, 2010. doi: 10.1111/j.1933-1592.2010.00382.x.

Larry Shapiro and Elliott Sober. Against proportionality. *Analysis*, 72(1):89–93, 2012.

Peter Spirtes and Richard Scheines. Causal inference of ambiguous manipulations. *Philosophy of Science*, 71(5):833–845, 2004.

Bram Vaassen. Halfway proportionality. *Philosophical Studies*, (9):1–21, 2022.

David C Van Essen, Kamil Ugurbil, Edward Auerbach, Deanna Barch, Timothy EJ Behrens, Richard Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtiss, et al. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.

James Woodward. Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology and Philosophy*, 25(3):287–318, 2010.

James Woodward. Explanatory autonomy: The role of proportionality, stability, and conditional irrelevance. *Synthese*, 198(1):1–29, 2018. doi: 10.1007/s11229-018-01998-6.

James Woodward. *Causation with a Human Face: Normative Theory and Descriptive Psychology*. Oxford University Press, 2021.

Stephen Yablo. Mental causation. *Philosophical Review*, 101(2):245–280, 1992.

Stephen Yablo. Wide causation. *Philosophical Perspectives*, 11:251–281, 1997.

Lei Zhong. Sophisticated exclusion and sophisticated causation. *Journal of Philosophy*, 111(7):341–360, 2014.