

On the Explanatory Depth and Pragmatic Value of Coarse-Grained, Probabilistic, Causal Explanations

David Kinney*

London School of Economics

March 27, 2018

Abstract

This article considers the popular thesis that a more proportional relationship between a cause and its effect yields a more abstract causal explanation of that effect, which in turn produces a deeper explanation. This thesis is taken to have important implications for choosing the optimal granularity of explanation for a given explanandum. In this article, I argue that this thesis is not generally true of probabilistic causal relationships. In light of this finding, I propose a pragmatic, interest-relative measure of explanatory depth. This measure uses a decision-theoretic model of information pricing to determine the optimal granularity of explanation for a given explanandum, agent, and decision problem.

1 Introduction

The special sciences frequently make probabilistic generalizations. To take a classic example, a patient who habitually smokes cigarettes is said to have a certain probability of developing lung cancer. These generalizations often facilitate causal explanations. In the case of smoking and lung cancer, if we ask *why* a given smoker develops lung cancer,

*I am highly indebted to Katie Steele and Luc Bovens for their feedback and support throughout the writing of this paper. I am also grateful to the following people for comments and conversations about this paper and its subject matter: Jonathan Birch, Philippe van Basshuysen, Hugh Desmond, Phil Dowe, Bryan Roberts, Jeremy Strasser, and David Watson. This work was presented to audiences at LSE, ANU, the 2016 conference of the Dutch Research School of Philosophy in Groningen, the 2017 meeting of the Munich-Sydney-Tilburg Philosophy of Science Group in Sydney and the 2017 meeting of the Society for Metaphysics of Science at Fordham University.

we can explain their affliction by saying that smoking was the likely cause. Similarly, if we observe a population in which more smokers develop lung cancer than non-smokers, then we can say that the higher incidence of lung cancer among smokers is causally explained by the higher prevalence of smoking.

In each of these cases, we could also tell a more complicated story. Instead of explaining a patient's lung cancer by citing the fact that they smoked, we could attempt to cite a long series of facts about individual strands of burnt tobacco producing carcinogenic particles that entered the smoker's lungs at different spatiotemporal points, eventually causing them to develop lung cancer. Clearly, there are pragmatic reasons for avoiding this kind of reduction. The level of detail required for the second kind of explanation would be incredibly cumbersome, and clearly not suited for the task at hand. In addition, many recent authors in philosophy of science, including Strevens (2008), Woodward (2010, 2016), Weslake (2010, 2013), Weatherson (2012), Franklin-Hall (2016) and Clarke (2017), have argued that coarse-grained explanations are to be preferred on explanatory as well as pragmatic grounds.¹ The central idea that each of these authors express is that there is a sense in which some coarse-grained causal generalizations provide *deeper* explanations than their more fine-grained counterparts. Weslake defines explanatory depth as "a measure in terms of which explanations can be assessed according to their explanatory value" (2010, 273). Going forward, I will adopt this definition; explanatory depth refers to the relative value of an explanation *qua* explanation.

Most of these authors consider, and accept to varying degrees, the thesis that higher-level causal explanations are deeper than lower-level explanations when and because the causes cited are *proportional* to their effects. The meaning of proportionality in this context will be made precise below, but the basic idea is as follows. When a cause and an effect are both represented by a variable taking some value, the cause is proportional to the effect to the extent that it is true that if the causal variable had taken a different value, then the effect variable would also have taken a different value. Notably, Weslake (2010, 2013) draws a link between the relative proportionality and the relative *abstractness* of an explanation, where abstractness is the ability of an explanatory causal model to be deployed in a wide range of scenarios. Abstractness, he argues, provides a crucial link between proportionality and explanatory depth; proportional explanations

¹Historically, similar positions have been held by Putnam (1979), Kitcher (1981), Garfinkel (1981), Jackson and Pettit (1992), Sklar (1993), and Batterman (2001).

are deeper than less proportional ones because they are more abstract.

However, the linkage between proportionality, abstractness, and explanatory depth has largely been developed in deterministic contexts. In what follows, I argue that in probabilistic contexts, it can be the case that a more proportional explanation is less abstract than a less proportional explanation of the same explanandum. In some of these cases, the more proportional explanation seems to be deeper, and in others, the more abstract explanation seems to be deeper. Thus, probabilistic cases of causal explanation muddy the waters when it comes to the relationship between proportionality, abstractness and depth. In these cases, I claim that we need to consider the pragmatic value of an explanatory model in order to make sense of its relative explanatory depth. Thus, I reject the strict bifurcation between pragmatics and explanatory depth, and conclude that pragmatics may play an indispensable role in assessing the relative depth of competing causal explanations.

The plan for this essay is as follows. In Section 2, I provide some background on probabilistic causal explanation and the formal notion of the granularity of explanation. In Section 3, I explicate the nature of the relationship between proportionality, abstractness, and explanatory depth, as it is developed in deterministic contexts. In Section 4, I turn to probabilistic contexts. I introduce a probabilistic account of proportionality put forward by Pocheville et al. (2017), and use it to show how proportionality and explanatory depth can come apart in cases where subtle differences in the nature of a cause lead to slight differences in the probability of its effect. I then show how abstractness and explanatory depth can also come apart in probabilistic cases. In Section 5, I introduce my proposed measure of the pragmatic value of a causal model, using the decision-theoretic concept of information pricing. I use this measure to argue that the pragmatic value of a causal explanation allows us to determine the optimal level of granularity for that explanation in a way that proportionality and abstractness are unable to do. In Section 6, I address some possible objections to my argument. In Section 7, I offer concluding remarks.

2 Background

The leading contemporary philosophical account of type-level causal explanation is the “interventionist” or “manipulationist” account developed by Pearl (2000) and Woodward

(2003). Put very briefly, this account holds that causal relata should be represented as variables in a structural graph, and that one variable in the graph is causally related to another if and only if at least some interventions on the value of the causal variable result in a change in the probability distribution over the effect variable. In this context, an ‘intervention’ is a change in the value of the causal variable that eliminates any other factors that may be relevant to the value of the causal variable. Events are represented by variables taking particular values, and one event explains another causally if and only if: (i) the explanans event can be represented as a value c_j of the variable C , and the explanandum event can be represented as a value e_i of the variable E , such that some interventions on C change the probability distribution over E , and (ii) an intervention that brings about c_j raises the probability of e_i , i.e. $p(e_i|\hat{c}_j) > p(e_i)$, where the hat symbol over the value c_j indicates that the variable C is set to c_j via an intervention.

To illustrate, consider a variable X , which takes the value x_1 if a person smokes and x_2 if they do not. Next, we can introduce a variable L which takes the value l_1 if a person has lung cancer, and l_2 if they do not. We want to say that x_1 causally explains l_1 , i.e. that smoking causally explains lung cancer. Interventions that change the value of X also change the probability distribution over L ; if we intervene so that a person becomes a smoker, then they are more likely to develop lung cancer, and if we intervene so that a person does not become a smoker, then they are less likely to develop lung cancer. So it is true that X causes L , according to the interventionist account. Thus, condition (i) for x_1 causally explaining l_1 is satisfied. Next, we can suppose that $p(l_1|\hat{x}_1) > p(l_1|\hat{x}_2)$, so that $p(l_1|\hat{x}_1) > p(l_1)$. Thus, condition (ii) is satisfied, and smoking causally explains lung cancer.²

This approach to causal explanation is contrastive. The relation of probability raising under intervention between $C = c_j$ and $E = e_i$ not only implies that the occurrence of the event represented by $C = c_j$ explains why the event represented by $E = e_i$ occurs. It also implies that the occurrence of the event represented by $C = c_j$ explains why the event represented by $E = e_i$ occurs *rather than* the set of events represented by $\neg(E = e_i)$. This is a special case of contrastive explanation, one in which the “fact” to be explained, i.e. the event represented by $E = e_i$, is contrasted with a “foil” that is its

²I am leaving to one side any instances of probability-raising non-causes and probability-lowering causes that may occur in cases of pre-emption or causal backups (see Fenton-Glynn 2015).

logical negation.³ This positive instance of contrastive explanation is underwritten by the fact that interventions bringing about $C = c_j$ render $E = e_i$ more likely to occur than its negation. It is worth noting that there is a substantial debate over whether contrastive explanations can ever be probabilistic, with Lewis (1986) perhaps the most prominent proponent of the view that chancy events cannot be explained contrastively, and Percival (2000) arguing that contrastive explanation is impossible when both fact and foil are chancy events. Though I do not have space here to address this debate in detail, for the sake of argument I am adopting Hitchcock’s (1996, 1999) position, *contra* Lewis and Percival, that there is no conflict between contrastive and probabilistic explanation.⁴

Using variables to represent types of phenomena in nature also provides a framework for stating two crucial necessary conditions for one variable to be a more fine-grained version of another. Let us stipulate that C' is a more fine-grained version of C , or, in other words, that C is a more coarse-grained version of C' , only if:

1. For any c_j that is a value of C , there is a c'_i that is a value of C' such that $C' = c'_i$ implies $C = c_j$.
2. There exists at least one value c_j and one value c'_i such that $C = c_j$ does not imply $C' = c'_i$.

Moving forward, when I stipulate that causal variables that are fine-grained and coarse-grained versions of each other, this implies that these two key constraints relating the variables are satisfied. Additionally, I assume that all variables form an exhaustive partition over the possible states of the world.

3 Depth, Proportionality, and Abstractness

The same event can be given a different causal explanation in virtue of changing the granularity of the causal variable. To see how this works, let Y be a variable that takes the value y_1 if a person smokes Marlboros, y_2 if they smoke another brand, and y_3 if

³See Lipton (1990) for more detail on the fact-foil distinction and its role in contrastive explanation.

⁴Although Hitchcock is, strictly speaking, concerned with the “explanatory relevance” of one event to another, where ‘explanatory relevance’ is defined so that event a explains event e just in case $p(e|a) \neq p(a)$, I take it that his defense of probabilistic contrastive explanation carries over to an account that conceives of causal explanation as probability-raising under intervention.

they do not smoke. Here, there are *at least some* interventions on the value of Y that change the probability distribution over L , i.e. those interventions that change the value of Y from either y_1 to y_3 or y_2 to y_3 , or vice-versa. It may also be the case that all interventions on Y change the probability distribution over L , if Marlboros are more or less carcinogenic than other brands, but this is not necessary for a causal relationship; all that is required is that at least some interventions on the causal variable lead to changes in the probability distribution over the effect variable. If we make the further plausible assumption that smoking Marlboros raises the probability of developing lung cancer relative to the prior probability of lung cancer, then we can say that if a patient who smokes Marlboros develops lung cancer, then her smoking Marlboros causally explains their lung cancer.

In a sense, this is exactly the result that we want. It is true that the fact that a patient smokes Marlboros causally explains why they develop lung cancer. However, there is also a sense in which the explanation is suboptimal. Instead of explaining the patient's lung cancer by citing the fact that they smoked Marlboros, we could say that the patient developed lung cancer because they smoked. This second, more coarse-grained explanation seems preferable. After all, the brand of cigarettes that a person smokes does not matter much, if at all, to their chances of getting lung cancer; the relevant detail here is that the patient smokes. Providing a rationalization for this kind of intuition is one goal of the literature on the explanatory depth of causal explanations.

As stated in the introduction, a popular idea in the literature on higher-level causal explanation is that more coarse-grained causal explanations can be chosen over more fine-grained explanations of the same causal phenomenon on the basis that the more coarse-grained explanation is more *proportional* than the more fine-grained one. This idea appears in Yablo (1992), and is endorsed to varying degrees by Craver (2007), Woodward (2010), Malaterre (2011), Weslake (2013), and Stegmann (2014). Woodward (2010) provides perhaps the most comprehensive definition of proportionality as a property of the relationship between cause and effect. For Woodward, proportionality is a property of causal relationships that comes in degrees. A causal variable stands in a proportional relationship to a given effect variable to the extent that the causal variable is described so that: (a) each possible intervention changing the value of the causal variable determines the value of the effect variable, and (b) all possible interventions changing the value of the causal variable correspond to changes in the value of the effect

variable (Woodward 2010, 296).

We can use a classic example from Yablo (1992) to illustrate Woodward's understanding of proportionality. Suppose that we have trained a pigeon to peck at only red targets, to the point that the pigeon will peck at any red target, and refrain from pecking at targets of any other color, as a matter of certainty. Now suppose that we present the pigeon with a scarlet target, leading it to peck. We can say that the pigeon pecked at the target either because the target was scarlet or because the target was red. Clearly, saying that the pigeon pecked because the target was scarlet fails to fully satisfy proportionality. Specifically, there is a failure to satisfy condition (b). We can intervene on the target to change it from scarlet to crimson, and the pigeon will still peck. By contrast, any intervention that changes the color of the target from red to some other color will make it the case that the pigeon does not peck. So the explanation 'the pigeon pecked because the target was red' perfectly satisfies (a) and (b), whereas 'the pigeon pecked because the target was scarlet' does not.

However, this does not answer the question of why we regard proportional explanations as being deeper than less proportional ones. Weslake (2010) argues that part of what is at work here is that we tend to regard more *abstract* explanations as being deeper than less abstract ones. For Weslake, when we compare two explanations, the more abstract explanation is one that we can use in a wider variety of cases, without changing the values of the variables used. Every Marlboro smoker is also a smoker, so the explanation 'smoking caused the patient's lung cancer' will be applicable in all cases in which the explanation 'smoking Marlboros caused the patient's lung cancer' is applicable. However, the reverse obviously does not hold; the explanation 'smoking Marlboros caused the patient's lung cancer' does not apply in all the cases in which 'smoking caused the patient's lung cancer' does.

Weslake (2013) argues further that there is a clear connection between the abstractness of a causal explanation and its proportionality. If we consider the pigeon example, it is clear that the explanation 'the pigeon pecked because the target was scarlet' is applicable to less cases than 'the pigeon pecked because the target was red', and that in any case in which the former applies, the latter also applies. Indeed, in deterministic cases a perfectly proportional causal explanation of a given explanandum will generally be more abstract than a less proportional explanation of the same explanandum. If all red targets cause the pigeon to peck, and all non-red targets cause the pigeon not

to peck, it follows that any red target that has some other property A will also cause the pigeon to peck, and any non-red target that has some other property B will cause the pigeon not to peck. Thus, a more proportional explanation is applicable wherever a less proportional one is also applicable. However, it follows from the truth of the more proportional explanation that the pigeon will also peck when presented with a red target that does not have property A , and not peck when presented with a non-red target that does not have property B . The explanation ‘the pigeon pecked because it was presented with a red target with property A ’ will not apply in these cases, but the more proportional explanation ‘the pigeon pecked because the target was red’ will. Thus, proportionality tracks abstractness in deterministic cases.

Moving further, there is also a strong intuitive connection between abstraction and explanatory depth. I take it that the force of this intuition lies in the thought, most famously expressed by Kitcher (1981), that explanation in general involves the subsumption of specific events under more general patterns, in this case patterns of causation. If such subsumption is part of the purpose of explanation, then it makes sense that we would prefer to give explanations that are able to subsume more events under a given type-level causal pattern. This preference for abstraction is linked closely to what Ylikoski and Kuorikoski call the “cognitive salience” of a given explanation. They write that “the kinds of inferences possible for limited cognitive systems such as humans directly effect what can be explained and understood by such cognitive systems” (2010, 214). Since more abstract explanations tend to use more coarse-grained variables, they allow us to use “the same words [to refer] to similarly structured forms of causal interaction in different fields,” thereby reducing the cognitive burden associated with causal explanation in many cases (214).

Weslake argues that an emphasis on abstraction as a good-making feature of an explanation stands in contrast with Hitchcock and Woodward’s (2003) claim that causal explanations are deeper when and because they are situated within an explanatory model that is more “invariant”, where an invariant model is one that answers more what-if-things-had-been-different questions about the target system. To use Weslake’s example, the ideal gas law can be used to explain the change in the pressure of a gas with a fixed temperature and changing volume. That is, we can use the ideal gas law to answer questions about what would happen under various changes in the volume of the gas’ container. However, a model containing a variable representing the microphysical state

of the gas, rather than its volume, answers just as many counterfactual questions as the more abstract, coarse-grained model. For Weslake, this shows the limitations of invariance as a dimension of explanatory depth; highly fine-grained explanations seem to have greater or equal invariance than coarse-grained explanations in cases where a coarse-grained explanation seems more appropriate. I concur with Weslake's argument here, although see Woodward (unpublished) for a response.

4 The Probabilistic Context

In probabilistic contexts, the picture presented above becomes considerably more complicated. So we will need to translate the concepts of proportionality and abstractness to accommodate probabilistic relationships between variables. In what follows, I will put forward what I take to be reasonable probabilistic extensions of the concepts of proportionality and abstractness. I will show that there are some cases in which a *less* proportional causal explanation seems to be deeper than a more proportional one. Similarly, I will show that in other cases, a *less* abstract explanation seems to be deeper than a more abstract competitor. Thus, I conclude that the tight relationship between proportionality, abstractness and explanatory depth comes apart in probabilistic contexts.

Let us begin by defining a probabilistic measure of proportionality. Clearly, the understanding of proportionality articulated above is not suited to probabilistic relationships, since it measures the proportionality of a causal relationship according to the extent to which the various states of the cause determine the state of the effect. So we will need a new measure for assessing the comparative proportionality of causal relationships where the relationship between cause and effect is probabilistic. One attempt at this task is put forward by Pocheville et al. (2017), who understand proportionality as follows. If refining the causal variable makes no difference to the conditional probability distribution over some effect variable, then the relationship between the coarse-grained causal variable and the effect variable is more proportional than the relationship between the fine-grained causal variable and the effect variable. However, if refining the causal variable makes *any difference at all* to the conditional probability distribution over the effect variable, then the relationship between the fine-grained causal variable and the effect variable is more proportional than the relationship between the coarse-grained

causal variable and the effect variable.

To see why this translation of proportionality into the probabilistic context is warranted, note that, in general, the concept of proportionality aims to capture the salient *difference-making* relationship between cause and effect (see List and Menzies 2009). In the pigeon case, when we say that the pigeon pecked at the target because it was scarlet, we fail to identify what is really making the difference with respect to whether the pigeon pecks. It is the target's redness, rather than its being any shade of red, that makes the difference. A perfectly proportional causal explanation is one that conveys all and only the difference-makers with respect to some effect. In probabilistic cases, it seems that we have no choice but to interpret this difference-making constraint probabilistically, and say that the proportional explanation is the one that identifies every factor that makes a difference to the probability distribution over the effect, and only those factors.⁵

We can show that Pocheville et al.'s criterion gets the right results in some intuitive cases. Consider the following case. Suppose that the probability of a smoker developing lung cancer is .1, and the probability of a non-smoker developing lung cancer is .01. Suppose that the brand of cigarette smoked does not make a difference to the patient's likelihood of lung cancer; both Marlboro smokers and non-Marlboro smokers have a .1 probability of developing lung cancer. Pocheville et al.'s approach renders the verdict that 'smoking caused the patient's lung cancer' is a more proportional explanation than 'smoking Marlboros caused the patient's lung cancer'. This fits with common sense as well; in cases where the brand of cigarette smoked makes no difference to a person's probability of developing lung cancer, the more proportional, more abstract, and deeper explanation is one that eschews any mention of cigarette branding.

So far, so good. However, once we let go of the constraint that refining the causal variable makes no difference to the probability of the effect, proportionality seems to come apart from explanatory depth. Suppose that smoking Marlboros results in a .1001

⁵In Pocheville et al.'s paper, they aim to define proportionality and other aspects of causal relationships using information theory, as formulated by Shannon and Weaver (1949). On their view, if the mutual information between an intervention on a causal variable and its effect is greater than that between another cause-effect pair, then the first relationship is more proportional. If two relationships have equal mutual information, the more proportional relationship is that with minimal entropy. This criterion ensures that any refinement of the causal variable that makes any difference to the conditional probability distribution over the effect variable will yield a more proportional causal relationship.

probability of lung cancer, and smoking other brands results in a .0999 probability of lung cancer. As before, smoking in general results in a .1 probability of lung cancer, and non-smoking results in a .01 probability of lung cancer. It follows from Pocheville et al.’s approach that the explanation ‘smoking Marlboros caused the patient’s lung cancer’ is more proportional than the explanation ‘smoking caused the patient’s lung cancer’. However, I take it that the latter explanation is deeper than the former. Certainly, in the actual practice of science it would seem overly specific to cite the brand of cigarettes that a person smokes when explaining why they developed lung cancer, simply because the person chose a very slightly more carcinogenic brand. What seems to matter here for explanatory purposes is that there is an important causal trend—smoking causes lung cancer—that a given case can be subsumed under. Thus, proportionality and explanatory depth come apart in this case.

Next, we can extend the notion of abstractness into the probabilistic setting. Recall that causal explanation is defined above as probability-raising under intervention; $C = c_j$ explains $E = e_i$ if and only if an intervention setting the value of C to c_j raises the probability that $E = e_i$. Next, recall that the most abstract explanation of a given explanandum event is one that subsumes the widest class of specific events under a given causal relation. Thus, ‘the patient’s smoking caused her lung cancer’ is a more abstract explanation than ‘the patient’s smoking Marlboros caused her lung cancer’. More generally, we can say that the most abstract explanation of a given explanandum is one that subsumes under one broad category all the ways in which the probability of the explanandum event could be raised via intervention. Typically, we think that abstractness in explanation tends to go hand-in-hand with explanatory depth; explanations that are not optimally abstract are usually overly detailed and therefore suboptimal.

However, there are cases that put pressure on the notion that explanations are deeper when and because they are more abstract. Consider a case where a person suffering from some bacterial disease receives an antibiotic containing penicillin and recovers. Suppose that in general, antibiotics containing penicillin have a .95 probability of curing patients with this disease. However, suppose that doctors could also prescribe the patient an antibiotic that does not contain penicillin, and that in these circumstances the patient would have recovered with probability .7. Without any antibiotics, the recovery rate is only .1. For some assignments of prior probabilities to each causal event, administering either antibiotic raises the patient’s probability of recovery. Thus, a more abstract rival

to the explanation ‘an antibiotic containing penicillin caused the patient to recover’ would be the explanation ‘an antibiotic caused the patient to recover’. However, it is not clear to me that in this case, the more abstract explanation is the better or deeper one. It seems to matter here that the patient’s recovery was aided by penicillin, and leaving this fact out seems to miss something important about the patient’s recovery. To use Franklin-Hall’s phrase, the more abstract explanation “overshoots” the optimal level of explanatory granularity (2016, 570). Thus, in some probabilistic cases, abstraction does not track explanatory depth.

In light of these arguments, we will need to look for another way of determining the optimal level of granularity for explanation in probabilistic cases. In the next section, I will introduce my strategy for measuring the pragmatic value of a causal model. I will then argue that by considering the pragmatic value of a causal model, we can arrive at a compelling justification of our intuitions regarding explanatory depth in the cases described above.

5 The Pragmatic Value of a Causal Model

It is clear from the arguments in the previous section that in probabilistic contexts, an explanation can be maximally proportional but overly detailed, because the differences that various refinements of the causal variable make to the probability distribution over the effect variable are too small. On the other hand, an explanation can be too abstract; sometimes details that make a difference to the probability distribution add explanatory depth, even if we can pick out the class of probability-raising factors for some explanandum without using these details. Weatherson (2012) calls this difficulty in finding the best level of detail for a given explanation the “Goldilocks Problem”; we need a level of granularity for our causal explanations that is not too detailed, not too abstract, but just right.

Franklin-Hall writes that “making sense of the explanatory—not merely the practical—superiority of high-level explanations in a physical world has been a kind of Holy Grail in the philosophy of science, long sought but never found” (2016, 555). My own thought is that a distinctly non-pragmatic dimension to explanatory depth may not exist. Indeed, when we rationalize the intuition that a given detail ought to be included or not included in the deepest explanation of some event, it seems that we invariably revert

to pragmatic language. If asked why we do not usually include the brand of cigarette smoked in an explanation of why a patient developed lung cancer, even when a patient smokes a slightly more carcinogenic cigarette, there is a strong urge to say that the slight difference in brands just does not *matter* enough to be worth documenting in an explanation. By contrast, in the antibiotics case it *does* seem to matter that penicillin was administered to the patient. This emphasis on the degree to which a given detail matters already shifts us into a consideration of the pragmatic value of an explanation. So it makes sense that we would appeal to pragmatics in assessing the overall superiority of a higher-level causal explanation as compared to a lower-level alternative.

There is a useful analogy here to the epistemic standard required for conviction in a criminal trial. Juries are asked to determine whether the defendant is guilty “beyond a reasonable doubt”. Of course, the nature of this standard is vague. If by ‘beyond a reasonable doubt’ we mean absolute certainty, then few if any defendants will ever be found guilty. So it must be that a juror’s credence in the defendant’s guilt must rise above some threshold probability. This generates the problem of where to set the threshold. It is at this point that pragmatics will begin to play a role. If the defendant is facing life in prison, then the threshold will have to be set very high, to account for the negative consequences of making a mistake. If the possible sentence is much less serious, then the jury’s standard will likely drop as well. I take it that a similar dynamic is at play in the pragmatic evaluation of explanatory models. There is some threshold for the degree of probabilistic difference-making between possible values of a causal variable such that those values should be included in our model. However, setting the threshold at ‘any difference at all’ is unrealistic and unintuitive; it would yield hopelessly complicated explanatory models. What we need to do is consider the pragmatic consequences of including or not including the detail in an explanatory model, and set our threshold for explanatorily relevant difference-making accordingly.

Thus, I will seek to develop a distinctly pragmatic criterion for determining the best level of description for a given probabilistic explanation. My approach can be summarized as follows. When a scientist or ordinary person explains a phenomenon causally, there is a sense in which they make a decision. They take a particular cause of some effect and chooses to embed it with a given causal model, as one value of a causal variable that may be defined at many different levels of granularity. I argue that in making this decision, the agent should choose the most coarse-grained causal variable

	Lung Cancer	No Lung Cancer
Approve	-100	10
Deny	0	0

Table 1: Utility Matrix For Life Insurance Decision

that does not elide any information that is of value to her. To justify this preference for coarse-grained causal explanations, recall both Weslake (2010) and Ylikoski and Kuorikoski’s (2010) insistence that it is a virtue of some explanatory models that we can use the same names to refer to similar causal factors, thereby providing a unified understanding of why some phenomenon occurred. My criterion takes this insistence on coarse-graining seriously, arguing that if there is no pragmatic cost to an agent from coarsening their description of a causal factor, then that agent should coarse-grain.

My proposal measures the value of information using the decision-theoretic concept of the fair price of causal information (FPCI), which has its roots in Blackwell (1951), Savage (1954) and Good (1967), and is also discussed in Resnik (1987). To illustrate the concept, consider the following example. Suppose that an analyst has to decide whether to approve or deny an individual for life insurance. In an obvious oversimplification, suppose that only one fact—whether or not the potential customer develops lung cancer—will determine whether offering the policy will lead to a positive or negative outcome for the insurance company. How much should the analyst pay to learn whether or not the potential customer smokes?

Let us begin by explicating the decision problem that the analyst faces. They have to choose between two actions: Approve or Deny. The utility that they get from performing each of these actions depends on whether or not the potential customer will develop lung cancer. Table 1 specifies the utilities of each action in each relevant state of the world. The analyst cannot learn whether the customer will develop lung cancer, but can learn whether or not the customer smokes. Suppose further that if a person smokes, then there is a .1 probability that they will develop lung cancer. If they do not smoke, then there is a .01 probability that they will develop lung cancer. Assume that each of these probabilities are “causal” probabilities, in the sense that the value of the causal variable is set via an intervention. Now suppose that the analyst learns that the potential customer does in fact smoke. Under these conditions, the expected utility of Approve

is $.1(-100) + .9(10) = -1$ and the expected utility of Deny is 0. So the analyst would choose to Deny, and expect a utility of 0. Next, suppose that the analyst learns that the person does not smoke. Under this supposition, the expected utility of Approve would be $.01(-100) + .99(10) = 8.9$, and the expected utility of Deny would be 0. So the analyst would choose to Approve, and expect a utility of 8.9. Finally, let us suppose that the analyst has the prior belief that a person is equally likely to smoke or not smoke. Under this supposition, the expected utility from learning the potential customer's smoker status is $.5(0) + .5(8.9) = 4.45$.

However, to calculate the fair price of information about whether or not the potential customer smokes, one also has to consider what the analyst would have done had they not learned the potential customer's smoker status. Her prior beliefs about whether the potential customer smokes and her posterior beliefs about whether the potential customer develops lung cancer (given the customer's smoker status) jointly imply that the potential customer has prior probability .055 of developing lung cancer, and therefore a prior probability .945 of not developing lung cancer. Thus, if the analyst does not learn whether the customer smokes, then the expected utility of Approve is $.055(-100) + .945(10) = 3.95$ and the expected utility of Deny is 0. Since the analyst always acts to maximize expected utility, in the absence of information as to the potential customer's smoker status, they would choose Approve. If we subtract the expected utility of Approve when smoker-status is unknown from the expected utility of learning whether the potential customer smokes, we get $4.45 - 3.95 = .5$. This is the fair price of information regarding whether or not the potential customer smokes; the analyst should only pay up to half a unit of whatever currency the utilities are expressed in to learn whether or not the potential customer smokes.⁶

In some cases, refining the causal variable adds no pragmatic value to a causal model. Consider the case from the previous section, in which a smoker has a probability .1 of developing lung cancer regardless of the brand that they smoke, and a non-smoker has a .01 chance of developing lung cancer. The utility matrix for the decision to approve or deny a customer for life insurance is the same as in Table 1. If we suppose that the brand smoked by the potential customer makes no difference to their probability of developing lung cancer, then we know that learning this information will not change the action with maximum expected utility, as compared to when the customer smokes. Thus, the FPCI

⁶See the appendix for an abstract model of the fair price of causal information.

with respect to whether the customer smokes Marlboros, another brand, or does not smoke will be the same as the FPCI with respect to whether or not the customer smokes. As there is no loss of pragmatic value at the more coarse-grained level of description, choosing the fine-grained variable over the coarse-grained variable introduces a level of complexity to the model that is worthless, from a pragmatic perspective, and therefore the refinement ought to be avoided.

We do not need to worry here about the prior probability distribution over either causal variable. If it is the case that an agent would pay more to learn the value of the fine-grained variable than they would to learn the value of its coarse-grained counterpart, then the prior probability distribution over causal variables determines to some degree *how much* more an agent should pay for information about the value of a more fine-grained variable as compared to its coarse-grained counterpart. However, changes in the prior distribution over the causal variable will never undo this inequality. Similarly, if information about the value of a fine-grained variable is worth no more to an agent than information about the value of a coarse-grained version of the same variable, this equality will be invariant across changes to the probability distribution over either variable.

In general, we want it to be the case that if refining the causal variable makes no difference to the probability distribution over the effect variable, then the coarse-grained explanation is to be preferred to the fine-grained explanation on pragmatic as well as explanatory grounds. This should hold regardless of the agent’s pragmatic interests. The FPCI measure defined here satisfies this desideratum. Let $\pi(\hat{C}, \mathbf{U}, E)$ denote the FPCI for a causal variable C , given a utility matrix \mathbf{U} and effect variable E . The following proposition is true (see the appendix for proof):

Proposition 1. *For any causal variables C and C' , effect variable E , utility matrix \mathbf{U} and probability distribution $p(\cdot)$, where C' and C are both causes of E and C' is a fine-grained version of C , if $p(e_i|\hat{c}'_i) = p(e_i|\hat{c}_j)$ for all sets of values e_i , c_j , and c'_i such that $C' = c'_i$ implies $C = c_j$, then $\pi(\hat{C}, \mathbf{U}, E) = \pi(\hat{C}', \mathbf{U}, E)$.*

This means that in the class of “easy cases” in which refining the causal variable makes no difference to the probability distribution over the effect, it is also the case that the coarse-grained explanation provides equal pragmatic value with less complexity, and can therefore be chosen over a fine-grained alternative. Of course, these are the same

easy cases in which a more proportional explanation is also more abstract, and also intuitively deeper. What we need to do now is consider those cases in which a more proportional or more abstract explanation seems less deep than a less proportional or less abstract rival, and determine whether an FPCI-based measure of the pragmatic value of an explanation is of any help.

First, we can consider the case where proportionality and depth come apart. Here, the probability of developing lung cancer from smoking Marlboros was .1001, and the probability of developing lung cancer from smoking other brands was .0999. If we again use the utility matrix in Table 1, one can check that learning whether the potential customer smokes Marlboros, another brand, or does not smoke will not make any difference to the analyst’s decision, as compared to simply learning whether or not the customer smokes. Thus, she will pay no more to learn the value of this more fine-grained variable than she would to learn to value of its coarse-grained counterpart. By considering the pragmatic value of causal information in assessing explanatory depth, we can conclude that that the coarse-grained explanation ‘the patient developed lung cancer because they smoked’ is deeper than the fine-grained explanation ‘the patient developed lung cancer because they smoked Marlboros’, even in cases where Marlboros are slightly more carcinogenic than other cigarettes.

Next, consider the case of treating a patient with antibiotics that may or may not contain penicillin. Recall that in this case, a less abstract explanation seems to be deeper. Again, we can consider the decision faced by a life insurance analyst, where the analyst’s utilities in the state of the world where the patient does or does not recover from the bacterial infection are given in Table 2. The relevant expected utilities for the analyst are as follows.

$$EU(\text{Approve}|\text{Penicillin}) = .05(-100) + .95(10) = 4.5 \quad (1)$$

$$EU(\text{Approve}|\text{Non-Penicillin Antibiotic}) = .3(-100) + .7(10) = -23 \quad (2)$$

$$EU(\text{Approve}|\text{No Drugs}) = .8(-100) + .2(10) = -78 \quad (3)$$

Once again the expected utility of Deny is always 0. Let us assume that a person has a prior probability .4 of receiving antibiotics with penicillin, a prior probability .4 of receiving antibiotics without penicillin, and a prior probability .2 of receiving no treatment. This entails that the prior probability of No Recovery is .32 and the prior

	No Recovery	Recovery
Approve	-100	10
Deny	0	0

Table 2: Utility Matrix For Life Insurance Decision

probability of Recovery is .68, so that the expected utility of Approve in the absence of information about the patient’s treatment is $.32(-100) + .68(10) = -25.2$. Thus, the fair price of causal information regarding the patient’s treatment is $(.4(4.5) + .4(0) + .2(0)) - 0 = 1.8$.

If we coarse-grain the causal variable and consider only whether or not the patient receives antibiotics, the model loses pragmatic value. One can check that if the analyst knew that the patient had received antibiotics, but not whether those antibiotics contained penicillin, then the analyst would choose Deny. Since the analyst would also choose Deny if they learned that the patient did not receive antibiotics, and would choose Deny in the absence of information, there is no value to learning whether or not the patient received antibiotics; the fair price of causal information is zero. Thus, the coarse-grained explanatory model of the patient’s recovery, i.e. the model that makes no mention of penicillin, is less pragmatically valuable than the fine-grained model that does mention penicillin. This goes some way towards providing a rationalization of why greater abstraction does not seem to result in greater explanatory depth in this case; the more abstract explanation conceals information that rises above some threshold of pragmatic usefulness. Thus, my pragmatic approach to explanatory depth proves useful in this case as well.

I do not claim that FPCI constitutes a totalizing measure of explanatory depth; I find it unlikely that such a measure could exist. However, I believe that FPCI addresses an important issue in the literature on explanatory depth. Woodward (2010), Ylikoski and Kuorikoski (2010) and Weslake (2010) all accept that the depth of a causal explanation at a particular level of granularity, as compared to explanations of the same phenomenon at finer or coarser grains, is sensitive to our interests as agents. However, this is often described in a very general way, without much mathematical precision. My proposal makes the interest-relative dimension of explanatory depth precise, by mathematically linking it to the utilities that an agent assigns to various outcomes in a given decision

problem. More precisely, my proposal quantifies the extent to which an explanatory causal model enables an agent to make the best decision possible, as defined by the agent's own utility function over outcomes, while minimizing cognitive burdens due to the complexity with which the explanans is described. This stands in contrast to purely epistemic measures of explanatory depth, which consider only the probability distribution over the possible values of an explanandum, given the possible values of the explanans.

6 Potential Objections

There are several potential objections to the view presented above. First, one could worry that it is strange to consider the value of information in cases where we already possess the information in question. In the smoking case, we may already know that the patient is a Marlboro smoker, and in the antibiotics case, we may already know that the patient has been prescribed antibiotics containing penicillin. So why does it matter what the insurance analyst would pay to learn this information in a counterfactual scenario where they did not already know it?

One way of replying to this worry is to say that our choice of granularity for an explanatory model guides future searches for explanatory information. For example, when we say of a given patient that smoking caused her lung cancer, this implies that when we search for an explanation of some other patient's lung cancer, we can begin by asking "did they smoke?" Clearly, the answer to this question will be valuable to us as both explanatory and pragmatic agents. Now suppose instead that we ask, "did they smoke Marlboros?" If the answer is "no", then we will need to ask another question, namely, "did they smoke any other brands?" Regardless of the answer to this question, it is clear that it would have been better just to ask, in the first place, "did they smoke?" By contrast, in the antibiotics case, it seems that the answer to the question "was the patient given penicillin?" will be explanatorily useful to us in understanding why the patient recovered, regardless of whether the answer is "yes" or "no." In this case, the introduction of a finer granularity of causal information seems to provide a meaningful direction of inquiry in future cases, rather than merely introducing unnecessary complexity to scientific practice.

Next, one could worry that in the examples that I have given above, the interven-

tionist theory of causation is actually doing very little work. For the life insurance analyst in question, information about the causes of lung cancer will not necessarily be any more valuable than information about non-causal correlates of lung cancer. To illustrate, it might be the case that information about how yellow a customer's fingers and teeth are may be as valuable to our insurance analyst as information about the customer's smoker status, even though yellow teeth and fingers are not causes of lung cancer. Further, depending on the structure of the broader causal model in which a given cause-effect relation is embedded, data about non-causal correlates may be *more* useful to an agent than data about the causes of some variable of interest. Finally, given that the interventions used in philosophical analyses of causation are usually understood as counterfactual, a real-life insurance analyst never *could* learn that such an intervention had been performed. So it is worth asking whether it makes sense to use an information-pricing model to measure the pragmatic value of a causal explanation.

In response, I concede that the example of an insurance analyst deciding how much to pay for information about a potential customer is not a perfect analogy, for exactly the reasons given above. However, the analogy still does serve its purpose. My project in this paper is to provide a strategy for determining the deepest level of causal explanation in a given case. After determining that competing causal explanations at varying levels of granularity are all causal explanations of the same phenomenon, I then use an information-pricing analysis to settle the question of the optimal depth. This second step is not, in itself, uniquely suited to causal as opposed to non-causal models. However, my analysis takes for granted that we are already talking about causal explanations. For example, on my account we would first use an interventionist criterion of causation to determine that smoking, and not yellow teeth and fingers, is a cause of lung cancer. Then, we would use an FPCI-based criterion to determine the optimal level of granularity for describing a person's smoker status when explaining why they have lung cancer. Variables denoting the color of a person's teeth or fingers do not make it to this second stage of analysis, since they are not causes of lung cancer in the first place.⁷

Finally, there is some force to the anti-pragmatic notion that the better explanation of some phenomenon should never depend on the context-specific utilities of the agent

⁷Having said this, an interesting avenue for future research would be an exploration of the overlap between information-pricing analyses of explanatory models and existing work on the expected utility of interventions and observations in inferring causal graphs (see Eberhardt 2007, 170).

doing the explaining. To put it bluntly, it seems strange to think that life-insurance premiums should influence what counts as the best explanation of a patient’s lung cancer. However, attempts up to this point to give some purely non-pragmatic account of the value of explanations have proven unsuccessful. As Franklin-Hall (2016) points out, the space of possible variables that we can use to represent a given causal explanans is enormous, presenting reasons for pessimism that a single epistemic criterion could rule out all but the best variables for a given explanation in the interventionist framework. The abstraction criterion proposed by Weslake (2010) solves some problems, but runs into others, both because it allows disjunctive variables—as Weatherston (2012) and Clarke (2017) have pointed out—and in probabilistic cases, as I have pointed out here. The same goes for proportionality, in part for reasons described above, and in part for reasons described by Shapiro and Sober (2012) that I have not had space to address. Finally, Franklin-Hall (2016) argues that analyzing explanatory depth in terms of the stability of an explanation, i.e. its robustness in the face of changes to background conditions, also allows for putatively deep explanations that use disjunctive variables. The use of such variables seems out of keeping with scientific practice, and thus stability also fails as an epistemic criterion for explanatory depth.

In light of this failure to find an epistemic criterion for explanatory depth, I believe that we ought to take seriously the idea that the notion of superiority that we attach to some explanations is fundamentally an artifact of our perspective as pragmatically-motivated agents, rather than some observer-independent feature of nature. As such, the intrusion of pragmatics into judgments of explanatory depth may be inevitable, especially in threshold cases where different explanatory desiderata such as proportionality and abstractness pull in different directions. My proposal provides a precise way to take into account an agent’s prudential interests in order to give an account of how these interests determine an optimal granularity for causal explanations.

There is some precedent for this kind of approach in the literature. Andersen (2017) follows Dennett (1991) in arguing that causal relata are high-level “patterns” formed of lower-level “pixels”, where patterns are realized by pixels, but are robust over a certain amount of noise in their pixel-level realizations. For example, a bar code composed of individual pixels may cause a scanner to receive a certain message when the bar code is scanned. This causal relationship between the bar code and the scanner may be robust over many possible changes in the individual pixels comprising the bar code. The degree

of noise that we will tolerate while still picking out the bar code as an important causal relatum is set via pragmatic concerns; as long as the bar code does its job, we can speak coherently of the bar code playing an explanatory causal role with respect to the transfer of information, and not worry about the particular arrangement of pixels. That this account of higher-level causation is sensitive to pragmatic concerns is explicated by Andersen when she writes that “the more efficiently we describe a pattern, the faster we can identify whether it occurs. If there are tasks for which speed is relevant, we might prioritize efficiency of description and accept reduced accuracy” (2017, 603). My proposal here makes a similar point; our pragmatic concerns clearly play some role in the kinds of events that we pick out as explanatorily salient, and the level of granularity with which we pick them out.

7 Conclusion

This work has considered the way in which proportionality and abstractness provide an account of explanatory depth in deterministic cases. I then show how this account unravels in the probabilistic context. I propose a novel measure of the pragmatic value of an explanatory framework, and show how it addresses the problematic cases considered. I argue that this measure enables us to account for the intuition that some explanations are better than others in probabilistic contexts, albeit via the essential consideration of the pragmatic context of an explanation.

Appendix

General Definition of FPCI

Begin with set of acts $A = \{a_1, a_2, \dots, a_r\}$, a set of effect states $E = \{e_1, e_2, \dots, e_n\}$, and set of cause states $C = \{c_1, c_2, \dots, c_m\}$. Let $u(\cdot)$ be a utility function that takes as its arguments an action a_k and a value of the effect variable e_i . Let $p(\cdot)$ be a joint probability function over the values of E and C . We can define the joint probability

matrix $\mathbf{P}_{\mathbf{E},\mathbf{C}}$, utility matrix \mathbf{U} , and effect vector \vec{P}_E as follows.

$$\mathbf{P}_{\mathbf{E},\mathbf{C}} = \begin{bmatrix} p(e_1|\hat{c}_1)p(\hat{c}_1) & p(e_1|\hat{c}_2)p(\hat{c}_2) & \dots & p(e_1|\hat{c}_m)p(\hat{c}_m) \\ p(e_2|\hat{c}_1)p(\hat{c}_1) & p(e_2|\hat{c}_2)p(\hat{c}_2) & \dots & p(e_2|\hat{c}_m)p(\hat{c}_m) \\ \vdots & \vdots & \vdots & \vdots \\ p(e_n|\hat{c}_1)p(\hat{c}_1) & p(e_n|\hat{c}_2)p(\hat{c}_2) & \dots & p(e_n|\hat{c}_m)p(\hat{c}_m) \end{bmatrix} \quad (4)$$

$$\mathbf{U} = \begin{bmatrix} u(a_1, e_1) & u(a_1, e_2) & \dots & u(a_1, e_n) \\ u(a_2, e_1) & u(a_2, e_2) & \dots & u(a_2, e_n) \\ \vdots & \vdots & \vdots & \vdots \\ u(a_r, e_1) & u(a_r, e_2) & \dots & u(a_r, e_n) \end{bmatrix} \quad (5)$$

$$\vec{P}_E = \begin{bmatrix} p(e_1) \\ p(e_2) \\ \vdots \\ p(e_n) \end{bmatrix} \quad (6)$$

Let $\max(\cdot)$ be a function that takes a matrix as its argument, and returns a vector containing the maximum value in each column of that matrix, with vertical vectors treated as one-column matrices. Let $\text{sum}(\cdot)$ be a function that takes a vector as its argument and returns the sum of each element in the vector. We can define the fair price of causal information for a causal variable C , utility matrix \mathbf{U} , and effect variable E as follows.

$$\pi(\hat{C}, \mathbf{U}, E) = \text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}})) - \max(\mathbf{U} \cdot \vec{P}_E) \quad (7)$$

Proof of Proposition 1

Proof. Begin by writing the formula for the FPCI of causal variables C and C' :

$$\pi(\hat{C}, \mathbf{U}, E) = \text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}})) - \max(\mathbf{U} \cdot \vec{P}_E) \quad (8)$$

$$\pi(\hat{C}', \mathbf{U}, E) = \text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}'}) - \max(\mathbf{U} \cdot \vec{P}_E) \quad (9)$$

Next, we can show that in this case, $\text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}'}) \geq \text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}}))$. To do so, we begin by expanding the matrices $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}}$ and $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}'}$.

$$\mathbf{U} \cdot \mathbf{P}_{\mathbf{E},\mathbf{C}} = \begin{bmatrix} p(\hat{c}_1) \sum_{i=1}^n p(e_i|\hat{c}_1)u(a_1, e_i) & \dots & p(\hat{c}_m) \sum_{i=1}^n p(e_i|\hat{c}_m)u(a_1, e_i) \\ \vdots & \vdots & \vdots \\ p(\hat{c}_1) \sum_{i=1}^n p(e_i|\hat{c}_1)u(a_r, e_i) & \dots & p(\hat{c}_m) \sum_{i=1}^n p(e_i|\hat{c}_m)u(a_r, e_i) \end{bmatrix} \quad (10)$$

$$\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'} = \begin{bmatrix} p(\hat{c}'_1) \sum_{i=1}^n p(e_i | \hat{c}'_1) u(a_1, e_i) & \dots & p(\hat{c}'_q) \sum_{i=1}^n p(e_i | \hat{c}'_q) u(a_1, e_i) \\ \vdots & & \vdots \\ p(\hat{c}'_1) \sum_{i=1}^n p(e_i | \hat{c}'_1) u(a_r, e_i) & \dots & p(\hat{c}'_q) \sum_{i=1}^n p(e_i | \hat{c}'_q) u(a_r, e_i) \end{bmatrix} \quad (11)$$

Next, consider any value c_j such that $C = c_j$ if and only if C' takes a value in some set $\{c'_l, c'_{l+1}, \dots, c'_{l+z}\}$. We can define submatrices $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}}[j]$ and $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'}[l : l + z]$ as follows.

$$\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}}[j] = \begin{bmatrix} p(\hat{c}_j) \sum_{i=1}^n p(e_i | \hat{c}_j) u(a_1, e_i) \\ \vdots \\ p(\hat{c}_j) \sum_{i=1}^n p(e_i | \hat{c}_j) u(a_r, e_i) \end{bmatrix} \quad (12)$$

$$\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'}[l : l + z] = \begin{bmatrix} p(\hat{c}'_l) \sum_{i=1}^n p(e_i | \hat{c}'_l) u(a_1, e_i) & \dots & p(\hat{c}'_{l+z}) \sum_{i=1}^n p(e_i | \hat{c}'_{l+z}) u(a_1, e_i) \\ \vdots & & \vdots \\ p(\hat{c}'_l) \sum_{i=1}^n p(e_i | \hat{c}'_l) u(a_r, e_i) & \dots & p(\hat{c}'_{l+z}) \sum_{i=1}^n p(e_i | \hat{c}'_{l+z}) u(a_r, e_i) \end{bmatrix} \quad (13)$$

The sum of every element in a given row of $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'}[l : l + z]$ equals the element in the corresponding row of $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}}[j]$. If $p(e_i | \hat{c}'_l) = p(e_i | \hat{c}_j)$ for all sets of values e_i , c_j and c'_l such that $C' = c'_l$ implies $C = c_j$, then $\sum_{i=1}^n p(e_i | \hat{c}'_{l+v}) u(a_k, e_i) = \sum_{i=1}^n p(e_i | \hat{c}_j) u(a_k, e_i)$ for any a_k . This implies that the submatrix $\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'}[l : l + z]$ is row-dominated. It follows from this that $\text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'}[l : l + z])) = \text{sum}(\max(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}}[j]))$ for any c_j and any set $\{c'_l, c'_{l+1}, \dots, c'_{l+z}\}$ such that $\{c'_l, c'_{l+1}, \dots, c'_{l+z}\}$ is the union of all values of C' that imply that $C = c_j$. It follows from this that $\max(\text{sum}(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}'})) = \max(\text{sum}(\mathbf{U} \cdot \mathbf{P}_{\mathbf{E}, \mathbf{C}}))$. Since $\max(\mathbf{U} \cdot \vec{P}_E)$ is unchanged by refining the causal variable, this implies that $\pi(\hat{C}, \mathbf{U}, E) = \pi(\hat{C}', \mathbf{U}, E)$. \square

References

- H. Andersen. Patterns, information, and causation. *Journal of Philosophy*, 114(11): 592–622, 2017.
- R. W. Batterman. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press, 2001.
- D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.
- C. Clarke. How to define levels of explanation and evaluate their indispensability. *Synthese*, 194(6):2211–2231, 2017.
- C. F. Craver. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press, Clarendon Press, 2007.
- D. C. Dennett. Real patterns. *Journal of Philosophy*, 88(1):27–51, 1991.
- F. Eberhardt. Causation and intervention. *Unpublished doctoral dissertation, Carnegie Mellon University*, 2007.
- L. Fenton-Glynn. A proposed probabilistic extension of the halpern and pearl definition of cause’. *British Journal for the Philosophy of Science*, page axv056, 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.
- A. Garfinkel. Forms of explanation: Rethinking the questions in social theory. *Philosophical Review*, 93(1):116–118, 1981.
- I. J. Good. On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4):319–321, 1967.
- C. Hitchcock. The role of contrast in causal and explanatory claims. *Synthese*, 107(3): 395–419, 1996.
- C. Hitchcock. Contrastive explanation and the demons of determinism. *British Journal for the Philosophy of Science*, 50(4):585–612, 1999.

- C. Hitchcock and J. Woodward. Explanatory generalizations, part ii: Plumbing explanatory depth. *Noûs*, 37(2):181–199, 2003.
- F. Jackson and P. Pettit. In defense of explanatory ecumenism. *Economics and Philosophy*, 8(1):1–21, 1992.
- P. Kitcher. Explanatory unification. *Philosophy of Science*, 48(4):507–531, 1981.
- D. Lewis. Causal explanation. *Philosophical papers*, 2:214–40, 1986.
- P. Lipton. Contrastive explanation. *Royal Institute of Philosophy Supplement*, 27:247–266, 1990.
- C. List and P. Menzies. Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy*, 106(9):475–502, 2009.
- C. Malaterre. Making sense of downward causation in manipulationism (with illustrations from cancer research). *History and Philosophy of the Life Sciences*, (33):537–562, 2011.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- P. Percival. Lewis’s dilemma of explanation under indeterminism exposed and resolved. *Mind*, 109(433):39–66, 2000.
- A. Pocheville, P. Griffiths, and K. Stotz. Comparing causes: an information-theoretic approach to specificity, proportionality and stability. In H. Leitgeb, I. Niiniluoto, E. Sober, and S. Paivi, editors, *Proceedings of the 15th Congress of Logic, Methodology and Philosophy of Science*, pages 93–102. London: College Publications, 2017.
- H. Putnam. *Philosophical Papers: Volume 2, Mind, Language and Reality*. Cambridge University Press, 1979.
- M. Resnik. *Choices: An Introduction to Decision Theory*. Univ of Minnesota Press, 1987.
- L. J. Savage. *The Foundations of Statistics*. Wiley Publications in Statistics, 1954.
- C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.

- L. Shapiro and E. Sober. Against proportionality. *Analysis*, 72(1):89–93, 2012.
- L. Sklar. *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge University Press, 1993.
- U. Stegmann. Causal control and genetic causation. *Nous*, 48(3):450–465, 2014.
- M. Strevens. *Depth: An Account of Scientific Explanation*. Harvard University Press, 2008.
- B. Weatherson. Explanation, idealisation and the goldilocks problem. *Philosophy and Phenomenological Research*, 84(2):461–473, 2012.
- B. Weslake. Explanatory depth. *Philosophy of Science*, 77(2):273–294, 2010.
- B. Weslake. Proportionality, contrast and explanation. *Australasian Journal of Philosophy*, 91(4):785–797, 2013.
- J. Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, 2003.
- J. Woodward. Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology and Philosophy*, 25(3):287–318, 2010.
- J. Woodward. The problem of variable choice. *Synthese*, 193(4):1047–1072, 2016.
- J. Woodward. Response to franklin-hall and weslake on stability and proportionality, unpublished.
- S. Yablo. Mental causation. *Philosophical Review*, 101(2):245–280, 1992.
- P. Ylikoski and J. Kuorikoski. Dissecting explanatory power. *Philosophical Studies*, 148(2):201–219, 2010.