

Causal control: A rationale for causal selection

Lauren N. Ross

Abstract

Causal selection has to do with the distinction we make between background conditions and “the” true cause or causes of some outcome of interest. A longstanding consensus in philosophy views causal selection as lacking any objective rationale and as guided, instead, by arbitrary, pragmatic, and non-scientific considerations. I argue against this position in the context of causal selection for disease traits. In this domain, causes are selected on the basis of the type of causal control they exhibit over a disease of interest. My analysis clarifies the principled rationale that guides this selection and how it involves both pragmatic and objective considerations, which have been overlooked in the extant literature.

1 Introduction. Causal selection has to do with the distinction we make between background conditions and “the” true cause or causes of some outcome of interest. We claim that “the” cause of a match lighting was the fact that it was struck, but not the presence of oxygen. A car crash may involve many causal factors—ice on the road, the speed of the car, the gas in the tank, and so on—but we are likely to explain it by citing few of these. In each case, just a small number of factors are selected as causes of the outcome, while most others are backgrounded. A longstanding consensus in philosophy views causal selection as lacking any objective rationale and as guided, instead, by arbitrary, pragmatic, and non-scientific considerations. This position is famously supported by John Stuart Mill, whose arguments have “won the field” and remain “echoed by contemporary authors” (Schaffer 2014). According to Mill, “Nothing can better show the absence of any scientific ground for the distinction between the cause of a phenomenon and its conditions, than the capricious manner in which we select from among the conditions that which we choose to denominate the cause” (Mill 1874, 238). Lewis (1973) agrees with this, claiming that we select causes because they are under our control, because we find them good or bad, or just because we want to talk about them. He states, “We sometimes single out one among all the causes of some event and call it ‘the’ cause, as if there were no others. Or we single out a few as the ‘causes,’ calling the rest mere ‘causal factors’ or ‘causal conditions’... I have nothing to say about these principles of invidious discrimination” (Lewis 1973, 558-559).

This position faces significant problems in the context of biomedicine, where scientists commonly identify “the” cause or causes of specific diseases. Scientists claim that the cause of scurvy is vitamin C deficiency, that the cause of tuberculosis is the tubercle bacterium, and that the cause of Huntington’s disease is a mutation in the *huntingtin* gene. There is widespread consensus on the selection of these causes in

[†]Acknowledgements: I would like to thank Jim Woodward, Ken Schaffner, Zina Ward, and Marc Ereshefsky for helpful comments on this paper. I would also like to thank the research group associated with the University of Calgary—including Brian Hanley, Marie Kaiser, Oliver Lean, Derek Turner, and Ken Waters—for valuable feedback and discussion.

the medical community. Furthermore, at first glance, these causes appear importantly relevant to their respective diseases in ways linked to possibilities of control. Targeting these causal factors has allowed for successful prevention and treatment of disease, to the extent of drastically reducing the incidence of some diseases and nearly eradicating others from the human population. If causal selection is arbitrary, why is there such widespread consensus on the selection of causes for some disease traits? What explains the apparent success of this selection in identifying factors that reliably allow for control over disease? Finally, if causal selection is unscientific, why do we view disease explanation as a scientific matter? These points identify significant problems for the mainstream philosophical position on causal selection and its ability to account for the selection of disease causes in biomedicine.

In this paper I argue that the rationale behind causal selection for disease is best understood in terms of the *causal control* that selected causes have over the disease of interest. I provide a novel account of the types of control that guide this process. In this domain, causes are selected on the basis of having (i) broad scope, (ii) probable, and (iii) stable control over disease. I suggest that this selection is pragmatic in the sense that it is relative to a *practical* goal—the goal of control. This goal is practical in the sense that it is useful, in general, for our navigation and operation within the world and, more specifically, for our practical aims of treating and preventing disease in patients. This is a different notion of “pragmatic” than is often found in the philosophical literature, where it is used to refer to considerations that are arbitrary, subjective and/or audience-relative (Achinstein 1984; Schaffer 2014). Furthermore, I suggest that causal selection for disease is also “objective” in the sense that once the goal of control is specified, there are objective facts and considerations about what means conduce to this goal. This paper provides an analysis of causal selection that clarifies these considerations.

I characterize causal selection as involving three main steps. The first step involves setting a contrastive focus, while the second and third steps involve selecting causes with respect to this focus. In order for a factor to be selected as “the” cause or one of “the” causes of an effect, it must pass both the second and third steps. Once the (1) contrastive focus is set, the remaining two steps involve selecting causes on the basis of the causal control they have over this focus. The second step involves (2) assessing whether a factor has *any* control over the effect of interest, and the third step involves (3) assessing what *type* of control the factor has over the effect. I clarify three types of control that guide this third and final selection step, which include control that is (i) specific, (ii) probable, and (iii) stable for the disease of interest. It makes sense that we select causes on these grounds, because they meet particular standards we have for disease explanation and they provide valuable targets for disease treatment and prevention. An important feature of my analysis is that it clarifies how unique features of disease traits make causal selection for disease much different, and arguably much easier, than causal selection for many other types of phenomena. One reason for this, is that scientists often impose constraints on what qualifies as a “legitimate” or “valid” disease category, where these constraints alone significantly narrow the number of candidate causes for any given disease. Differences between disease examples and the examples commonly examined in the philosophical literature on causal selection can help clarify why mainstream philosophical views have overlooked the principled

rationale that guides causal selection for disease.

This paper is organized as follows. Section 2 discusses important features of disease traits and how they set the relevant contrastive focus. With this contrastive focus specified, I provide a minimal criterion for causal control that captures the second step of causal selection. This criterion draws on Woodward’s (2003) interventionist account of causation and it identifies “interventionist causes” for the contrastive focus. Section 3 examines the third and final selection step, where particular factors are selected from a pool of candidate causes. I clarify the types of control that guide the selection of causes in this final step. Throughout this analysis I clarify how causal selection for disease is influenced by both pragmatic and objective considerations and I return to this topic in the the concluding section.

2 Disease traits and interventionist causes. Human diseases are a unique type of biological phenomenon. Modern medicine aims to define disease traits causally, in terms of particular causal etiologies. While there is a sense in which all diseases are produced by a multitude of causal factors, often very few factors are selected as “the” cause or one of “the” causes of any given disease. For “monocausal” diseases we cite single factors, as in the cases of scurvy, tuberculosis, and Huntington’s disease. For other diseases we cite longer lists of causes. For example phenylketouria (PKU) is explained by appealing to both a gene mutation and a dietary factor.¹ When little is known about the causal etiology of a clinically accepted disease category—as is common with psychiatric and other disorders—the “legitimacy” and “validity” of the disease is often disputed or, at least, considered an open question.² In this paper, I focus on diseases for which we have some sufficiently understood causal etiology. These are the cases where we commonly select disease causes and where there is consensus on this selection.³

Human diseases, even those with well understood or simple causal etiologies, often involve many symptoms. For example, although tuberculosis is explained by appealing to a single bacterial factor, patients with this disease often present with a wide variety of symptoms. These patients can exhibit symptoms that include: dry cough, blood-tinged sputum, night sweats, weight loss, and fatigue, to name a few. Although we distinguish among these symptoms we do not view them as individual diseases, but as features of a single disease process. Additionally, symptomatic presentation can vary significantly across patients with the same disease. Patients with the same disease may present with completely different combinations of symptoms or with similar symptoms that vary in degree. For an example of the former, one patient with tuberculosis may

¹PKU is characterized by disordered metabolism of the amino acid phenylalanine, which results in impaired neurological development. This disease is caused by both the ingestion of phenylalanine and a mutation in the gene for the phenylalanine hydroxylase enzyme.

²The worry is not that the symptoms and signs of the disorder are not real (in the sense of being experienced or exhibited by the patient), but that the disorder category will change with further clarification of the etiology (Kincaid and Sullivan 2014). For more on this see Schaffner’s discussion of validity and etiopathological validity in the context of psychiatry and general biomedicine (Schaffner 2012).

³If we know little about the causal etiology of a clinically accepted disease category, we typically do not consider ourselves in a position to identify its causes.

exhibit all of the symptoms mentioned above, while another may present with only a dry cough. In the case of the latter, two patients with tuberculosis may present with a dry cough, but the severity of their cough may differ.

Tuberculosis has a relatively simple causal etiology in the sense that it is explained by appealing to a single causal factor. It also has a specific causal etiology, in the sense that all cases of this disease have the same cause. What this reveals is that a disease can have a simple and specific causal etiology, *without* this implying that it always presents in patients with a single, uniform symptomatology. Unfortunately, the view that particular diseases manifest in patients with uniform symptomatology is common in the philosophical literature.⁴ This view incorrectly posits uniform symptomatology as a kind of “standard” for medically accepted disease categories. If distinct diseases always presented in such a uniform manner, diagnosis in the clinical setting would be much easier and straightforward than it is. Instead, patients with the same disease often present with highly divergent symptom profiles and clinicians are trained to diagnose in light of this challenge.

While we do not expect most diseases to have *simple* causal etiologies, we often assume that they have *specific* causal etiologies. In this sense, to say that disease D has a specific causal etiology means that all instances of D are produced by roughly the same causal factors.⁵ (I discuss this further in section 3). As the notion of specific causal etiology pertains to many instances of a given disease—as opposed to a single or token instance—it represents a type-level consideration. This type-level focus is present in many of our claims about what “the” cause of a particular disease is. In this paper, I focus on causal selection for type-level disease traits. This can be thought of as causal selection that focuses on answering the following question: “What is the cause of disease D in the human population?” When we answer this question we often focus on the binary contrast of disease “absence” and “presence.” Of course, a patient either has disease D or she does not—she cannot be in both of these states or in neither of them. We expect disease causes to explain *at least* this contrast, in part because it is often the ultimate contrast we want control over. If a causal factor only explains positive degrees of disease pathology, without also explaining disease absence, we view it as an incomplete and unsatisfying explanation. Our interest in completely curing and preventing disease—i.e. ensuring the complete absence of disease—motivates our interest in identifying causes of this binary contrast. If these causes also account for varying states of a disease—in addition to disease absence—we view this as an added advantage, but not a necessary criterion for the factors we select as disease causes.

The first step of causal selection involves setting a contrastive focus. So far, I have identified two important features of disease traits that set this focus: disease traits are (1) type level phenomena, which (2) are often represented as taking on the values “present” or “absent.” In this paper, I argue that the rationale behind causal selection

⁴For examples of this, see: (Poland 2014; Blaxter 2015; Kincaid and Sullivan 2014; Murphy 2014).

⁵This is *not* the same as claiming that the causal factors in question *only* produce disease D and not other diseases—this refers to the specificity of *effects*, given some cause. The notion of specific causal etiology I discuss involves the specificity of *causes* given some effect (a particular disease). These two types of specificity may be distinguished as (2.1) specificity of effect (given some cause) and (2.2) specificity of cause (given some effect).

for disease is best understood in terms of the *causal control* of factors over this contrastive focus. In the second step of causal selection, a factor is selected as a candidate cause if it has *some* control, or some minimal amount of control, over this contrast. What does it mean for a factor to have *some* causal control? Consider a minimal interventionist criterion for causal control, which is met by factors I call “interventionist causes” and inspired by Woodward (2003):

(i.c.) interventionist cause: a factor C has causal control over disease D if and only if there are circumstances S such that if some (single) intervention that changes the value of C (and no other variable) were to occur in S, then the value of D or the probability distribution of D would change, *for the contrastive focus in question*.

The notion of an intervention ensures that when variable C is manipulated, it allows for a change in the value of D in a way that excludes confounders or other variables that may causally influence D. One advantage of the interventionist framework is its ability to capture the motivation behind some of our experimental methods for identifying causal relationships. If we want to determine whether substance X causes disease Y, we might design an experiment in a model organism where we manipulate values of X (and only X) to see if this brings about a change in the occurrence of disease Y. In this experiment, we are likely to keep potential confounding factors constant (e.g. diet, exercise, etc.) in order to ensure that changes in X—and not changes in these other factors—cause the outcome we measure.⁶ The interventionist criterion (i.c.) involves a counterfactual claim: it maintains that C has causal control over D in the sense that *if* there was a change in C, this *would* produce a change in D. This criterion does not require that such an intervention on C is actually performed, in the sense of being manipulated or even manipulable with current technology or by human means. This makes sense of the fact that we view gene variants as the cause of many human diseases, despite our inability to intervene on them in human patients. Technological limitations and ethical restraint prevent us from manipulating these gene variants in humans. However, we still maintain that some genes cause particular diseases, in the counterfactual sense indicated: we mean that *if* such gene variants were manipulated, this *would* change the disease status of the subject.⁷ We support this claim on the basis of evidence acquired from various sources and not just from actually performing the intervention in question. My analysis relies on a notion of causal control that is counterfactual in the same sense.

In order for a factor to be selected as a disease cause, it must meet (i.c.) and have *some* causal control over the contrastive focus. This is the second step and

⁶For more of these details, see (Woodward 2003).

⁷This highlights important differences between my position and Gannet’s analysis of pragmatic considerations that influence genetic explanation (Gannet 1999). She views causal selection as guided by our success with *actual* manipulation of candidate causes. I think her position neglects the importance of counterfactual (or hypothetical) information in causal selection. Acknowledging the importance of this information helps explain why we cite genes and other factors as disease causes, despite our limitations in actually manipulating them. Furthermore, Gannet and I both claim that causal selection in biology is pragmatic, but we disagree about why. I suggest that causal selection is pragmatic in the sense that it relates to the practical goal of control, as opposed to our ability to actually manipulate causal factors.

the first “cut” of the causal selection process. Notice that the causes we select for scurvy, tuberculosis, and Huntington’s disease all meet this interventionist criterion (i.c.). Dietary vitamin C has causal control over scurvy in the sense that manipulating levels of this dietary factor provides control over the occurrence and nonoccurrence of the disease.

It might seem that the (i.c.) criterion is overly inclusive, in the sense that it picks out a huge number of factors. Recall the car crash example, for which numerous factors were identified as causally relevant. Lewis claims that the causes of the crash include: the ice on the road, the blind corner, and even the birth of the driver’s paternal grandmother (Lewis 1986, 215-6). If all these factors meet the (i.c.) criterion, it might not seem like a very helpful first cut. The ordinary life examples discussed in the philosophical literature are different from disease examples, in a way that can misrepresent causal selection in biomedicine. Disease traits are very narrow phenomena, both by our choosing and for reasons that have to do with their manifestation in living organisms. First, as we often define disease traits in terms of specific causal etiologies, our own characterization of them significantly constrains their relevant causal factors from the outset. (I discuss this more in section 3). Second, disease phenotypes are narrow in the sense that their causes are expected to control the disease absence/presence contrast, and only this contrast, in living human patients. This is a very narrow and fragile contrast for a factor to have control over, in a way that is not characteristic of many non-biological examples. If we want to prevent a match from lighting, or put it out once lit, there are many ways to do this. We could chop the match into unrecognizable pieces, pour corrosive chemicals on it, or throw it in the ocean. There are far fewer candidate causes for disease, because we expect such factors to have control over disease traits without killing or harming the patient. We can destroy the match to prevent it from lighting, but we do not want to destroy the patient to eliminate disease. There is a sense in which there is no disease in a dead patient, but clearly this is not the type of control that we want, or expect, disease causes to have. These features of disease traits significantly reduce the number of candidate causes that we consider in our search for “the” cause or causes of a given disease.

Once the contrastive focus of disease presence/absence is specified, the (i.c.) criterion clarifies “objective” considerations that guide selection. Given the goal of control, there are objective facts about which factors have control over this focus and which factors do not. Consider the disease scurvy again, but now with regard to the candidate causes “oxygen” and “dietary vitamin C.” “Oxygen” does not meet the (i.c.) criterion for this disease, because manipulating oxygen does not allow for control over the “absence” and “presence” of scurvy in *living* human patients. Of course, manipulating oxygen in certain ways can kill a patient—this would happen if oxygen levels were set to a very low value or if oxygen were completely removed from the patient’s environment. Manipulating “oxygen” does have some control over whether a patient lives or dies, but this is *not* the dominant contrastive focus for disease explanation. Instead, the dominant focus is disease absence/presence in *living* patients. “Dietary vitamin C” meets the (i.c.) criterion for this contrast in the case of scurvy, while “oxygen” does not. This explains why we cite dietary vitamin C, and not oxygen, as “the” cause of scurvy. However, while “oxygen” does not meet the (i.c.) criterion for scurvy there

is still a sense in which it is relevant or “necessary” for the incidence of this disease. “Oxygen” is relevant to this disease in the sense that its presence is required for the causal control that vitamin C levels have over scurvy. This is due to the fact that oxygen is a requirement for human life and thus, it is required for attaining both values of the contrast in question. This clarifies how we can view a factor like oxygen as importantly relevant or necessary for disease, while still distinguishing it from factors we select as “the” cause or causes of disease.

Thus, if a factor lacks causal control over the contrastive focus for a disease trait it fails to pass the second causal selection step. This reveals part of the rationale behind causal selection for disease traits and how the medical community can reach consensus on “the” causes of some diseases. This consensus is partly explained by our expectation that disease causes should have some causal control over the disease of interest. The interventionist criterion (i.c.) captures this basic standard and the second step of the causal selection process. In the next section I discuss the third step of this process, where we select among candidate interventionist causes for a given disease.

3 Selecting among interventionist causes. In order for a factor to be selected as “the” cause or one of “the” causes of a disease, meeting the interventionist criterion (i.c.) is necessary, but not sufficient. Some factors meet (i.c.) for particular diseases without being selected as “the” cause of the disease. For example, consider sleep deprivation and the seasonal flu. Changing sleep duration, such that a patient is sleep deprived, has some causal control over the flu in the sense that it can increase susceptibility to infection by decreasing immune functioning (Bryant, Trinder, and Curtis 2004). However, we do not consider sleep deprivation “the” cause (or even one of “the” causes) of the seasonal flu. We reserve this designation for the particular flu virus. What explains this selection of one interventionist cause over another? Notice that targeting the flu virus (e.g. with vaccination⁸) provides a very different type of control over the occurrence of the flu than does targeting patients’ amounts of sleep. Targeting the virus provides (3.1) control over many to all cases of this particular flu, (3.2) a high likelihood of preventing it in each case, and (3.3) control across a wide variety of genetic and environmental conditions present in the patient population. Compared to the flu virus, targeting levels of sleep significantly underperforms in all of these areas. I discuss these three types of causal control—which I refer to as causal control that is (3.1) specific, (3.2) probable, and (3.3) stable—and I clarify how they guide our selection of disease causes.

3.1 Specific causes and causal control of broad scope. The factors we select as disease causes exhibit very particular types of causal control over disease. The first type of causal control I discuss relates to the assumption of specific causal etiology. This assumption captures our default view that type-level disease traits have specific causes. In other words, for a given disease D we often expect that most or all instances of D are the result of a similar causal process. In the philosophical literature, the notion of specific causal etiology for disease is often referred to as the “causal signature”

⁸This is the vaccine that is recommended annually, for everyone 6 months and older.

(Murphy 2014, 105), “disorder-specific pathophysiology” (Caspi and Moffitt 2006, 586), or “shared causal process” (Zachar 2014, 87) for a disease trait. Sometimes we refer to single factors as the specific causal etiology for a disease, like the *huntingtin* gene mutation for Huntington’s disease. Other times we refer to multiple interacting factors as the specific causal etiology for a disease, like the gene variant and dietary factor for the disease PKU. The important feature of specific causal etiology is not how many factors cause an instance of disease D, but that the *same* factors produce all or most instances of disease D.⁹

Just because we expect diseases to have specific causal etiologies, does not mean that all clinically useful disease categories meet this standard. Sometimes we start with clinically useful categories—like Parkinson’s disease—only to find out later, that these categories fail to track specific causal processes. (This is unsurprising, because we often create these categories before we have clear information about their causes.) For example, our best evidence suggests that Parkinson’s disease has a *non-specific* causal etiology, in the sense that (3.1.1) different combinations of causal factors cause instances of the disease on different occasions. This contrasts with a situation of *specific* causal etiology where (3.1.2) several factors interact to produce a disease, but where every case of the disease is produced by the same interacting factors. An example of (3.1.2) is PKU, because the same two interacting factors cause every instance of the disease.

What does the assumption of specific causal etiology have to do with selecting disease causes and causal control? This assumption clarifies our aim of selecting causal factors that are specific for a disease, in part because these factors are likely to provide control of *broad* scope over all or most instances of the disease in question. Consider the situations above, where a disease D has either (3.1.1) a non-specific causal etiology or (3.1.2) a specific causal etiology. One advantage of selecting specific causes is that they can often be manipulated to control all or many cases of the population-wide disease of interest. This is because, in situation (3.1.2), all cases of the disease are produced by the same causal factors. Alternatively, if we select causes that are non-specific for a disease, as in the case of (3.1.1), these factors are likely to have causal control of narrow scope, in the sense that they influence a smaller percentage of the total cases of disease D. This is because, in situation (3.1.1), different causal factors produce different instances of the same disease D, so targeting the factors that produce any one instance of the disease, is unlikely to provide control over the other instances that have different causes. From the standpoint of treating and preventing disease, identifying factors that are causally specific for particular disease traits is extremely valuable. It often identifies factors that we can target to control and explain a large percentage of all cases of a given disease in the population.

The notion of specific causal etiology plays a more complicated role in causal selection than I have indicated so far. One complication is that this notion influences both the factors we select as disease causes *and* those phenomena that we consider

⁹In this paper, I use a notion of specificity that refers to the number of cause and effect *variables* that participate in a type-level causal relationship, as opposed to the *values* of these variables. The latter (*value*) sense of specificity has received much more attention in philosophy of biology (Woodward 2010; Waters 2007; Griffiths 2006). For more on this distinction, see (Woodward 2015).

to be “legitimate” diseases to begin with. If a disease category has no known causal etiology, or if it has a non-specific causal etiology, the “legitimacy” of the category and whether it represents a “true” disease are often called into question. This is helpfully illustrated with Parkinson’s disease, which we understand as having a non-specific causal etiology. As mentioned above, current research suggests that different combinations of causal factors produce different cases of Parkinson’s disease. Researchers view this causal non-specificity as having “splintered the unitary conception of this disease” and as “challenging traditional conceptual frameworks” for understanding Parkinson’s disease (Shulman, De Jager, and Feany 2011, 214, 193). In this situation medical researchers suggest either (1) continuing the search for some shared causal process (i.e. some specific causal etiology) or (2) dividing up the disease category on the basis of the distinct causal processes. Both of these options restore causal specificity by either (1) finding it for the pre-established disease category or (2) creating it by redefining the disease category. This reveals how causal specificity is both a guiding rationale for causal selection and a standard that influences how we define disease traits. This suggests that causal selection is more of a back-and-forth process than just a search for causes given a fixed contrastive focus. We may start with a disease category and search for its causes, only to redefine the category on the basis of what we find.

3.2 Probable causal control. A second type of causal control that guides causal selection for disease is what I call *probable* causal control. This causal control refers to the *probability* with which each outcome of the contrastive focus is produced when selected factors are manipulated. Consider a light switch C, which can take the values ‘up’ or ‘down’ and a light E, which can take the values “on” or “off.” In the first case, turning the switch “up” results in a 99% probability of the light bulb being “on” and turning the switch “down” results in a 99% probability of the light bulb being “off.” In a second scenario, turning the switch “up” has only a 60% probability of causing the light to turn “on” and turning the light switch “down” only has a 60% probability of turning the light “off.” In both cases the switches have some causal control over the state of the light, but their control differs with regard to how *probable* each outcome of the contrast is with interventions on the switch.¹⁰

When we select disease causes we prioritize factors that have a high degree of probable causal control over disease. By targeting these factors, as opposed to factors with less probable causal control, we increase the likelihood of getting a particular outcome. This has a clear advantage for our treatment and prevention measures and for explaining disease outcomes. If we want to control whether a light is “on” or “off” we will prefer the first switch-bulb system to the second, because we are more likely to get the outcome we want by manipulating the switch. In cases where we reach consensus on disease causes, they often have probable causal control over the disease trait in question. Consider diseases like scurvy, tuberculosis, and Huntington’s disease. In each of these cases, when the disease cause is present (or properly introduced) in a patient, her likelihood of acquiring the disease approaches 100%. Similarly, when the cause is absent (or properly avoided) the likelihood of disease absence also approaches 100%.

¹⁰This notion of probable causal control shares similarities with the suggestion by Lu et al. that we look for causes with high “power” (Lu, Yuille, Liljeholm, Cheng, and Holyoak 2008).

A patient with the *huntingtin* gene mutation is almost certain to acquire Huntington's disease and very unlikely to get this disease without it (cases of this disease without the mutation are unheard of). Alternatively, consider our attempts at causal selection for schizophrenia, which is a psychiatric disorder associated with a multitude of causally relevant gene variants. In this case, no single variant (or set of variants) confers a high probability of disease occurrence and nonoccurrence, although some provide a low probability of this sort. Failure to meet the standard of probable causal control partly explains why the medical community does not view such variants as "the" cause of this disease and why the search to better understand its etiology continues.

When a single causal factor provides a low degree of probable causal control over a disease trait, we often search for interacting causes that increase this type of control. Consider PKU again, which we explain by citing two causal factors: a gene variant and a dietary factor. These are interacting causes for this disease, because they both meet the (i.c.) criterion and they each influence the causal control that the other has over the disease. The gene variant only causes PKU when the dietary factor is present, and vice versa. One reason for selecting both of these causes in explaining PKU is that together they provide more probable causal control over the disease than a single factor alone. Of course, PKU is a relatively simple disease in the sense that we explain it by citing only two interacting causes. One significant challenge associated with disease explanation is that many diseases appear to be causally complex in the sense that they have a multitude of interacting causes, which must all be accounted for to provide a high degree of probable causal control over the disease. This can make disease explanation and causal selection much more difficult, because a larger number of causal factors have to be identified.

To say that causes with probable causal control are privileged in causal selection is different from claiming that causation requires this type of control. Some probability-raising accounts of causation support this later position, by maintaining that causes are factors that result in a high probability of their effects.¹¹ A well known objection to these accounts cites the low probability of general paresis, or late state neurosyphilis, among untreated patients with syphilis (Scriven 1959). Approximately one-third of patients with untreated syphilis end up with general paresis, yet we still view the syphilis bacterium as "the" cause of this low probability outcome. Probability-raising accounts struggle to make sense of why we explain general paresis by citing the syphilis bacterium, since the cause confers a low probability of the occurrence of the effect, compared to a causal factor that approaches a 100% likelihood of producing an outcome. My analysis clarifies this confusion and explains why this example is not problematic for my position. First, we more often view general paresis as a set of *symptoms* produced by the disease syphilis, as opposed to a distinct *disease* itself. It is true that we view the syphilis bacterium as the cause of general paresis, but we rarely expect disease causes to have specific or probable causal control over disease *symptoms*. This is because individual symptoms can be found in many different diseases (and thus, have many different causes) and they can have variable presentation across cases of a partic-

¹¹Hempel's inductive statistical (IS) account of explanation has been viewed as supporting this position (Hempel 1965).

ular disease (where there is, presumably, the same causal etiology).¹² When we focus on the proper disease target “syphilis,” we do identify the syphilis bacterium as “the” cause of this disease. Our reasons for doing this include the facts that this bacterial cause exhibits specific and probable causal control over the disease trait. Furthermore, my analysis clarifies what features of causal relationships we privilege in some contexts, but not what features of relationships make them causal. To say that factors with less probable causal control are not privileged in causal selection is not to deny that they are still causal relationships.

3.3 Stable causal control. A final type of causal control I discuss is causal control that is stable. Stability is a feature of causal relationships that has been discussed extensively by Woodward (2003, 2006, 2010). Recall that the interventionist criterion (i.c.) refers to “circumstances S” in which a cause C has causal control over an effect D. Stability refers to the extent to which the causal control of C over D holds in a range of other circumstances S_i , which differ from circumstances S (Woodward 2010). Consider a broad set of conditions, which include the various genetic backgrounds of the current human population and the range of environmental surroundings they live in.¹³ If a cause C only has causal control over effect D in a very narrow range of these circumstances, this causal control is *unstable*. An example of this would be a situation where manipulating dietary vitamin C levels only controlled scurvy incidence in people with a gene variant for brown hair, who also live in the state of Florida.¹⁴ In this case, the causal control of C over D only holds in patients with a narrow range of genetic and environmental conditions, relative to all patients and environments in the world. Alternatively, if C has causal control over D in a very broad range of these circumstances, this causal control is considered *stable*. An example of this would be if manipulating dietary vitamin C levels controlled scurvy incidence in patients of a wide variety of genetic backgrounds, who live in many different environments. The second situation of stable causal control is clearly more useful for the purposes of treatment and prevention of disease—it allows for measures that have the potential to treat a larger number of patients who live in many different types of environments. The causes we select for scurvy and tuberculosis have stable causal control in the sense captured in the second scenario. We can target these causes to treat and prevent these diseases in patients with diverse genetic backgrounds, who live in a wide variety of environments.

There is an important complication involved in assessing the stability of causal relationships. The degree to which a cause variable has stable (or unstable) causal control depends on the range of circumstances S_i that are considered. Clarifying exactly how broad or narrow these circumstances are is difficult, because this determination often seems context-dependent. This is acknowledged by Woodward, who states that the range of circumstances that matter for assessing stability are those that “do not depart too much from the actual state of affairs or that do not seem too far-fetched or that are not judged to be unimportant or irrelevant for subject-matter-specific reasons” (Woodward 2006, 11).

¹²I discuss these points in section 2.

¹³For more on causal selection in the context of population-wide outcomes, see Waters (2007).

¹⁴This is not a true scenario. It is intended to clarify the relevant sense of unstable causal control.

What are the circumstances S_i that matter for determining stability in the context of disease traits? First, there is an important sense in which stability is relative to particular reference classes, which group patients on the basis of factors like age group and sex, depending on the disease of interest.¹⁵ For example, given a virus that causes cervical cancer, we do not assess the stability of this cause with respect to all patients, but only with respect to patients of the female sex. The same can be said for pediatric, geriatric, and other diseases, which are present in particular patient populations.¹⁶ When diseases are restricted to particular populations, we often assess the stability of causal relationships with respect to this restriction. Second, as Woodward has indicated, the circumstances that matter for assessing stability are those circumstances that *actually* occur in the contexts we are interested in. In the context of disease, we care about the range of conditions that are *actually* present in human patients and their environments in the world. We care about these circumstances because they are the circumstances in which we want to control disease. Whether a factor meets this standard depends on the time-frame of interest. Future changes in our genetic make-up and the environments we live in can alter the circumstances S_i that we use in assessing the stability of causal factors for disease. Third, we restrict the conditions included in circumstances S_i to those that include basic biological requirements for human life. Just because there are some low oxygen environments on Earth, does not mean that we assess the stability of disease causes with such circumstances in mind. All disease causes break down in this type of environment, because it is incompatible with human life. Since we cannot sustain human life in these circumstances, we are not focused on controlling disease in them.

In cases where we identify “the” cause or causes of disease traits, these factors often have control that is specific, probable, and stable, for the disease of interest. It makes sense that we privilege these factors, because they provide types of control that serve our interests in explaining, treating, and preventing disease. For a given disease trait, targeting such factors is likely to provide (3.1) control over many to all instances of the disease, (3.2) a high likelihood of preventing each instance, and (3.3) control across a wide variety of genetic and environmental conditions that are present in the patient population.

4 Conclusion. What does this analysis suggest about the role of “pragmatics” in causal selection for disease traits? In the philosophical literature “pragmatic” is commonly used to imply that something is arbitrary, subjective, or audience-relative. Causal selection for disease does not appear to be “pragmatic” in any of these senses. The significant consensus on causal selection for some disease traits and my analysis of the principled rationale that guides this selection suggests that it is not simply an arbitrary procedure. The factors we select as disease causes provide special types of causal control that an arbitrary selection method would not explain. Furthermore, it isn’t clear that subjective or audience-relative preferences, whatever they may be,

¹⁵Boorse provides a helpful basic characterization of a reference class as “an age group of a sex of a species” (Boorse 1977, 555).

¹⁶Providing an account of reference classes and their role in stability assessments is outside the scope of this paper.

could capture the sense in which causal selection for disease is relative to a practical goal, viz. the goal of control. Information relevant to control can provide us with the means to change disease outcomes of real patients. It isn't just a contrived way of selecting causes that we want to talk about, that strike our fancy, or that we just happen to view as important, without good reason. The factors we select as "the" cause or causes of particular diseases often provide viable targets that we can use to *make a difference* in the disease outcomes of real patients.

This suggests that causal selection for disease is pragmatic in the sense that it is relative to the *practical* goal of control. Once the goal of control is specified, there are objective facts and considerations about what means conduce to it. These considerations include whether factors meet a minimal interventionist criterion (i.c.) in the second step of causal selection, and whether factors provide (1) broad scope, (2) probable, and/or (3) stable causal control in the third and final step of causal selection. These may not be the exclusive types of control that guide this process, but they help to clarify why we often single out few factors as "the" causes of particular diseases.

References

- Achinstein, P. (1984). The pragmatic character of explanation. pp. 1–19. PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association.
- Blaxter, M. (2015). *Health*. Polity Press.
- Boorse, C. (1977). Health as a Theoretical Concept. *Philosophy of Science*, 1–33.
- Bryant, P. A., J. Trinder, and N. Curtis (2004). Sick and tired: does sleep have a vital role in the immune system? *Nature Reviews Immunology* 4(6), 457–467.
- Caspi, A. and T. E. Moffitt (2006). Gene–environment interactions in psychiatry: joining forces with neuroscience. *Nature Reviews Neuroscience* 7, 1–8.
- Gannet, L. (1999). What’s in a cause?: The pragmatic dimensions of genetic explanations. *Biology & Philosophy*, 349–374.
- Griffiths, P. E. (2006). Proximate and ultimate information in biology. In *Festschrift for Philip Kitcher*. Oxford University Press.
- Hempel, C. (1965). *Aspects of scientific explanation*. And Other Essays in The Philosophy of Science. The Free Press.
- Kincaid, H. and J. Sullivan (2014). Classifying psychopathology: mental kinds and natural kinds. In H. Kincaid and J. Sullivan (Eds.), *Classifying Psychopathology: Mental Kinds and Natural Kinds*. The MIT Press.
- Lewis, D. (1973). Causation. *The Journal of Philosophy* 70(17).
- Lewis, D. (1986). Causal explanation. In *Philosophical Papers*. Oxford University Press.
- Lu, H., A. L. Yuille, M. Liljeholm, P. W. Cheng, and K. J. Holyoak (2008). Bayesian generic priors for causal learning. *Psychological Review* 115(4), 955–984.
- Mill, J. S. (1874). *A system of logic* (Eighth ed.). Harper & Brothers Publishers.
- Murphy, D. (2014). Natural kinds in folk psychology and in psychiatry. In H. Kincaid and J. Sullivan (Eds.), *Classifying Psychopathology: Mental Kinds and Natural Kinds*. The MIT Press.
- Poland, J. S. (2014). Deeply rooted sources of error and bias in psychiatric classification. In H. Kincaid and J. Sullivan (Eds.), *Classifying Psychopathology: Mental Kinds and Natural Kinds*. The MIT Press.
- Schaffer, J. (2014). The metaphysics of causation. *Stanford Encyclopedia of Philosophy*.
- Schaffner, K. F. (2012). A philosophical overview of the problems of validity for psychiatric disorders. In K. Kendler and J. Parnas (Eds.), *Philosophical Issues in Psychiatry II*, pp. 1–32. Oxford University Press.
- Scriven, M. (1959). Explanation and prediction in evolutionary theory. *Science*, 1–7.
- Shulman, J. M., P. L. De Jager, and M. B. Feany (2011, February). Parkinson’s disease: genetics and pathogenesis. *Annual Review of Pathology: Mechanisms of Disease* 6(1), 193–222.
- Waters, C. K. (2007). Causes that make a difference. *The Journal of Philosophy*, 1–30.
- Woodward, J. (2003). *Making things happen*. Oxford University Press.

- Woodward, J. (2006). Sensitive and insensitive causation. *The Philosophical Review* 115(1), 1–50.
- Woodward, J. (2010). Causation in biology: stability, specificity, and the choice of levels of explanation. *Biology & Philosophy* 25(3), 287–318.
- Woodward, J. (2015). The problem of variable choice. *Synthese* 194(4).
- Zachar, P. (2014). Beyond natural kinds: Toward a "relevant" "scientific" taxonomy in psychiatry. In H. Kincaid and J. Sullivan (Eds.), *Classifying Psychopathology: Mental Kinds and Natural Kinds*. The MIT Press.