

# Normative commitments, causal structure, and policy disagreement

Georgie Statham

## Abstract

Recently, there has been a large amount of support for the idea that causal claims can be sensitive to normative considerations. Previous work has focused on the concept of actual causation, defending the claim that whether or not some token event  $c$  is a cause of another token event  $e$  is influenced by both statistical and prescriptive norms. I focus on the policy debate surrounding alternative energies, and use the causal modelling framework to show that in this context, people's normative commitments don't just affect the causal claims they are willing to endorse, but also their understanding of the *causal structure*. In the context of the alternative energy debate, normative considerations affect our (implicit) understanding of the causal structure by influencing our judgements about which variables *should* be held fixed, and therefore which variables should be relegated to the background of a causal model. In cases of extreme disagreement, normative commitments can also affect which causal structure we think should be instantiated. Thus, focusing on a new context has revealed a previously unexplored sense in which normative factors are incorporated into causal reasoning.

Intuitively, we might think that causation is a natural relation; in particular, that whether or not  $c$  is a cause of  $e$  is a matter of fact, and therefore not influenced by any normative commitments. However, an increasingly large body of work questions this assumption, arguing that normative commitments do make a difference to at least some kinds of causal judgments. So far, the focus has been on the concept of *actual* (or token) causation—that is, the concept that is used to make claims of the form ‘event  $c$  caused event  $e$ ’. Both traditional philosophical analyses<sup>1</sup> and empirical studies<sup>2</sup> suggest that we are more likely to cite abnormal than normal events as causes, where the relevant notion of ‘normal’ includes both statistical and prescriptive norms.

I build on this work, showing that—and how—normative factors influence a different kind of causal judgment, which occurs when we are considering how to (and whether it is possible to) *bring something about*. In this context, normative commitments can affect people's (implicit) understanding of *causal structure*, as well as the causal claims

---

<sup>1</sup>For example Hall (2007); Halpern and Pearl (2005); Hitchcock (2007); McGrath (2005); Menzies (2004, 2007, 2009).

<sup>2</sup>For example Alicke et al. (2011); Hitchcock and Knobe (2009); Knobe (2010); Systma et al. (2012).

that they endorse. For example, consider policy debates. Here, different individuals often have different normative commitments; they frequently also make conflicting causal claims. I use the *causal modelling framework* to diagnose the disagreement between relatively well-informed people on opposing sides of one particular policy debate: that surrounding alternative energy. I show that in this context, people’s normative commitments do—and should—affect their implicit understanding of the causal structure.

The question that is central to the alternative energy debate is often expressed as follows: Is it possible to replace fossil fuels with alternative energy sources? This is a causal question, in the sense that it asks whether it is possible to bring about a certain state of affairs. After clarifying the question—and therefore what is at stake in the alternative energy debate—I use the causal modelling framework to show that people who give different answers to the above question can often be characterised as (implicitly) reasoning with different causal models. These differences in the understanding of causal structure arise for a variety of reasons, one of which is a difference in normative commitments.

In order to make causal judgements, we have to decide which variables are allowed to vary and which are held fixed. Our normative commitments influence where we draw this line—in general, we don’t (and shouldn’t) consider interventions on variables that we think *should* be held fixed. These variables are therefore relegated to the background. Thus, people with different normative commitments end up (implicitly) reasoning with different causal models, and generating different causal judgements.

Causal models generally have a limited *invariance range*: they represent a causal structure that is instantiated in some contexts, but not in others. In most cases, there are possible manipulations that exceed this invariance range, and in doing so, change the causal relationships that hold between the variables in the model. We can think of such interventions as altering the causal structure itself.<sup>3</sup> In these cases, normative considerations can affect not just which variables we think should be held fixed, but also which causal structure we think should be instantiated. Again, this affects our causal judgements.

It is generally accepted that the construction of causal models is to some extent subjective.<sup>4</sup> For example, which model is most appropriate will depend on the causal enquiry.<sup>5</sup> In the situations described here, however, the causal enquiry is held fixed, and yet normative commitments still influence which causal model is appropriate. This therefore repre-

---

<sup>3</sup>Of course, there is a sense in which the causal structure remains the same before and after such an intervention, and it is generally possible to represent both situations using a single, larger, causal model. The point is that when the intervention range is exceeded, the causal structure represented by the original causal model is no longer instantiated. There is therefore a sense in which the causal structure *has* changed.

<sup>4</sup>See Hitchcock (2007); Halpern and Hitchcock (2010); Woodward (2016).

<sup>5</sup>Questions about which causal model is appropriate should not be understood as pertaining to the metaphysics of causation, but to a different, methodological, question about how we should go about causal reasoning. For discussion, see Woodward (2015).

sents a previously unexplored sense in which our understanding of causal structure—and causal reasoning more generally—is influenced by mind-dependent factors.<sup>6</sup>

## 1 Previous work on causation and the normative

According to a traditional, and intuitive, understanding of causation, what causes what is completely determined by the physical facts. Thus, as Sarah McGrath notes, causation is ‘commonly held to be a paradigmatic example of a natural and so entirely non-normative relation’ (2005, 125).

Recently, however, a number of authors have argued that normative considerations do (legitimately) influence causal judgements. The claim (which is backed up by empirical research) is that we are more willing to cite *abnormal events* as causes, where, importantly, the relevant notion of ‘normal’ includes both descriptive and prescriptive norms.<sup>7</sup> Consider the following examples:

1. The heavy rain caused the flood.
2. My ear infection caused my dizziness.
3. Sophie’s alarm failing to go off caused her to be late for work.
4. The driver’s speeding caused the crash.

Examples 1–4 are all instantiations of the concept of *actual causation*—that is, they are all judgements of the form ‘*c* caused *e*’, where both *c* and *e* are token events.<sup>8</sup> In addition, the causes (and effects) are all deviations from some norm.

In 1, it rains more than is statistically normal for a particular region—that is, the heavy rain is a deviation from a *statistical norm*. Healthy humans don’t have ear infections. Thus, in 2, my ear infection is a deviation from what we might call a *norm of proper functioning* of human bodies. In 3, Sophie’s alarm failing to go off is a deviation from the norm of proper functioning of alarm clocks—that is, of a certain kind of artifact. Finally,

---

<sup>6</sup>It has previously been observed that normative factors make a difference to causal judgements in the sense that we judge *x* to be a cause of *y* if *y* counterfactually depends on *x* *relative to a default situation* (Hall, 2007; Hitchcock, 2007; Halpern and Hitchcock, 2010). However, the examples used to illustrate this phenomenon tend to start by representing the relevant situation using a particular causal structure, and then show that in order to account for our causal judgements, we have to add default values to this structure, where the default values are determined by both statistical and prescriptive norms. What I emphasise is that normative factors can (legitimately) influence the first step—that is, what we take to be the appropriate causal structure itself. As discussed in footnote 13, this can be seen as extending a suggestion from Blanchard and Schaffer (2017).

<sup>7</sup>This approach originates in Hart and Honoré (1959). For recent statements, see the works cited in Footnote 1.

<sup>8</sup>Hitchcock and Knobe point out that it is the concept of actual causation that has been the primary target of interest in the philosophy of causation (2009, 587). Judgements of actual causation can be understood as answers to the question: which event(s) was causally responsible for a particular effect? See Woodward (2011).

by speeding, the driver is certainly breaking the law—a *legal norm*—and potentially also a *moral norm*, if we think that he is recklessly endangering other people’s lives.<sup>9</sup>

The view that I have just outlined has been prominently defended by Christopher Hitchcock and Joshua Knobe (2009). Their aim is to determine the purpose of the concept of actual causation—that is, they ask what is the point of having a concept of causation that is sensitive to the expansive notion of ‘normal’ introduced above? In answer, they argue that the concept of actual causation is importantly linked to our ability to intervene in the world.

In general, while the causal structure identifies all of the factors that *could* be manipulated (either singly or in combination) to effect a change in the outcome, the actual causes are the factors that *should* be manipulated. (Hitchcock and Knobe, 2009, 590, italics in the original)

According to Hitchcock and Knobe, the reason we are more likely to cite events that are deviations from the normal as actual causes is that these tend to be ‘appropriate targets for intervention’ (2009, 591).<sup>10</sup> For example, they consider the case of a student who has failed a test, and wants to make sure this doesn’t happen again. They point out that the following counterfactuals are all true, and therefore all correspond to possible strategies (at least in theory).

- i. I would not have gotten an F if the teacher had been eaten by a lion.
- ii. I would not have gotten an F if the Earth’s gravitational pull had suddenly decreased.
- iii. I would not have gotten an F if I had had less to drink the night before the test. (2009, 591)

Only iii) is really relevant, and it is only this counterfactual that identifies a deviation from the normal (drinking too much is a deviation from a prudential norm). Additionally, this is the only option that identifies an appropriate target of intervention. Thus, while it may well be possible for the student to avoid failing her next test by somehow ensuring that her teacher gets eaten by a lion, we can see why it might be useful to have a concept that focuses on the fact that she got drunk the night before the test—that is, on factors that are deviations from the normal.

Most of the time people agree about which norms are relevant to a particular situation.

---

<sup>9</sup>An obvious objection to the kind of account described above is that it confuses causation with *blame*. Perhaps normative considerations only enter into our causal judgements to the extent that we are failing to distinguish causation and blame. Empirical studies have shown that judgements of actual causation are affected by negative evaluations such as blame (see e.g. Alicke et al. (2011)). However, this doesn’t exhaust the role played by the normative. Other studies have shown that we are more likely to cite deviations from the norm as actual causes, even when the outcome is positive (see e.g. Hitchcock and Knobe (2009)).

<sup>10</sup>See also Halpern and Hitchcock (2015).

But what if they don't?<sup>11</sup> One implication of Hitchcock and Knobe's account is that if people don't agree on the norms (or on which norms have priority), they will endorse different causal claims.

For example, in 2011, there were serious riots in London. At the time, I heard a lot of discussion of these riots, and in particular, much disagreement about what caused them. Some people asserted something like sentence 5, whereas others denied 5, and asserted something more like 6.

5. The culture of hooliganism caused the riots.

6. The state of high unemployment caused the riots.

According to the approach under consideration, disagreement about the causes of the riots can be explained by the fact that different individuals had a different understanding of which norms were relevant.<sup>12</sup>

In summary, existing work on the concept of actual causation describes one way in which our normative commitments plausibly affect our causal judgements, and also one way in which disagreement about the relevant norms can result in disagreement over causal claims. As discussed above, my goal is to extend this existing work, and show that normative commitments can affect our (implicit) understanding of causal structure.<sup>13</sup> For this, we need a way of representing causal structure. Thus, I now introduce the causal modelling framework.

## 2 Using causal models to represent causal structure

The causal modelling framework is a powerful system for representing causal structure. Visually, this structure is represented using causal graphs, which consist of a set of variables and arrows, or directed edges, each of which represents the existence of a direct causal relationship between two variables. For example, the graph in Figure 1 represents a causal structure that is instantiated in many river catchments.

Figure 1 shows that the amount of rainfall in the catchment area, the extent to which slopes are vegetated, and the steepness of the slopes, are all causally relevant to the river level. The river level then determines whether or not there is a flood.

---

<sup>11</sup>Halpern and Hitchcock consider examples in which there is disagreement about the relevant norms in their (2015).

<sup>12</sup>Perhaps it is more likely that there was general agreement about which norms held, but that different people prioritised these differently. This would also explain the disagreement over the cause of the riots.

<sup>13</sup>One way of understanding this project is as an elaboration of an idea from Blanchard and Schaffer (2017). In opposition to the work described above, they argue that we don't need to incorporate the deviant/default distinction into our account of actual causation. Rather, we need to pay more attention to what it takes to be an *apt causal model*. It is an implication of their argument (which they don't discuss) that there will be cases in which normative commitments influence which causal model we consider to be appropriate—that is, our understanding of the causal structure.




Figure 1: Causes of flooding

Of course, we generally want to know more than just that one variable is causally relevant to another variable; we also want to quantify this causal relevance. For example, we want to know how much rainfall in a particular river catchment is likely to result in a flood. This quantitative information is incorporated into causal models using *structural equations*; these express the value of each effect variable in the model as a function of its direct causes. For example, the causal model represented in Figure 1 would include structural equations expressing *RL* as a function of *R*, *V* & *S*, and *F* as a function of *RL*.

The causal modelling framework provides the basis for a particular theory of causation—the interventionist theory—which has recently been popularised in philosophy by James Woodward (2003). On this theory, variable *X* is causally relevant to variable *Y* if and only if there is a possible intervention on *X* that would make a difference to the value of *Y*, where ‘intervention’ is a technical term that is characterised using causal models.<sup>14</sup> For example, the amount of rainfall is causally relevant to the river level, because if we were to manipulate the amount that it rained in a particular region (assuming this were possible), this would make a difference to the river level.<sup>15</sup>

In the remainder of the paper, I use the causal modelling framework—and the associated interventionist theory of causation—to represent lay people’s implicit understanding of the causal structure underpinning the alternative energy debate. In doing so, I assume that human causal inference makes use of mental models that can (at least roughly) be represented using this framework.<sup>16</sup> This assumption is supported by recent work

<sup>14</sup>Roughly, an intervention variable on *X* with respect to *Y* has to be a cause of *X*, and has to affect *Y* (if at all) only via *X*. A random controlled trial is therefore a paradigmatic example of an intervention: the whole point of this experimental design is to ensure (as best as possible) that confounding factors are controlled for—that is, that any effect on the dependent variable is due to the independent variable. For Woodward’s precise characterisation of the notion of an intervention, see 2003, 98–99.

<sup>15</sup>Notice that causal models encode a set of counterfactuals. For example, Figure 1 asserts that there is a possible intervention on the amount of rainfall (*R*) that makes a difference to the river level (*RL*). This entails that there is a true counterfactual with the following form: if it were to rain *x* amount (rather than *x'* amount), the river level would be *y* (rather than *y'*).

<sup>16</sup>Thanks to an anonymous reviewer for pressing me to clarify the role played by the causal modelling framework in my argument.

in psychology. Many psychologists have adopted the causal modelling framework as a normative standard against which to assess human causal reasoning. It is also assumed that causal reasoning makes use of mental models that, although not necessarily identical to causal models, can be described using this framework.<sup>17</sup> Because my argument is founded on assumptions shared by recent work on causal cognition, it has implications for this field. At the very least, it makes a prediction that could be confirmed or disconfirmed by empirical studies.

Before leaving this section, there is one further feature of causal models that is particularly important in the context of my argument. Since any given causal model represents only part of the world, when constructing one of these models, there is always a decision to be made about where to draw the line between the variables that are part of the system being represented and those that are not—that is, between the foreground and the background. For example, in Figure 1, the variables  $R$ ,  $V$ ,  $S$ ,  $RL$ , and  $F$  comprise the foreground; everything else is in the background.

Unlike foreground variables, which can take a range of values within a single causal model, background variables are held fixed at some value, and are not considered as potential targets of intervention.<sup>18</sup> Generally, these variables are held fixed at the value that they normally take. For example, in Figure 1, the magnitude of the earth’s gravitational field is part of the background, and is held fixed at its actual value.<sup>19</sup> I assume that human causal inference makes use of mental models in which there is a foreground/background distinction with these features.

In the flood example, it is obvious that the earth’s gravitational field should be part of the background, rather than the foreground (and that it should be held fixed at its actual value). However, determining which background conditions to hold fixed is not always so uncontroversial. When we don’t agree on the relevant norms (and therefore which events are deviations from the normal), we are also likely to disagree about whether

---

<sup>17</sup>See e.g. Gopnik and Schulz (2007); Lagnado (2011). Lagnado explicitly claims that ‘successful causal inference presumably requires the capability to represent networks of directed relations between variables’ (2011, 139).

<sup>18</sup>A clarification is required here. It is not the case that *all* variables that are in the background (as opposed to the foreground) should be held fixed. In particular, background variables that are causally intermediate between a cause and its effects should be allowed to vary freely. For example, consider Figure 1. We could move the variable  $RL$  into the background, and represent  $R$ ,  $V$ , and  $S$  as direct causes of  $F$ . In this scenario, if we were to hold  $RL$  fixed, then intervening on any of  $R$ ,  $V$ , or  $S$  would not make a difference to the value of  $F$ , and so these variables would not come out as causally relevant. Thus, if the model under consideration is to correctly represent the causal relationships between these four variables, we had better allow  $RL$  to vary. More precisely, then, the requirement is that background variables that are not on a causal pathway between any of the foreground variables are held fixed at some value.

<sup>19</sup>Note that the background/foreground distinction differs from the normal/deviant distinction in that the former holds between variables, whereas the latter holds between values of variables. For example, imagine waking up tomorrow morning to discover that the earth’s gravitational field has become much stronger overnight. This would represent a deviant value of the variable {magnitude of the earth’s gravitational field}. We would probably respond by (either implicitly or explicitly) moving this variable from the background to the foreground of many causal models, including Figure 1.

a particular variable should be held fixed, and therefore whether it should be part of the background. As I will show, this is important, because drawing the line between the foreground and the background in different places results in different causal models, which can generate different causal judgements.

Using the debate surrounding alternative energies as a case study, I show that people with different normative commitments will often (implicitly) assume different causal models. Furthermore, this phenomenon explains some policy disagreement.

### 3 Applying causal models to the alternative energy debate

Let's consider the public debate about alternative energies, and, in particular, the question: Can we effectively replace fossil fuels with renewables?<sup>20</sup> In answering this question, we are not making a judgement of actual causation. Nevertheless, the question is causal in the sense that those who answer 'yes' think that we can bring about a certain state of affairs—that is, that it is possible to cause this state to come into being—whereas those who answer 'no' think that we can't.

My hypothesis is that people on different sides of the alternative energy debate are implicitly reasoning on the basis of different understandings of the causal structure that underpins the debate—that is, with different causal models—and that this partially explains why they give different answers to the same causal question. We might initially think that this causal disagreement can be easily resolved, because it is a matter of fact, not of value. However, I show that people's understanding of the causal structure is legitimately influenced by their normative commitments, and therefore that separating fact from value is not at all straightforward.

Before proceeding, I need to clarify the question at the heart of the alternative energy debate. We might worry that the question 'Can we effectively replace fossil fuels with renewables?', actually conflates two different questions, one technological (and causal) and one political (and normative). The first asks: Is it technologically possible to convert to alternative energies? The second: What energy policy *should* we adopt?

It is clearly true that the alternative energy debate does have both technological and political aspects. However, it is not clear that these issues should be separated. There is definitely a sense in which it is *possible* to immediately stop using fossil fuels and rely solely on renewables; it is just that the consequences of doing so would be devastating. It is also clear that the political question can't be answered without taking the state of technology into account. Thus, I think that the question most people are interested answering lies somewhere in between the purely technological and purely political. What

---

<sup>20</sup>For my purposes, it is not important to give a precise definition of the term 'renewables'. However, assume that it refers primarily to wind and solar.



we want to know—and what we disagree about—can be rephrased as follows: Is it possible to convert to renewables while maintaining our current standard of living (or with only a small reduction in standard of living)? This is an *economic* question. Importantly, it is also causal: as noted above, it asks whether it is possible to bring about a certain state of affairs.

Even after the original question is restated, we might worry that it still has normative content. We might expect that people will disagree on what it is to ‘maintain our current standard of living’, and that these differences will arise partly as a result of people’s different values. For example, consider two people who agree that converting to renewables is likely to result in both higher unemployment and a significant reduction in air pollution. These individuals might disagree about whether it is possible to convert to renewables while maintaining our current standard of living, because they disagree about what it is that constitutes a *good standard of living*. In particular, they disagree about the relative contribution that employment levels and air pollution make to standard of living.<sup>21</sup> In this kind of situation, differences in judgements about the possibility of converting to renewables may be attributed to the fact that people are asking whether it is possible to bring about different states of affairs.<sup>22</sup> This sort of disagreement is important. However, it is important to recognise that there are other sources of disagreement that also affect the alternative energy debate.

To see this, notice that by operationalising the notion of standard of living in terms of some economic variable (for example GDP), it is possible to restate the question at issue in a way that is purely empirical. It then becomes: Is it possible to convert to renewables while maintaining the current GDP (or with a reduction in GDP of less than  $x$ )?<sup>23</sup> Although most people engaged in the alternative energy debate do not state the question in these terms, I take it that most people would accept that maintaining GDP (at least approximately) is a basic requirement of any energy policy. I therefore assume that most advocates of alternative energies (to the extent that they have thought about economic issues at all) *do* think that it is possible to convert to renewables while maintaining GDP. On the other side of the debate, however, there are many advocates of fossil fuels who clearly think that it is *not* possible to switch to renewables without damaging the GDP.

---

<sup>21</sup>Another way of putting this point is to say that the two individuals discussed above disagree about the utility values that should be assigned to the values of the variables {employment level} and {air pollution}, and therefore how much these contribute to standard of living.

<sup>22</sup>There are also likely to be situations in which two individuals both agree that it is *possible* to convert to renewables while (roughly) maintaining our current standard of living, but disagree about whether this is *desirable*. See §5 for further discussion.

<sup>23</sup>Two clarifications. First, I leave aside the question of exactly what it would take to ‘convert to renewables’. Perhaps generating 90% of a country’s energy requirements using renewables would fulfil this criterion, or perhaps not. I suspect that different people have different targets in mind. Second, there is also a distributional issue. It may turn out that it is possible to convert to renewables without any reduction in a country’s GDP, but that the conversion nevertheless has unacceptable consequences for certain groups of people. I also set this issue aside for the purposes of this paper.

To reiterate, although our interest in whether or not it is possible to convert to renewables while maintaining the current GDP arises from normative concerns, the question itself is empirical. And people on opposing sides of the alternative energy debate disagree about the answer. My goal is to explain why this is the case.

## 4 An initial causal model

Just as there are many questions that we might ask with respect to the alternative energy debate, there are also many aspects of the debate that we could potentially represent using causal models.<sup>24</sup> For example, we could try to construct a quantitatively accurate model of a particular part of the energy economy, in order to accurately predict the effects of interventions. Or we could aim to analyse the causal reasoning of economists, scientists, and engineers working in this area. I am not engaged in either of these projects. Rather, my aim is to represent the causal structure that is assumed by relatively well-informed people (who are not necessarily experts) on different sides of the alternative energy debate, in order to diagnose their differing answers to the question: ‘Is it possible to convert to renewables?’ I demonstrate that people’s normative commitments make a difference to their implicit understanding of the causal structure underpinning the debate, and therefore to their assessment of the feasibility of converting to renewables.

My source of arguments for and against alternative energies was the website ProCon.org. This site compiles arguments for and against a range of issues that are controversial in the United States, including the alternative energy debate. The models constructed below are based on the arguments cited in response to the questions ‘Can alternative energy effectively replace fossil fuels?’<sup>25</sup> and ‘Should the US subsidise alternate energies?’<sup>26</sup> As discussed in the previous section, I assume that in the context of the US, the question ‘Can alternative energy effectively replace fossil fuels?’ can be paraphrased as ‘Is it possible for the US to convert to renewables while maintaining its current GDP (or with a reduction in GDP of less than  $x$ )?’<sup>27</sup>

It is generally assumed that renewables will only be able to (effectively) replace fossil fuels if they are competitive on price (although see §8 for a possible exception to this assumption). Thus, the answer to our question is determined by the relative price of

---

<sup>24</sup>Strictly speaking, I do not provide causal models, but causal graphs. To turn these into causal models, you would need to add a set of structural equations.

<sup>25</sup>ProCon.org (2017a) This site cites 19 arguments for, and 18 against, the above question.

<sup>26</sup>ProCon.org (2017b) This cites 7 arguments both for and against government subsidies of alternative energies. I included the responses to this second question because government subsidies came up in response to the question ‘Can alternative energy effectively replace fossil fuels?’, but were addressed more comprehensively in response to the more specific question.

<sup>27</sup>I assume that for it to be possible to convert to renewables is for there to be some practically realistic intervention (or set of interventions) that would result in this conversion. In terms of the causal models discussed below, this restricts us to considering interventions on the exogenous variables.

renewables and fossil fuels ( $P_R$  and  $P_F$ ).

A number of variables were mentioned as making a difference to the price of renewables. For example, as the technology of renewables ( $T_R$ ) improves, we can expect that the price of renewables will drop. Similarly, as the amount of energy that is produced by renewables ( $E_R$ ) increases (and therefore production of renewable technologies themselves increases) the price will also drop. Reliability is also a big issue for renewables ( $R_R$ ). We are assuming that this sector consists predominantly of wind and solar; the reliability issue arises because it is not always windy, nor is it always sunny. However, advances in the technology of renewables (and especially energy storage) can be expected to make at least some difference to the reliability of renewables, and therefore to their price.

Although reliability is not generally taken to be an issue for fossil fuels, changes in both the technology ( $T_F$ ) and the amount of energy produced from fossil fuels ( $E_F$ ) were mentioned as relevant to the price of fossil fuels.

Some people argued that environmental and health costs of either fossil fuels or renewables should be factored into their price. The variables  $EH_F$  and  $EH_R$  represent any mechanisms designed to do this (e.g.  $EH_F$  could represent a carbon tax).<sup>28</sup> Finally, government subsidies ( $GS$ ) were also cited as making a difference to the effective price of both kinds of energy source. Compiling all these variables into a causal graph results in Figure 2.

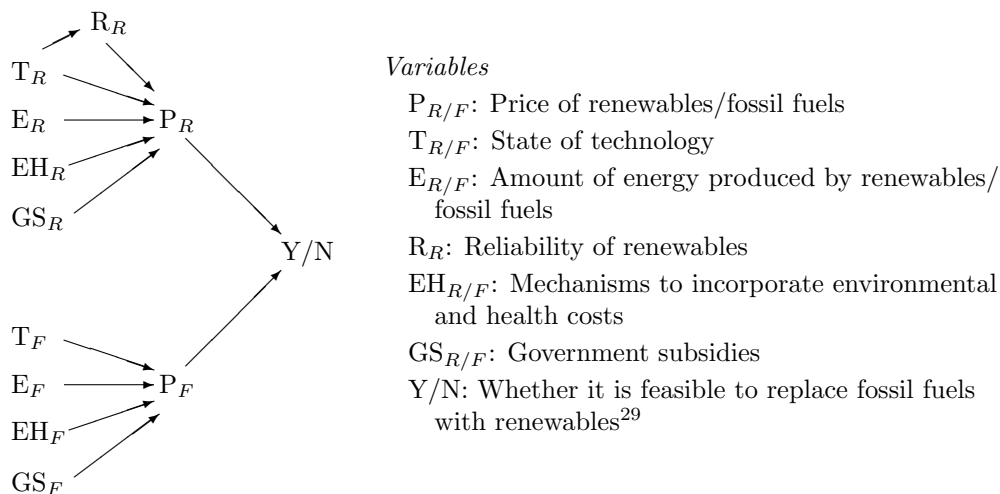


Figure 2: Causal structure of alternative energy debate

<sup>28</sup>To dispel any worries that  $EH_F$  and  $EH_R$ , as characterised above, are not apt for including in a causal model, notice that they are distinct from the other variables in the model, capable of being intervened on, and that we can assign them a set of possible values that represent incompatible states of affairs.

<sup>29</sup>It may seem strange to have a yes/no variable as an effect in this model. An alternative way of understanding this variable is as follows: Assume that the other variables in the model all represent the

It is important to reiterate that I do not take Figure 2 to be a comprehensive, complete, or even particularly accurate representation of the causal structure that underlies the alternative energy debate. However, with one caveat that is discussed in §8, I think that people on both sides of the debate would agree with the causal structure represented here. That is, they would agree that manipulating any of the variables represented in the graph would make a difference to the price of renewables or fossil fuels, and therefore to the feasibility of replacing fossil fuels with renewables.

The fact that it is possible to construct a causal graph that would be accepted by people on both sides of the debate might seem to indicate that, with respect to this policy debate at least, there is general agreement on the causal facts—that is, it might seem to negate the hypothesis that our normative commitments influence our causal judgements in this kind of context. In the next section, I argue that this is not the case, and that advocates of renewables and fossil fuels can be characterised as implicitly assuming different causal models. In §6, I argue that the reason for this is at least partly due to the fact that they have different normative commitments.

## 5 Models for and against

In order to extract the implicit causal structure assumed by advocates of renewables and fossil fuels, I again turned to the arguments for and against the questions ‘Can alternative energy effectively replace fossil fuels?’ and ‘Should the US subsidise alternate energies?’ cited on ProCon.org.

First, I considered just the arguments in favour of converting to renewables, and constructed a causal graph that included only those variables that were mentioned as being causally relevant in one of these arguments. I then repeated this procedure using the arguments against renewables. This generated the causal graphs in Figure 3a and b respectively. I now discuss the differences between these two graphs, and show that these occur for two reasons. First, disputants disagree about the *strength* of some of the causal relationships; second, they disagree about *which variables are in the foreground and which are in the background*.

Advocates of both fossil fuels and renewables accept that over time, as more energy is produced by renewables ( $E_R$ ) and the technology of renewables improves ( $T_R$ ), the price of alternative energy will go down. However, there is less agreement over the reliability of renewables ( $R_R$ ). Those who advocate for renewables tend to assume that advances in the technology of renewables will solve the reliability issue (thus pushing the price of renewables down significantly further). Those who support fossil fuels, on the other hand, tend to assume that there will be no significant improvement in the reliability of renewables—that is, they assume that the strength of the causal relationship between

---

situation at some time  $t = 0$ .  $Y/N$  can then be understood as representing the maximum proportion of energy that could be produced by renewables at some later time  $t = 1$ .

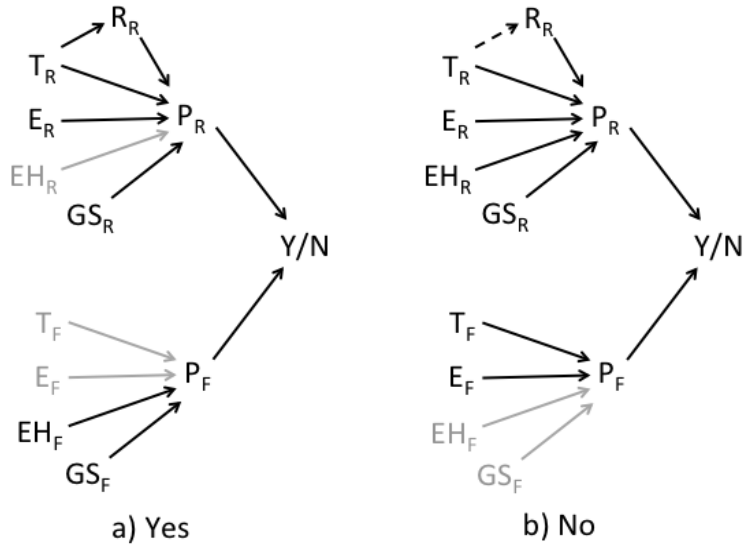


Figure 3: Causal structure assumed by those a) for and b) against renewables

$T_R$  and  $R_R$  is too weak to have a much effect. Thus, in this case, both sides agree on the structure of the causal graph, but disagree on some of the associated structural equations. The directed edge between  $T_R$  and  $R_R$  in Figure 3b is dashed to represent this difference.

Advocates of renewables tend to consider fossil fuel technology ( $T_F$ ) to be stable, and therefore background this variable. (Background variables are greyed out in Figure 3.) However, some supporters of fossil fuels argue that coal has become cleaner over the last twenty years, and that this trend will continue—that is, they include  $T_F$  in their causal model. Proponents of fossil fuels also point out that as more energy is produced by renewables, and therefore less by fossil fuels, the price of fossil fuels is likely to drop. Thus,  $E_F$  is explicitly included in their causal model. Advocates of renewables, on the other hand, tend not to mention this variable.

Perhaps unsurprisingly, both sides think we should be taking into account the environmental and health costs of the other. Thus, proponents of renewables argue that the environmental and health costs of fossil fuels should be factored into the price of this energy source, and vice versa. Currently neither of these cost is factored into the price of energy in the US. Thus, advocates of renewables are proposing that we intervene on  $EH_F$  but leave  $EH_R$  unchanged; vice versa for advocates of fossil fuels. Assuming that we tend to leave variables that we think should remain unchanged in the background, but place variables that should be intervened on in the foreground, this entails that each side is backgrounding the environmental and health costs of its favoured technology and foregrounding the costs of the other.

Finally, there is also disagreement about government subsidies. At the moment, both fossil fuels and alternative energies are subsidised by the US government. Those who advocate for renewables argue that it is crazy to be subsidising fossil fuels, and that subsidies on renewables should be increased. Thus, they consider both  $GS_F$  and  $GS_R$  as candidates for intervention. Supporters of fossil fuels, on the other hand, argue that we shouldn't be subsidising renewables, but don't mention any change in the subsidies on fossil fuels. Thus, they include  $GS_R$ , but background  $GS_F$ .<sup>30</sup>

To reiterate, my contention is that, for the reasons just discussed, the proponents of each side of the alternative energy debate are implicitly assuming different causal models, as illustrated in Figure 3.

Some of the differences between Figure 3a and 3b are about straightforwardly empirical matters. For example, we have seen that advocates of renewables and fossil fuels disagree about the strength of the causal relationship between  $T_R$  and  $R_R$ : the former tend to think that advances in the technology of renewables will make more of a difference, more quickly, to the reliability of renewables than the latter.

Since the extent to which advances in the technology of renewables will make a difference to the reliability of alternative energy is an empirical question, it seems likely that the reason that opponents of renewables tend to ignore this relationship is due to motivational bias—for example, the unconscious omission of information that doesn't help one's cause.<sup>31</sup> Some of the other differences between Figures 3a and 3b are also empirical and can also be explained by motivational bias. For example, whether using more renewables will affect the price of fossil fuels.

In the next section, however, I argue that some of the differences in people's implicit causal models can be attributed to their different *normative* commitments, in a way that can't simply be attributed to motivational bias. The upshot is that normative factors make a difference to our understanding of causal structure in a way that has not previously been discussed.

Before going on, it is necessary to respond to an objection.<sup>32</sup> Couldn't it be that people on both sides of the debate take Figure 2 to be the correct representation of causal structure, and disagree only about which interventions should be carried out, given this structure? For example, consider two people who completely agree on the likely effects of converting to renewables. In particular, like the two individuals discussed in §3, they agree that switching to renewables would lead to higher unemployment, as well as a decrease in air pollution. These two people may both accept that it is *possible* to convert to renewables (even while maintaining GDP), but disagree about whether we *should*, because they disagree about whether doing so would be an improvement over

---

<sup>30</sup>Some advocates of fossil fuels argue that the government should only subsidise renewables if these are close to being competitive with fossil fuels on price. That is, they argue that there should be a causal arrow going from  $P_R$  to  $GS_R$  (or perhaps from  $Y/N$  to  $GS_R$ ). I leave out this postulated causal link.

<sup>31</sup>For a review of work on motivational bias, see Kunda (1990).

<sup>32</sup>Thanks to two anonymous reviewers for pressing this objection.

the status quo.

I think it is likely that this kind of disagreement does exist. However, there are two reasons to think that this doesn't explain the disagreement cited on ProCon.org. First, the passages on this website are intended to be answers to the question 'Can we effectively replace renewables with fossil fuels?' That is, the people cited here explicitly disagree about whether it is *possible* to bring about a particular state of affairs, not whether we *should* bring about this state of affairs.

Second, mere disagreement about which interventions should be carried out doesn't explain the arguments cited on ProCon.org. For example, imagine an advocate of fossil fuels who thinks that we shouldn't incorporate the environmental and health costs of fossil fuels into the energy economy. Now say that she assumes that the causal structure is as given in Figure 2. She is therefore explicitly aware of the fact that if we were to incorporate the environmental and health costs of fossil fuels into the energy economy, this would increase the feasibility of converting to renewables. When defending her position, you would expect such a person to give some argument for *why* we shouldn't incorporate the environmental and health costs of fossil fuels. After all, it is on this point that she disagrees with her opponents. However, there is no evidence of this kind of argument. In fact, none of the advocates of fossil fuels mention the environmental and health costs of this energy source at all. Thus, while the kind of disagreement considered in this objection may well occur, I take the evidence from ProCon.org to be more consistent with the claim that people on different sides of the debate are implicitly working with a different understanding of the causal structure.<sup>33</sup>

## 6 Normative commitments and causal models

Let's consider the effect of environmental and health costs on the price of renewables and fossil fuels. Recall that the variables  $EH_R$  and  $EH_F$  represent any mechanisms that are designed to incorporate the environmental and health costs of renewables and fossil fuels into their price. There are a number of questions that are relevant to these variables, including: What are the environmental and health costs of renewables and fossil fuels?<sup>34</sup> Should these costs be incorporated into the energy economy? And, if yes, how? Answers to the first and third question influence the *values* we assign to  $EH_R$  and  $EH_F$ , whereas answers to the second affect whether we think these variables should be intervened on at all.

---

<sup>33</sup>It is possible that some of the apparent differences in people's understanding of the causal structure are actually due to rhetoric—that is, that people on opposing sides of the debate agree on the causal structure, but only emphasise those aspects of the structure that support their argument. However, it seems implausible that rhetoric can fully account for the differences represented in Figure 3.

<sup>34</sup>This first question may appear to be empirical; certainly it is partly empirical, and it is possible to operationalise environmental and health costs such that they can be assessed empirically. However, it seems unlikely that there is a value neutral choice of operationalisation. If not, there is a genuinely normative question here.

Because we're trying to understand the way that normative commitments affect our understanding of causal structure (as opposed to the values that we assign to variables), let's focus on the second question. This is obviously value laden. Whether we think that environmental and health costs should be incorporated into the energy economy will depend on what we are trying to achieve, as well as what we think will work. For example, it will depend on the value we place on both the environment itself and the welfare of future generations.

Someone who thinks that the environmental and health costs of fossil fuels shouldn't be incorporated into the energy economy thinks that we shouldn't intervene on this variable—that is, that it should be left at its current value. In general, if a variable shouldn't be intervened on, there is no need to include it in the foreground of a causal model: it only makes sense to include variables that we think should be varied (or that it is acceptable to vary).<sup>35</sup> Thus, to think that the environmental costs of fossil fuels shouldn't be incorporated into the energy economy is to think of  $EH_F$  as part of the background.<sup>36</sup>

To generalise, because of their different normative commitments, people on different sides of a policy debate will often have different ideas about which variables we *should* manipulate. Since variables that shouldn't be intervened on are automatically relegated to the background of a causal model, this entails that where we draw the line between the foreground and the background can be influenced by our normative commitments. Normative disagreement can therefore result in individuals implicitly assuming different causal models.

Importantly, different causal models can result in different causal judgements—in this case, in different judgements about whether it is feasible to convert to renewables. For example, say Tom thinks that the environmental and health costs of fossil fuels should be taken into account, but Ann thinks they shouldn't be. Tom therefore puts  $EH_F$  in the foreground, where it can take a range of values, including, for example, 'carbon tax'. Ann, on the other hand, leaves this variable in the background, where it takes its current value: 'none'. In this scenario, Tom will judge that the price of fossil fuels could—and should—be higher than Ann does. He is therefore more likely to think that it is feasible to convert to renewables while maintaining the GDP—that is, that there are interventions that we can carry out to bring about this state of affairs.

Note that Tom and Ann do not disagree on the strength of any causal relationship, or on any other empirical question. They may well even be in agreement on both what the environmental and health costs of both fossil fuels and renewables are, and how we

---

<sup>35</sup>Another way of putting this point is to note that if you think that a variable shouldn't be intervened on, you think that there is only one value of that variable that represents a *serious possibility*. You therefore don't need to represent other values of this variable in your causal model. For further discussion, see Blanchard and Schaffer (2017, 197–198).

<sup>36</sup>Whether or not we think that the government should subsidise renewables and/or fossil fuels (and by how much) is another place in which there is a genuine difference in normative commitments that affects which variables are in the foreground and which are in the background.



should price these into the economy if we are going to do so. Nor do they disagree about the existence of any of the individual causal links in Figure 2—that is, they agree on the causal structure *once the boundary between the foreground and the background has been specified*. However, because they consider different sets of possible interventions to be *available*, they automatically draw the line between the causal system and the background in different places, and thus reason on the basis of different causal models, in which the range of possible values of  $EH_F$  is different. As a result, they generate different answers to the question: Is it possible for the US to convert to renewables while maintaining its current GDP?

The above discussion implies that talk of our ‘understanding of the causal structure’ is ambiguous. This phrase can refer to i) a system of counterfactuals that an individual would accept on reflection (e.g. as represented in Figure 2) or ii) the model she is implicitly reasoning with (e.g. Figures 3a and b). I come back to this point in §10.

Since this is not an empirical study, I haven’t identified individuals fulfilling the descriptions of Tom and Ann. Nevertheless, since normative disagreements like theirs certainly do exist, it seems reasonable to conclude that normative commitments do affect people’s implicit understanding of causal structure—and therefore their causal judgements—in the way I have described.

At this point you might be thinking ‘Sure different people implicitly reason with different causal models, but that’s just because they’re not reasoning with the *correct* causal model. We could resolve this disagreement about causal structure by being more careful about the causal models that we use to reason with’. In the next section, I show why this suggestion doesn’t succeed.

## 7 Resolving disagreement

In order to construct a causal model, we first have to decide where to draw the line between the system of interest and the background. Since background variables are held fixed, but foreground variables can take a range of values, this decision can make a difference to causal judgements. But, we might think, isn’t there a right way of deciding where to draw this line? If so, we should be able to resolve any disagreement over the causal structure. I now consider two suggestions as to how this might work, both of which are ultimately unsuccessful.

First, we could stipulate that any variable that is causally relevant and that does in fact vary in the context of interest (or is likely to vary in the context of interest) must be explicitly included in the causal model that is used to reason with. This is, indeed, what we should aim to do, although as Jennan Ismael points out, there are reasons to think that this is a practical question, that don’t necessarily have one, correct answer.

In decision contexts, what we hold fixed is a partly causal question that

depends on what we can expect to be fixed in the hypothetical circumstances in which the choice will take place ... There is no simple recipe for making these judgments. They are causal judgements, but ones that demand ... practical wisdom beyond mere scientific knowledge. Ismael (2013, 229)

Ismael is referring to decisions that individuals make about their lives. For example, when deciding what to study at university, we have to judge what kind of life we will want in the future, as well as what interests us at the time.

We should expect that Ismael's comments also apply to policy debates—that is, that here, too what we should hold fixed depends on what we can expect to be fixed in future hypothetical situations (and that there is no simple recipe for making these judgements, either).

My point is that, although pursuing the above strategy may well resolve some disagreement (for example that surrounding the reliability of renewables), my discussion of the alternative energy debate—and particularly the environmental and health costs associated with this debate—shows that considering which variables are likely to be held fixed (and to vary) is not sufficient. In the case of the environmental and health costs, the question is not only whether these are *likely to vary*, but also whether they *should* be varied—that is, the issue is normative, as well as empirical.

In principle, it is a good idea to relegate variables that should be held fixed to the background when constructing a causal model. For example, we saw in §3 that when people who are engaged in the alternative energy debate ask whether it is possible to replace fossil fuels with renewables, they are (usually) really asking whether it is possible to do so *while maintaining our current standard of living*. That is, people on both sides of the debate agree that we should hold our standard of living fixed when thinking about the causal structure underpinning the debate. In general, there are always going to be many variables that are causally relevant in a given situation, so it makes sense (and is probably essential) to narrow the focus to those that it is reasonable to intervene on.

Nevertheless, we might think that one way of resolving disagreement about the causal structure (if not about the resulting causal judgements) is to include every variable that is taken to be a suitable target of intervention by *some* people in the relevant debate. This is therefore the second suggestion for determining the appropriate causal model.

Considering Figures 2 and 3 shows that this second suggestion does have some promise. For example, let's return to Ann, an advocate of fossil fuels, who begins by (implicitly) reasoning on the basis of the model in Figure 3b. She thinks that we shouldn't incorporate the environmental and health costs of fossil fuels into the energy economy, and so has implicitly backgrounded  $EH_F$ .

Now say that Ann is presented with Figure 2. She is likely to agree that this model correctly represents the causal structure underpinning the alternative energy debate. For example, she agrees that if we *were* to introduce a mechanism to take the environmental

and health costs of fossil fuels into account, this *would* make a difference to the price of fossil fuels. Importantly, however, this doesn't imply that she now thinks that Figure 3b is incorrect (she still thinks that  $EH_F$  *should* be held fixed). Rather, it shows that the same causal system can be represented by more than one causal model.

Figure 2 is useful because presenting individuals on opposing sides of the debate with this causal structure can reveal the source of (some of) their disagreement. Once everyone has agreed that Figure 2 is an acceptable causal model, it becomes clear that the source of their disagreement is normative: they disagree about which interventions we *should* carry out. Thus, it is possible to diagnose some disagreement by explicitly introducing a causal model that includes all the variables that are taken to be a suitable target of intervention by either side. In the next section, however, I show that this strategy is less successful when the disagreement is more extreme.<sup>37</sup>

## 8 Changes in the economy

There is one causally relevant factor that is mentioned in the arguments for renewables cited on ProCon.org that I haven't been able to include in the causal graphs so far. This is raised in the following passage:

The transition will entail costs—not just money and regulation, but also changes in our behaviour and expectations ... a truly all-renewable economy may be very different from the American economy we know today. The renewable economy will likely be slower and more local; it will probably be a conserver economy rather than a consumer economy. Richard Heinberg, Post Carbon Institute (ProCon.org, 2017a)

The causal graphs in Figures 2 and 3 presuppose that the energy economy stays constant in the sense that it is assumed that in order for it to be feasible for renewables to replace fossil fuels, renewables have to be competitive on price.<sup>38</sup> However, as soon as people start talking about a radically different economy, we have to abandon this assumption.

Should we hold the economy constant? This is clearly a normatively laden question: Heinberg—and others who suggest transitioning to an economy of renewables—are partly motivated by their ideas about what our society *should* be like. The suggestion that we move to a different energy economy therefore provides a second example in which normative commitments influence our implicit understanding of the causal structure. In one

---

<sup>37</sup>Note that normative factors can also make a difference to our assessment of the *value* of certain variables (e.g. environmental and health costs). Incorporating all the variables that are considered to be suitable targets of intervention by some people engaged in the relevant debate will also fail to resolve this kind of disagreement.

<sup>38</sup>The economy is not assumed to be constant in the sense that any specified variable (or set of variables) is held fixed. Rather, it is assumed that we don't exceed the invariance range of the graph in Figure 2.

sense, this is a more extreme version of the phenomenon exemplified by environmental and health costs (in §6). Again, normative disagreement affects what we think should be held fixed.

However, this second example is more extreme in that the causal structure of the energy economy envisaged by Heinberg would presumably be very different to that in existence today, and therefore to that represented by Figure 2. For example, a conserver economy would be more localised, meaning that the distribution of resources would become more important, and the overall market price of certain forms of energy less important.<sup>39</sup> Thus, whether it is possible to convert to renewables may no longer be primarily determined by the relative price of renewables and fossil fuels. In other words, the variable  $Y/N$  would have a different set of direct causes. Thus, Heinberg is not just questioning whether a particular variable should be held fixed, but suggesting that we should intervene to change the causal structure of the energy economy itself.<sup>40</sup> That is, he doesn't think that Figure 2 is a correct representation of the causal structure that *should* underpin the alternative energy debate.

It is an open question whether the alternative economy envisaged by Heinberg is a realistic possibility. Thus, there is likely to be disagreement about whether it is even *possible* to alter the causal structure in the way that he suggests.

The larger point, however, is that there are situations in which we are able to control which causal structure is instantiated, and that in such situations, our normative commitments can influence which causal structure we think *should* be instantiated.

## 9 Summary

As we have seen, a number of authors have recently claimed that normative factors make a difference to judgements of actual causation, in that actual causes tend to be deviations from either descriptive or prescriptive norms. Hitchcock and Knobe argue that the purpose of this normatively loaded concept of causation is to discover 'appropriate targets for intervention' (2009, 591). Recall that they claim that 'the actual causes are the factors that *should* be manipulated' to change the outcome (2009, 590).

I've focused on a different kind of causal judgment—judgements about which states we can bring about—and shown that these are also sensitive to which factors we think *should* be manipulated. Unlike judgements of actual causation, however, here our normative

---

<sup>39</sup>The causal structure of a conserver economy may be quite different in different locations—even between different cities in the same country.

<sup>40</sup>The kind of intervention described above is often referred to as a *structural intervention*. For example, see Malinsky (2017). The structural interventions Malinsky considers involve changes to the parameters of structural equations. However, there are also situations in which we can affect the causal structure more dramatically. For example, some manipulations change *which* variables are causally relevant to a given effect (as described above). In other work, I have argued that this is actually a very common practice, which is essential to our ability to control the world (2017).

commitments affect our causal judgements by influencing our implicit understanding of the causal structure itself.

I have argued that, because we don't (and generally shouldn't) entertain interventions on variables that we think should be held fixed, these variables are included in the background, rather than the foreground, of a causal model. Thus, people with different normative commitments end up (implicitly) reasoning with different causal models, and therefore effectively talking past each other.

Implicit disagreement about causal structure can often be resolved by making the causal model explicit, and including every variable that is taken to be a suitable target for intervention by some people in the relevant debate (e.g. Figure 2). In extreme cases, however, normative considerations can affect which causal structure we think should be instantiated.

I end by commenting on the value of the causal modelling framework to policy debates.

## 10 The value of causal models

In general, there are many systems of counterfactuals that are true of a given part of the world. For example, the three causal graphs in Figures 2 and 3 all represent different causal structures—and thus different systems of counterfactuals—that are relevant to the alternative energy debate.<sup>41</sup> Furthermore, these graphs all represent systems of counterfactuals that are at least plausibly true—that is, corresponding to every directed edge (or chain of directed edges) in these graphs, it is at least plausible that there is some intervention on  $X$  that would make a difference to  $Y$ . Even those who think that the economy should be changed would presumably agree on the causal structure in Figure 2 *given that we hold the economy fixed*.

However, in order to make a causal judgement—for example, ‘ $X$  causes  $Y$ ’, or ‘It is possible to bring about  $Z$ ’—you have to decide what is held fixed (at what value) and what is allowed to vary—that is, you have to pick one set of variables and background conditions, and reason on the basis of these. This is because, as we have seen, holding different variables fixed (at different values) results in different counterfactuals being true.<sup>42</sup>

Different decisions about where to draw the line between the system and the background correspond to different causal models. When people implicitly reason with different causal models in this sense, what they are really (implicitly) disagreeing about is which model is most *appropriate*. We can understand person  $A$  as saying that it is possible to bring about  $Y$ , *assuming a particular division between foreground and background*, and

---

<sup>41</sup>See footnote 15 on the connection between causal structure and counterfactuals.

<sup>42</sup>For more discussion, see Halpern and Hitchcock (2010).

person *B* as saying that it is not possible to bring about *Y*, *assuming a different division*. As a result, they end up talking past each other.

Causal models are useful because they can make explicit the causal structure that is being assumed. As I have shown in the case of the alternative energy debate, this allows us to identify points of disagreement, and to consider whether these are empirical, or due to genuine normative disagreement. Empirical disagreements can then be resolved, and normative disagreements can at least be recognised for what they are. Thus, my contention is that the use of causal models could improve the public debate surrounding the use of alternative energies in particular, as well as policy debates more generally.

### Acknowledgements

Thank you to Claire Benn, Sharon Berry, Casper Storm-Hansen, Moshe Justman, York Hagemayer, Christopher Hitchcock, Enno Fischer, participants at the Linguistic Perspectives on Causation Workshop at the Language Logic and Cognition Center, Hebrew University of Jerusalem (June 2017) and the Polonsky Academy Seminar (May 2017), and especially to three anonymous reviewers from this journal.

### References

- Mark D. Alicke, David Rose, and Dori Bloom. Causation, norm violation, and culpable control. *Journal of Philosophy*, 108:670–696, 2011.
- Thomas Blanchard and Jonathan Schaffer. Cause without default. In Helen Beebe, Christopher Hitchcock, and Huw Price, editors, *Making a Difference: Essays on the Philosophy of Causation*, pages 175–214. Oxford University Press, Oxford, 2017.
- Alison Gopnik and Laura Schulz, editors. *Causal Learning: Psychology, Philosophy, and Computation*. Oxford University Press, Oxford, 2007.
- Ned Hall. Structural equations and causation. *Philosophical Studies*, 132:109–136, 2007.
- Joseph Halpern and Christopher Hitchcock. Actual causation and the art of modelling. In Rina Dechter, Hector Geffner, and Joseph Y. Halpern, editors, *Probability and Causality: A Tribute to Judea Pearl*, pages 383–406. College Publications, 2010.
- Joseph Y. Halpern and Christopher Hitchcock. Graded causation and defaults. *British Journal for the Philosophy of Science*, 66:413–457, 2015.
- Joseph Y. Halpern and Judea Pearl. Causes and explanations: A structural-model approach. Part 1: Causes. *British Journal for the Philosophy of Science*, 56:843–887, 2005.
- H. L. A. Hart and Tony Honoré. *Causation in the Law*. Clarendon Press, Oxford, 1959.
- Christopher Hitchcock. Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, 116:495–532, 2007.

- Christopher Hitchcock and Joshua Knobe. Cause and norm. *Journal of Philosophy*, 106: 587–612, 2009.
- Jennan Ismael. Causation, free will, and naturalism. In Don Ross, James Ladyman, and Harold Kincaid, editors, *Scientific Metaphysics*, pages 208–235. Oxford University Press, Oxford, 2013.
- Joshua Knobe. Person as scientist, person as moralist. *Behavioral and Brain Sciences*, 33:315–329, 2010.
- Ziva Kunda. The case for motivated reasoning. *Psychological Bulletin*, 108:480–498, 1990.
- David Lagnado. Causal thinking. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*, pages 129–149. Oxford University Press, Oxford, 2011.
- Daniel Malinsky. Intervening on structure. *Synthese*, published online, 2017.
- Sarah McGrath. Causation by omission: A dilemma. *Philosophical Studies*, 123:125–148, 2005.
- Peter Menzies. Difference-making in context. In John Collins, Ned Hall, and L. A. Paul, editors, *Causation and Counterfactuals*, pages 139–180. MIT Press, Cambridge, MA, 2004.
- Peter Menzies. Causation in context. In Huw Price and Richard Corry, editors, *Causation, Physics, and the Constitution of Reality*, pages 191–223. Clarendon Press, Oxford, 2007.
- Peter Menzies. Platitudes and counterexamples. In Helen Beebe, Christopher Hitchcock, and Peter Menzies, editors, *The Oxford Handbook of Causation*, pages 341–367. Oxford University Press, Oxford, 2009.
- ProCon.org. Can alternative energy effectively replace fossil fuels?, 2017a. URL <http://alternativeenergy.procon.org/view.answers.php?questionID=001244>.
- ProCon.org. Should the us subsidize alternative energies?, 2017b. URL <http://alternativeenergy.procon.org/view.answers.php?questionID=001251>.
- Georgie Statham. The manipulation of chemical reactions: probing the limits of interventionism. *Synthese*, 194:4815–4838, 2017.
- Justin Systma, Jonathan Livengood, and David Rose. Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attributions. *Studies in History and Philosophy of Science Part C*, 43:814–820, 2012.
- James Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, Oxford, 2003.

James Woodward. Psychological studies of causal and counterfactual reasoning. In Christoph Hoerl, Teresa McCormack, and Sarah R. Beck, editors, *Understanding Counterfactuals, Understanding Causation*, pages 16–53. Oxford University Press, Oxford, 2011.

James Woodward. Methodology, ontology, and interventionism. *Synthese*, 192:3577–3599, 2015.

James Woodward. The problem of variable choice. *Synthese*, 193:1047–1072, 2016.