

Broken Brakes and Dreaming Drivers: The Heuristic Value of Causal Models in the Law

Enno Fischer*

Accepted for publication in
European Journal for Philosophy of Science

Abstract

Recently, there has been an increased interest in employing model-based definitions of actual causation in legal inquiry. The formal precision of such approaches promises to be an improvement over more traditional approaches. Yet model-based approaches are viable only if suitable models of legal cases can be provided, and providing such models is sometimes difficult. I argue that causal-model-based definitions benefit legal inquiry in an indirect way. They make explicit the causal assumptions that need to be made plausible to defend a particular claim of actual causation. My argument concerns the analysis of legal cases involving a combination of double prevention and causal redundancy. I show that discussions among legal theorists about such cases sometimes suffer from ambiguous assumptions about the causal structure. My account illustrates that causal models can act as a heuristic tool for clarifying such assumptions, and that causal models provide a framework for more accurate analyses of legal cases involving complex causal structure.

Keywords: Actual Causation, Causal Models, Legal Cases, Causal Redundancy, Double Prevention

*Institut für Philosophie I, Ruhr-Universität Bochum, Universitätsstr. 150, 44801 Bochum. enno.fischer@ruhr-uni-bochum.de

1 Introduction

A rental company failed to discover and repair defective brakes in a car that was rented to a driver. The driver failed to try to use the brakes to avoid running into a pedestrian. The injury to the pedestrian would have been avoided if and only if the rental company had repaired the brakes and the driver had tried to use them. What is the actual cause of the pedestrian's injury? This is a difficult question because neither the company's failure nor the driver's failure are difference-makers in the current situation. Even if the driver had applied the brakes, the car could never have stopped because the brakes were faulty. Even if the company had checked and repaired the brakes, the car could not have stopped because the driver did not even try to apply the brakes.¹

Recently, there has been an increased interest in applying causal models in the context of legal inquiry, specifically, to study cases with complex causal structure and to answer questions of actual causation.² An important worry that comes up regarding model-based definitions of actual causation is whether the model dependence might undermine the definitions' potential to adjudicate between opposing views. Asked whether an action was in fact a cause of the harm would yield different answers depending on which of the (likely opposing) models is selected. Unless there are objective criteria to select a model, the framework seems to provide little against such a situation of impasse.

In response to worries of model dependence, causal modelers, such as Joseph Halpern, have argued that model-dependence is not a drawback but a potentially useful feature of such definitions because it "moves the question of actual causality to the right arena—debating which of two (or more) models of the world is a better representation

¹The discussion is inspired by the legal case *Saunders System Birmingham Co v Adams* 117 So 72 (Alabama, 1928). Among legal theorists the braking case and similar cases involving the non-use or misuse of missing or malfunctioning safety devices has attracted considerable discussion and there is no agreement on the causal verdict regarding such cases (Wright, 1985, 2001, 2011, 2022; Fischer, 1992, 2006; Stapleton, 2008; Abrams, 2022). There has also been philosophical interest in the braking case as an instance involving preemptive omission (Metz, 2022).

²There have been various formalizations of legal criteria for causation such as Richard Wright's (1985) NESS condition (Baldwin and Neufeld, 2004; Halpern, 2008; Bochman, 2018; Beckers, 2021), and suggestions to formalize concepts of responsibility and blame (Chockler and Halpern, 2004; Halpern, 2016). Several case studies have been provided in which this approach has been applied to clarify legal cases (Halpern, 2015; Chockler et al., 2015; Liepina et al., 2018) and other cases involving disagreement over causes (Statham, 2020; Fischer, 2021b; Hanley, 2022). Moreover, there is a growing body of empirical research suggesting that the kind of explanatory and evidential reasoning featuring in juror decision making can be framed in terms of causal models (Lagnado and Gerstenberg, 2017).

of those aspects of the world that one wishes to capture and reason about. This, indeed, is the type of debate that goes on in informal (and legal) arguments all the time" (2015, 91f). However, this response does not explain the real gain of employing the formal machinery of causal models. If questions of actual causation are simply transformed to questions of proper representation that are similar to those coming up in informal arguments, why should we care to employ the formal framework?

Here I provide an analysis of the ways in which legal inquiry can benefit from causal models as a heuristic device. Even if they sometimes do not provide conclusive criteria for actual causation, model-based definitions of actual causation still inform us about the structural features of a case that need to be made plausible if we wish to defend a particular claim of actual causation. Moving the discussion to the level of causal models can represent genuine progress because it can help us identify those aspects of a legal case that matter for the evaluation of specific causal hypotheses. In this sense causal models provide guidance towards the evidence that matters for solving a legal case.

More specifically, I will explore three regards in which the heuristic value of causal models can play out in the analysis of cases like the braking case. First, there is a question whether the driver and the rental company are causally on a par or whether one of the factors has causal priority. Causal modelers have developed strategies for distinguishing such cases. So far, these strategies have mostly been applied to address conceptual questions concerned with appropriate definitions of actual causation. Here I suggest that the techniques can also be applied to shed new light on the structure of legal cases with causal redundancy. More specifically, it is a consequence of causal-model-based definitions that instances of preemption have a certain signature in causal models that differs from the signature of symmetric overdetermination. Whether a causal model exhibits the preemption signature or not can help see more clearly whether one factor has causal priority over another factor or not.

Second, causal models may be used to explore ways in which the assumed symmetry between two factors in a model can be broken by including further evidence. Given the current description of the case, it seems that the driver and the rental company are causally on a par. What kinds of circumstances would overturn this judgment? This

is an important consideration because inquiry is typically driven by a specific causal hypothesis. The framework helps to highlight under what circumstances the hypothesis may have to be overturned.

Third, causal models may be employed for comparing cases with supposedly similar structure. The braking case is often seen in analogy with other failure-to-warn cases. But sometimes causal judgments appear to diverge between the cases. On my account this should act as a warning signal that the causal structure may be relevantly different after all. On the level of causal models this is reflected by a disturbance of isomorphisms between the respective models.

There is now a growing body of empirical research on the viability of causal models as potentially representing forms of causal reasoning in legal and moral decisions (e.g. Lagnado et al. (2013); Lagnado and Gerstenberg (2017)). More recently, empirical research has suggested that employing causal models as a tool for representing conflicting explanations in jury decisions can have significant impact on the outcomes of such decision-making (Liefgreen and Lagnado, 2023).

Here I am not aiming to address the empirical question of whether legal decision makers *do* reason in accordance with the causal modelling framework. Nor am I aiming to provide additional empirical evidence of the impact that employing the modelling framework may have. Instead, I am interested in exploring reasons why legal decision-making *should* employ such models to support the causal analysis of cases with complex structure, the main point being that causal models can act as a useful heuristic tool for systematically analyzing relevant causal relations. More precisely, I will argue so for cases involving the kind of causal redundancy exhibited by the braking case and other failure-to-warn cases. While the argued points may extend to a broader class of legal cases, an explicit analysis of such cases will have to be subject to future research.

Here is an outline of the argument. In section 2 I suggest a simple model of the braking case, according to which the rental company's failure and the driver's failure are causally on a par. The model appears to capture the main features of the situation's causal structure. Yet one might worry that the causal model is symmetrical regarding the rental company and the driver, and that the result of the causal analysis is thus

built into the model from the get-go. In sections 3 and 4 I address these concerns by examining more closely a well-known proposal to treat the braking case as an instance of preemptive prevention (Wright, 1985, 2001, 2011). According to this proposal, only the driver's actions are an actual cause because they preempted the rental company's failure. I show that this approach to the braking case faces important challenges. In sections 5 to 7 I address the heuristic value of causal models. In section 5 I discuss the idea that preemption cases (including preemptive prevention) have a certain signature in causal models and show that it is difficult to realize the signature in causal models of the braking case. In section 6 I look beyond the given description of the braking case, and show that models can be used to generate hypotheses regarding further evidence that might overturn the causal evaluation of the braking case. In section 7 I extend the analysis to a number of other cases that involve the non-use or misuse of missing or malfunctioning safety devices and address the heuristic value of comparing the structure of these cases. In section 8 I sum up the results of the discussion.

2 A Simple Model for the Braking Case

In the law, the concept of actual causation is typically considered to capture the factual and norm-independent aspects of causal inquiry (Wright, 1985, 2001, 2011; Fumerton and Kress, 2001).³ A standard way to spell out this concept employs the 'but-for' criterion: would the harm have occurred but for the defendant's actions? This criterion amounts to counterfactual dependence. It asks us to imagine a counterfactual scenario in which the defendant had not acted in a certain way and asks whether the harm to the victim would still have occurred.

Yet the but-for criterion is too restrictive. Sometimes agents are responsible for

³In Anglo-American jurisdictions 'actual causation' is usually referred to as 'cause-in-fact' or 'factual cause'. Causal modelers have used the concept of actual causation differently, capturing considerations of normality (Hitchcock and Knobe, 2009; Halpern and Hitchcock, 2015). It should also be noted that among legal theorists there is a long-standing discussion as to whether concepts of actual causation can be defined in a principled way and whether they are free of norms and context-dependent considerations (Malone, 1956; Hart and Honoré, 1959). The concept of actual causation is sometimes contrasted with a concept of proximate causation. Here as well, there is controversy as to whether this concept is context dependent and norm dependent. Some authors also refer to this as 'legal cause' or 'scope of liability.' For an overview see e.g. Moore (2009); Knobe and Shapiro (2021).




Figure 1: A simple model for the braking case. The pedestrian is hit by the car because the brakes did not operate. The brakes’ operating was prevented by the rental company’s failure to repair them and the driver’s failure to apply them.

consequences that do not depend counterfactually on their actions, as the braking case illustrates. Here the but-for test is thought to fail because it is assumed that the outcome depends neither on the rental company’s actions nor on the driver’s actions. Legal scholars have considered a range of strategies for dealing with this kind of problem. Some authors (Fischer, 1992; Stapleton, 2008, 2015) suggest that in cases involving causal overdetermination cause-in-fact criteria have to be substituted by normative or policy-dependent solutions. Here we will be concerned with approaches like Wright’s (1985; 2001; 2011; 2022) that treat the braking case as a case of factual causation, and the discussion will be limited to such questions of factual causation.⁴

The braking case involves two structural features. First, it is an example of double prevention. The braking process would have prevented the pedestrian’s being hit by the car. But in the current scenario the pedestrian is being hit because other factors prevent the braking process. Second, the case involves causal redundancy. The braking process was prevented by two independent factors: the driver’s failure and the rental company’s failure, and each one of the two failures is commonly assumed to be individually sufficient for the outcome.⁵

Figure 1 shows a model of the braking case that captures both these features. Such

⁴For another approach to the braking case see Abrams (2022) who (unlike Fischer and Stapleton) argues that causal overdetermination is a matter of factual causation and instead argues that the braking case fails to be a matter of factual causation because it involves omissions.

⁵More generally, double prevention means that there is an effect that would have been prevented if the preventing process would not have been prevented. Two factors c_1 and c_2 are causally redundant with regard to effect e , if both c_1 and c_2 are realized and each would have been sufficient for e . Typical forms of causal redundancy include symmetrical overdetermination and various forms of preemption. In what follows I assume that omissions or failures can act as preventers and can stand in relations of causal redundancy. For an alternative treatment of the omissions in the braking case see Abrams (2022).

a causal model consists of a set of variables and a set of structural equations that summarize the dependencies between the variables. The variables are represented by upper case letters (X, Y, \dots) and typically take on one of two or more possible values (lower case letters x, y, \dots) that represent possible events or states of affairs in the model's target system. In our example RB represents the rental company's actions, with $RB = 1$ if the rental company checks and repairs the brakes and $RB = 0$ if the rental company fails to do so, as is the case in the current situation. Analogously, the model includes a variable AB that represents the driver's actions, with $AB = 1$ if the driver attempts to apply the brakes, and $AB = 0$ if he fails to do so, as is the case in the current situation. It is assumed that the brakes operate if they have been repaired and if the driver applies them: $BO = RB \wedge AB$. Finally, the pedestrian is hit ($PH = 1$) if the brakes do not operate ($PH = \neg BO$), as is the case in the current situation.⁶

What's the actual cause according to this model? There has been a lively debate about the concept of actual causation among authors who employ formal causal models and no agreement has yet been found (Pearl, 2009; Hitchcock, 2001; Woodward, 2003; Halpern and Pearl, 2005; Halpern, 2016; Weslake, forthcoming; Beckers and Vennekens, 2018; Andreas and Günther, 2021; Fischer, 2023, forthcoming). While these proposals differ in their details, the starting point of most approaches in this tradition is to spell out causation in terms of counterfactual conditionals. In situations involving redundancy, actual causation is thought to come down to counterfactual dependence "under certain contingencies" (Halpern and Pearl, 2005, 844). For example, under the contingency that the rental company had checked and repaired the brakes it is the case that the accident depends counterfactually on the driver's failure to operate the brakes; and this is why the driver's failure is an actual cause. Likewise, under the contingency that the driver had tried to apply the brakes, there is a counterfactual dependence of the pedestrian's being hit on the rental company's failure; and this is why the rental company's failure is an actual cause.

The exact result of this kind of analysis depends on the definition that one chooses.

⁶It is assumed that the driver does not know that the brakes are faulty, otherwise we would have to include another causal route leading from the company's failure to the driver's failure because knowledge of the faulty brakes would affect the driver's decision whether to apply the brakes.

According to a popular definition by Halpern and Pearl (2005), both factors are actual causes. According to Halpern's (2016) more recent 'modified' definition, both factors are *parts of* an actual cause. This move to 'parts of and actual cause' helps simplify definitions of actual causation and helps account for a number of problematic cases, but it requires explicating the relevant sense of parthood.

What matters for our purposes, is that any definition that identifies the driver's failure as actual cause based on this model will have to also identify the rental company's failure (and vice versa). This is because the causal model is symmetric with respect to the corresponding variables. Symmetry here means simply that the model is invariant under exchanging variables *RB* and *AB*: if *RB* and *AB* switched roles in the causal model we would end up with the same causal model (for more on such symmetry considerations see section 6).

3 Is the Simple Model Too Simple?

Here is an immediate objection to the analysis suggested above: the causal model treats the driver's failure and the rental company's failure completely on a par. We could simply exchange the variables *RB* and *AB* in the graph and the corresponding structural equations and end up with the same model. It is therefore not at all surprising that if one of the failures is identified as an actual cause that the other failure is an actual cause as well. In fact, it seems that I have built the result of my analysis into the model. No one who thinks that there are relevant causal differences between the failures will be convinced by my analysis.

This concern is an instance of a broader concern that seems to put pressure on the method I am suggesting here. A widely acknowledged feature of model-based definitions is their model dependence. What counts as an actual cause in a target system depends on certain choices that we have to make in modelling that target system. This may give rise to a worry regarding the applicability of the framework in the legal context. Asked whether an action was in fact a cause of the harm would yield different answers depending on which of the (likely opposing) models is selected. In order to resolve such

disagreement, it seems, we would need objective criteria for choosing the best model.

Causal modelers have developed several strategies that one might think could be employed to respond to this worry. They have studied conditions under which causal verdicts remain stable even if the model is changed, and they have tried to provide general criteria for appropriate models (Halpern and Hitchcock, 2010; Halpern, 2016; Woodward, 2016). But in many cases the criteria leave model choices underdetermined. In what follows we will see that the framework of causal modelling benefits inquiry in a more indirect way. The framework can be seen as suggesting structural features of the target system that need to be realized to make a particular causal claim plausible. These suggestions, I will argue, can play a heuristic role for collecting evidence that is relevant for evaluating the defendant's causal role.

4 Wright's Approach to Causal Priority

But before addressing the heuristic role in more detail, let's look at the idea that the driver has causal priority. The most detailed defense of this idea has been provided by Wright.⁷ The basic idea of Wright's account is that "*a particular condition was a cause of (condition contributing to) a specific consequence if and only if it was a necessary element of a set of antecedent actual conditions that was sufficient for the occurrence of the consequence*" (1985, 1790, emphasis original). The acronym NESS summarizes the core idea that a cause is a "necessary element of a sufficient set" (1790) of the conditions for the effect. The NESS account goes back to Hart and Honoré (1959) and in the philosophical literature Mackie's (1974) INUS condition is a better-known development of it.

The NESS account essentially weakens the well-known but-for criterion. The cause does not need to be necessary for the effect per se, but only needs to be a necessary element of a set of conditions that are jointly sufficient for the effect. The driver's failure is not necessary for the accident's occurrence. But the driver's failure to operate the

⁷Wright has discussed the case in a series of papers and his approach has changed in response to criticism (Wright, 1985, 2001, 2011). Here I focus on the main argument as given in the 2011 paper as it contains the most detailed discussion. In his most recent discussion Wright has changed position, now arguing that the rental agency has causal priority over the driver (Wright, 2022). For a discussion of this recent shift in Wright's view see below.

brakes is a necessary element of one set of antecedent actual conditions (including the car's approaching the pedestrian with sufficient speed, there not being road structures protecting the pedestrian...), and that set is taken to be sufficient for the pedestrian's being hit. So, Wright's account identifies the driver as actual cause. But, likewise, the company's failure to repair the brakes can be construed as a necessary element of a set of conditions that are sufficient for the pedestrian's being hit.

Therefore, Wright argues that we should look more closely at the braking process. This is a process that consists of several stages, including the braking pedal being depressed by the driver, braking fluid being transmitted to impose pressure on the brake cylinders, and the friction created by the braking pads which leads to slowing the rotation of the wheels (2011, 318). What matters in cases like these, according to Wright, is where the braking process is first interrupted (because this supposedly affects where a complete set of sufficient conditions can be instantiated along the causal path). The first interruption, Wright (2011) argues, is the driver's failure. The rental company's failure would come into play only at a later stage of the braking process. But since the braking process never reaches that stage, the rental company's failure is not an actual cause at all. Thus, it seems, Wright construes the case as instantiating a variety of preemption, such that the driver's failure preempts the rental company's failure.

In preemption cases two (or more) causal processes have the potential to bring about the outcome. One process causes the outcome (the preempting process), while the other process does not cause the outcome (the preempted process). Philosophers and legal scholars have discussed a variety of preemption cases, involving early preemption (where the preempted process is interrupted before the preempting process is complete), late preemption (where the preempted process is interrupted through or after completion of the preempting process), and trumping preemption (where one process trumps another process, such as the order of a higher ranked officer trumping the order of a lower ranked officer).

Wright himself sees the braking case in analogy to cases of preemptive prevention (McDermott, 1995; Collins, 2000).⁸ Consider a ball being thrown towards a window.

⁸More specifically, Wright (2011, 316f) characterizes both the braking case and the fielder cases to be

There is a fielder who catches the ball and another fielder who would have caught the ball if the first fielder had failed to do so. Typically, the first fielder is thought to prevent the window's being hit, thereby preempting the second fielder's attempt of prevention. This is captured by Wright's account because it was the first fielder who interrupted the causal process of the ball's approaching the window.

Things get more complicated, though, if we replace the second fielder by a solid brick wall which would block the ball if the remaining fielder failed to catch it. Wright argues that in this kind of scenario it is still the first interrupting factor that has to be seen as the actual cause (or actual preventer) of the result, regardless of the details of the backup. So, he is committed to arguing that the fielder prevented the window from shattering even if the window was protected by a solid brick wall. This is implausible.

Why do we identify the fielder more easily as a preventer when the backup is another fielder? Presumably, that is because we can imagine a situation in which not just the first fielder but also the second fielder fails to catch the ball, such that the window would be hit after all. Given what we know about the two fielders, it is a realistic scenario that both fail to catch the ball. This is different if the backup is a solid brick wall. No matter what the remaining fielder does, the brick wall will guarantee that the window will not be hit by the ball. It is simply not a realistic scenario that the ball will get past the wall.

What matters for our discussion of Wright's approach is that, arguably, the braking case has more similarity with the scenario involving the wall. Given what we know about the faulty brakes, it is simply too unrealistic to assume that they would have worked if the driver had tried to apply them. Thus, it would be problematic to argue for the driver's causal priority on account of the situation's similarity to other instances of preemptive prevention, if that is what Wright attempts to achieve by invoking these examples.⁹

In summary, Wright approaches the braking case in analogy to instances of supposed

discussed below as examples of "overdetermined negative causation." Whether and to what degree the analogy holds may be contested, see Fischer (2021a).

⁹In a more recent paper Wright (2022) agrees with Ingeborg Puppe (see e.g. Puppe (2013)) that the brakes have causal priority because their failure has temporal priority over the driver's failure. But again: construing the braking case in analogy to relevant instances of preemptive prevention would not make this claim more plausible. We wouldn't say that the wall prevented the window from being hit if the ball was caught by the fielder before it could reach the wall.

preemptive prevention. Preemptive prevention means that one factor preempts another in preventing the outcome and, thus, the causal symmetry between the competing factors is broken. But here I have argued that the analogy to such cases fails to achieve this argumentative goal.

5 The Signature of Preemption

I have argued that Wright's analogy fails to support his claim of the driver's causal priority. But maybe this is a problem of the specific analogy that Wright chooses to explain his argument. In this section we will put the framework of causal models to work to explore in more detail why such claims of causal priority seem to be difficult to establish in the context of the braking case. The main idea is that preemption cases have a specific signature in causal models: a counterfactual dependence between the assumed cause and the effect can be restored if certain other variables are kept fixed in their current value. But we will see that, given the current description of the braking case, it seems difficult to realize this signature. While this might not fully disprove the driver's causal priority, it gives at least prima facie reason against it.

Consider how causal modelers address instances of preemption. For ease of exposition consider the following case of early preemption (Hitchcock, 2001).¹⁰ Two assassins, Suzy and Billy, both set off to a mission to kill a victim. Suzy shoots the victim, and the victim dies. Seeing that Suzy has killed the victim, Billy refrains from pulling the trigger of his gun. In this example we typically identify Suzy's actions as actual cause but not Billy's actions. How do we capture this?

A very simple model of the preemption case is provided in figure 2A. According to this model, both Billy and Suzy set off to their respective missions (*BM* and *SM*) and the victim dies (*VD*). This model, of course, fails to capture the causal difference between Billy and Suzy. It does not represent the fact that Suzy's pulling the trigger leads to Billy's refraining from doing so.

Figure 2B displays a more accurate way to model the assassination case. $SM = 1$ and

¹⁰The argument carries over to other forms of preemption, if modeled in analogy.

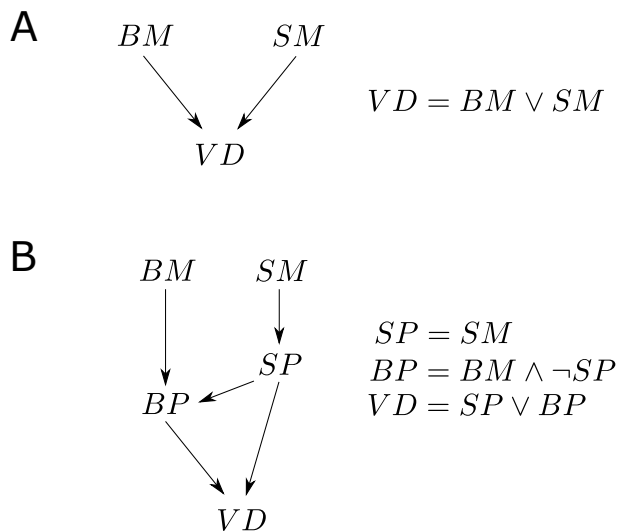


Figure 2: A standard case of causal preemption. A: Simplified model omitting relevant detail. B: More sophisticated model including relevant detail.

$BM = 1$ represent the fact that Suzy and Billy both set off to their mission, respectively. Suzy pulls the trigger of her gun if she sets off to the mission ($SP = SM$) and Billy pulls the trigger of his gun if he sets off to the mission and Suzy has not pulled the trigger of her gun ($BP = BM \wedge \neg SP$). Finally, the victim dies if either Suzy or Billy pull the trigger of their guns ($VD = SP \vee BP$).

Let's apply Halpern's (2016) definition of actual causation. According to this definition, a set of variables with value assignments $\vec{X} = \vec{x}$ is an actual cause of φ in a causal setting (described by a model M with assignments to the model's input variables) if three conditions hold. First, in the current causal setting the cause and the effect have to occur. Second, we need to be able to establish a counterfactual dependence of φ on $\vec{X} = \vec{x}$. To do so we may hold fixed some of the variables in M in their actual values (the values that the variables take on in the original/actual scenario). Third, the set of variables \vec{X} needs to be minimal, that is, no proper subset of $\vec{X} = \vec{x}$ may fulfil the first two conditions.

Thus, according to the more sophisticated model and Halpern's definition, Suzy is an actual cause because (1) she fires her gun in the current situation and the victim dies, (2) keeping fixed the fact that Billy does not pull the trigger of his gun one can establish a




Figure 3: The braking case represented as an instance of causal preemption. The driver’s actions are thought to preempt the rental company’s failure. But the model appears to misrepresent the causal structure.

counterfactual dependence of the victim’s death on Suzy’s actions, and (3) Suzy’s firing her gun is the minimal set of events fulfilling conditions (1) and (2). Billy, by contrast, is not an actual cause because there is no variable that could be held fixed in order to establish such a counterfactual dependence.

Note that the apparent situation of symmetric overdetermination was turned into a preemption situation by including more detail in the causal model. This suggests a more general strategy for looking whether a supposed case of symmetric overdetermination may be construed as an instance of preemption: look whether the symmetry between apparently overdetermined factors can be broken such that a counterfactual dependence between the supposed preempter and the effect can be restored if an off-path variable can be held fixed in its current value.

Let’s try this strategy in the braking case. Simply adding an additional edge linking AB and PH in the simple model, as suggested by Baldwin and Neufeld (2004), does not help. Keeping fixed the fact that the brakes do not operate $BO = 0$, there is no way the pedestrian’s being hit depends on the driver’s attempting to apply the brakes.

The situation looks slightly better if we replace variable BO by a variable BF representing the brakes’ failure (figure 3). The brakes fail if the rental company does not repair them, and the driver attempts to apply them: $BF = \neg RB \wedge AB$. In the current situation the brakes are faulty (they have not been repaired) but they do *not* fail because the driver does not try to apply them. In this new causal model, the pedestrian’s being hit depends on the brake’s failure and the driver’s actions: $PH = BF \vee \neg AB$.

This model would give the right verdict according to Halpern's definition.¹¹ First, the driver's failure is an actual cause of the pedestrian's being hit. In the actual scenario the brakes do not fail ($BF = 0$). Keeping fixed this fact, we can establish a counterfactual dependence of the pedestrian's being hit on the driver's actions. Second, the rental company's failure is not an actual cause. Whatever variables we hold fixed, there is no way to establish a counterfactual dependence of PH on RB in this model.

But the problem here is that the model is not an appropriate representation of the current situation. First, one might worry that it is a questionable framing to say that brakes only fail if someone tries to apply them. Note that it is exactly this framing that would allow Wright's account discussed in the foregoing section to achieve the driver's causal priority. Second, it is unclear what causal mechanism would correspond to the direct link between the driver's failure to operate the brakes and the pedestrian's being hit. It seems like any causal process that would link these two variables would have to go through the variable that describes whether the brakes are functional.

In this section I have examined whether the signature of preemption can be realized in causal models of the braking case. The difficulties encountered here are not conclusive evidence against some form of causal priority. But they put additional pressure on accounts like Wright's account. More generally, the discussion illustrates an aspect of the heuristic value of causal models: they can help shed new light on legal cases if we employ them to look for specific signatures, such as the signature of preemption.¹²

6 Can the Symmetry Be Broken?

While the foregoing discussion assumed a certain framing of the case, the discussion in this section and the following, will be concerned with ways in which the provided

¹¹It would also reproduce this verdict on other definitions of causation, in particular, on a suitably formalized definition of Wright's NESS criterion, such as the one provided by Beckers (2021).

¹²Could one construct an asymmetric model by 'merging' the models in figures 1 and 3? The resulting model would include both variables BO and BF and potentially a direct link from BO to PH that could bypass BF . The resulting model would then potentially establish causal priority for the driver. However, the resulting model would be problematic. First, the relation between BF and BO would be a conceptual rather than a causal relation. This has consequences for interpreting any potential link from BO to PH bypassing BF : any model that suggests that the brakes operate to prevent the pedestrian's being hit while it also states that the brakes are faulty seems unpalatable. Thanks to an anonymous reviewer for pointing out this possibility.

framing may have to be improved. Here another benefit of the causal modelling framework is that it allows for a more systematic exploration of the various ways in which a causal evaluation of the braking case and similar cases may be overturned. A key role in such thinking is played by symmetry considerations. Symmetry means that replacing one variable for another and changing the structural equations accordingly does not change the model. If a model is invariant regarding such exchanges of variables, then the causal judgement needs to be the same with regard to both variables. This is because definitions such as Halpern's definition of actual causation usually take only structural features of the model into account but not intrinsic properties of the individual variables.¹³

The simple model of the braking case discussed in section 2, for instance, is symmetric with regard to the rental company's failure and the driver's failure. That is, exchanging the corresponding variables and adjusting the structural equations accordingly would yield the same causal model. And since Halpern's definition of actual causation is insensitive to such changes, there won't be any change in causal judgement. This implies that the two factors represented by the variables (the driver's failure and the rental company's failure) should be taken to be causally on a par.

Let's distinguish two ways causal inquiry can proceed. First, inquiry can proceed from causal model to causal judgement. Here one would first take the target system and try to provide an appropriate model of it. At a second stage one would include information about the actual values of the model variables in the current situation and apply a definition of actual causation to find out the actual cause. We shall call this an open causal inquiry. Second, inquiry can proceed from causal hypothesis to causal model. Starting with a causal hypothesis one may consider causal models that would agree with this hypothesis. At a second stage one would collect evidence to find out whether that model is an apt representation of the target system. We shall call this a hypothesis-driven inquiry.

In the law, hypothesis-driven inquiry is more common. Typically, legal inquiry does

¹³An exception are definitions of actual causation that take into account a distinction between default values and deviant values of variables (Halpern and Hitchcock, 2015).

not start with a model of the target system to address an open question of causation. Instead, causal questions are more specific as they ask whether the defendant's wrongful action actually caused the harm in question.

Let's focus on the causal role of the driver's wrongful action in the braking case. Can the driver's wrongful action be considered the sole actual cause of the accident? Are there features of the causal structure that would imply that the driver's wrongful action was not a cause of the accident at all? Starting from the symmetric model that I have provided in the foregoing discussion, there are various ways such claims could be established, by showing that the symmetry between variables RB and AB is disturbed.

First, challenging one of the arrows feeding into variable BO would undermine the symmetry. This would amount to questioning, for instance, whether the brakes really were faulty when the car was provided to the driver. Alternatively, one might want to challenge whether the driver really failed to apply the brakes at the right moment.

Second, the foregoing discussion of preemption cases suggests that the symmetry could be disturbed by adding an extra direct link between the driver's actions and the variable representing whether the pedestrian is being hit by the car. Figure 4 represents such a situation. This model is similar to the initial model suggested in section 2, but here the variable representing the driver's failure is replaced by a three-valued variable DA that describes the driver's actions and allows for the further option of diverting the car ($DA = 0$: the driver attempts nothing to prevent the accident; $DA = 1$: the driver attempts to brake; $DA = 2$: the driver diverts the car).¹⁴ According to this model, the brakes operate if the rental company repairs them, and the driver applies them. But according to this model, the accident can also be prevented if the car is diverted. This causal route is independent of the brakes' proper functioning. The structural equation

¹⁴The driver could, of course, divert the car *and* try to use the brakes. This would require a different model, according to which the pedestrian's being hit simply depends on the driver's diverting and simultaneous braking. Note also that the case may be usefully compared with a case where the driver is placed in a trolley approaching the victim. Suppose that the driver does not know that the brakes of the trolley are broken but nevertheless he does not try to operate them. Since braking seems to be the only (apparent) way to prevent the harm, it is more difficult in this instance to break the symmetry in the model.




Figure 4: The driver would be the sole actual cause if he had additional options for control over the outcome. The dashed edge indicates the hypothetical nature of the suggested link between *DA* and *PH*.

for the brakes' operating then is:

$$BO = \begin{cases} 1 & \text{if } RB = 1 \wedge DA = 1 \\ 0 & \text{else,} \end{cases}$$

and the structural equation for the pedestrian's being hit is:

$$PH = \begin{cases} 0 & \text{if } BO = 1 \vee DA = 2 \\ 1 & \text{else.} \end{cases}$$

According to this model, only the driver's failure is an actual cause of the pedestrian's being hit, but not the rental company's failure. Holding fixed that the brakes do not operate in the current situation, we can establish a counterfactual dependence of the pedestrian's being hit on the driver's actions (exploiting that the driver can divert the car).¹⁵ For the rental company's failure such a dependence cannot be established in this model. There is no corresponding way the accident can be prevented in the current model, if we keep fixed the fact that the driver fails to apply the brakes.

Finally, the symmetry could be disturbed by adding another edge between the two candidate causes that are supposed to be independent in the original description of the case. For example, one potential reason for the driver's not trying to use the brakes

¹⁵The pedestrian's being hit will depend counterfactually on the driver's failure to divert or the driver's actions in general. It will not depend on the failure to brake, as before. The difference may matter for legal purposes because one of the driver's failures may be negligent whereas the other is not.

could have been that he learned of them being faulty right before the accident happened.

This list, of course, is not complete as further variables may need to be included in the causal model, potentially giving rise to further options for disturbing the symmetry. Yet the list illustrates how playing around with the graphical representation can lead to a more systematic way of generating hypotheses regarding the driver's causal status.

One might wonder what exactly the benefit of employing the formal framework is in such cases. Specifically, one might object that for pointing out the above possibilities the formal framework is not strictly necessary. Any of the potential challenges to the initial causal hypothesis may have been raised without the formal machinery. For instance, the suggestion that the driver's option for diverting the car would overturn the judgement about the driver's causal status is quite commonsensical.¹⁶

The causal model framework, thus, seems to play a largely expository role. However, the importance of such an expository role should not be underestimated. In a trial, legal decision makers need to evaluate the plausibility of competing explanations. Crucially, such evaluations do not just involve inferences from causal structure to judgement of actual causation. Primary subject of inquiry is the evidential status of potentially competing accounts of the causal structure of the legal case. Especially in the context of jury trial this is a complex and multi-stage procedure. Large amounts of evidence have to be taken into account. The evidence is often presented over several days and in a disconnected form due to the question-and-answer format and the fact that different witnesses testify with regard to aspects of the case in no particular order (Hastie et al., 1983).

I have argued that in cases involving the kind of causal redundancy encountered in the braking case, symmetry considerations can act as a useful guide for exploring relevant causal features of the case. But isn't legal inquiry rather driven by normative questions, especially, regarding the driver's duties? Several potential duties come to mind. Considering defendant's duty to pay attention, we look at what would have happened had the defendant paid attention. Considering defendant's duty to brake, we

¹⁶Thanks to an anonymous referee for pointing this out.

look at what would have happened had defendant tried to operate the brakes.¹⁷

Such normative considerations certainly play an important role in guiding legal inquiry. Crucially, however, not any kind of normative consideration will be useful in guiding such inquiry. The kinds of wrongful action that are being considered here will have to be relevant for the harm under consideration. And such considerations of relevance will often be based on preliminary assumptions about the causal relation of the wrongful action and the harm in question. Therefore, employing symmetry considerations such as those discussed above and using causal models in an expository way, may help raise relevant normative questions in the first place.

In this context, causal models will guide inquiry by helping structure that evidence and highlighting those parts of causal structure that are particularly relevant for deciding the overall causal judgement. In cases like the braking case this heuristic role will be supported by considerations of symmetry. If a model is symmetric with regard to two variables, then the overall causal judgement should treat the corresponding factors on a par, or otherwise there should be included further structural features that account for the difference.¹⁸

7 Extending the Analysis

The braking case is typically discussed as an instance of a broader class of cases that involve the non-use or misuse of malfunctioning or missing safety devices. But whether specific causal judgements in the braking case carry over to the other cases, I will argue, crucially depends on the details of such cases. This showcases another heuristic use of causal models and their graphical representations, this time by highlighting analogies and disanalogies between seemingly similar cases.

More specifically, the heuristic value derives from another instance of symmetry

¹⁷Thanks to an anonymous referee for raising this point.

¹⁸Again, one may argue that sometimes causal judgement is (legitimately) affected by aspects intrinsic to a specific variable, such as an assignment of a default or deviant value (Halpern and Hitchcock, 2015). Such default and deviant values have been argued to reflect considerations of normality that play important roles in causal judgement. However, legal scholars such as Wright would object that such considerations of normality would not concern supposedly factual considerations of actual causation (Wright, 1985). It should be noted, though, that it is an open question whether a neat line can be drawn between such considerations of actual causation and considerations of normality (Fischer, 2021a).

considerations—this time not within an individual model but between models. If two cases are aptly modelled by two models that are structurally isomorphic, then the same overall causal judgements should apply to the two cases. Structural isomorphism means that replacing the variables of one model by the variables of the other model and adjusting the structural equations accordingly should yield the same model. And the overall causal judgements applying to the models should be the same because a definition such as Halpern’s definition employed above is sensitive only to such structural features.¹⁹ Conversely, if the intuitive causal judgements with regard to two structurally similar cases come apart, then this should be a warning signal that the cases may be less similar after all, and that additional evidence is needed.

First, consider the case in which a hotel failed to provide a fire escape but the deceased hotel guest failed to look for a fire escape and dies.²⁰ This case appears to be structurally isomorphic to the braking case because, first, it involves double prevention: taking the fire exit would have prevented the hotel guest’s death. But other factors prevented the hotel guest’s taking the fire exit. In addition, the case involves redundancy because taking the fire exit is prevented by the absence of the fire exit and also by the hotel guest’s failure to look for one.

So, it seems a plausible model of the hotel case would be similar to the one presented for the braking case in section 2. According to this model both the hotel owner and the deceased guest would be actual causes of the death. Are there ways to disturb the symmetry between the two factors? The hotel guest would be the sole actual cause, for instance, if on the level of the model one could include a direct causal link between the guest’s actions and his death. And, in fact, the details of the case suggest that such action would have been possible if the hotel guest had not locked himself in his room or had left the building through one of the windows as other guests did. So, it seems that arguing for the causal priority of the hotel guest would be possible on similar grounds as for arguing for the causal priority of the driver.

Second, there is a case in which a prefabricated fire place was incorrectly installed by

¹⁹See Hall (2007); Halpern and Hitchcock (2015); Blanchard and Schaffer (2017); Fischer (2021b) for discussion of such isomorphisms.

²⁰*Weeks v. McNulty* 101 Tenn. 495, 48 S.W. 809 (1898).

a carpenter, but the manufacturer also had failed to include relevant warning instructions.²¹ The basic structural features seem to be again the same. First, the damage could have been prevented if the carpenter had read a warning-label containing information that the prefabricated fireplace has to be installed in a certain way. But the carpenter's reading such instructions was prevented. Second, the carpenter's reading the label was prevented by two factors each of which would have been sufficient: the manufacturer did not include the correct warning instructions and the carpenter did not intend to read such instructions.

From this one would have to conclude that the fire-place case is best modelled by a simple symmetric model as proposed in section 2, and that both the manufacturer and the carpenter are actual causes. But in this case the symmetry between the two factors is maybe more easily disturbed: the carpenter would have the status of the sole actual cause if we included an additional edge connecting the carpenter's actions directly to the outcome. That is, it should be examined what degree of control the carpenter has over the way the fireplace is installed, and whether the carpenter could have had independent ways to know that the fireplace has to be installed in a certain way in order to prevent fire.

Finally, consider the following hypothetical failure-to-warn case discussed by Fischer (2006). Suppose a painter throws an aerosol spray paint can into the trash and the can explodes while the trash is being burned. Suppose also that the company that manufactured and sold the can placed a label on the can that warned about the risk of inhaling paint fumes. Fischer's case also emphasizes that this warning label is very conspicuous. But the label contains no warning that exposing the can to fire could cause the can to explode. The painter is assumed to throw the can into the bin without having read the warning label. Fischer observes that a majority of test subjects presented with this case had the intuition that the failure to warn had causal priority.

Like the other failure-to-warn cases this case has a causal structure that is similar to that of the braking case (see figure 5): the burning of the can ($BC = 1$) was not prevented because the painter did not recognize the danger ($RD = 0$). This is because

²¹*Safeco Insurance Co. v. Baker*, 515 So. 2d 655 (Ct. App. 1987), cert. denied, 519 So. 2d 130 (La. 1988).




Figure 5: In the painting case, damage could have been prevented if the painter had recognized them. But this did not happen because the paint can was not correctly labelled, and the painter failed to read the label. The painting case is ambiguous because placing a conspicuous label may have attracted the painter’s attention (dashed edge).

the can was not appropriately labelled ($AL = 0$) and the painter failed to read the label ($RL = 0$). But the intuition that the failure to warn had causal priority should act as a warning signal that the case may not be isomorphic to the case (as modelled by the simple model in figure 1 after all. The formal framework suggests that including an additional direct causal route from the manufacturer’s failure to the painter’s failure would overturn the causal judgement. In the resulting causal model, the manufacturer would clearly have causal priority over the painter. And indeed, such an additional edge might be reasonably argued for, because it is plausible that a conspicuous warning sign of explosion would have attracted the painter’s attention.

In conclusion, the modelling framework plays a heuristic role also in the comparison between seemingly similar legal cases. A key role is again played by symmetry considerations. If two cases are modelled by structurally isomorphic models, then the cases should give rise to (structurally) the same causal judgements. If the intuitive judgements, however, seem to come apart, then this should be a warning signal to the effect that the cases may not be structurally isomorphic after all, and that additional evidence is needed to make the corresponding causal claims plausible.

Here is a potential worry. Like any other kind of formal representation, causal-model-based representations necessarily leave out details of legal cases. Specifically, the highly abstract representation in terms of variables and structural equations may sometimes

appear ill-suited to represent the details of defendant's actions or failures to act. Don't we face the risk that causally relevant information gets lost?²²

The use of causal models and considerations of symmetry and isomorphism proposed here is that of a heuristic device. The advantage of the advocated strategy is that it may help to highlight differences between candidate causes within one case (and multiple candidate factors across cases) by examining structural similarities and differences. In this context the formal framework is highly flexible, and, arguably, more useful than other kinds of frameworks (see e.g. Hitchcock (2007)), but certainly it also faces important limitations. This would be a problem for the account advocated here if it aimed at definite criteria for actual causation in the law. Yet, I argue, that the benefit is rather that of a useful heuristic device. Such a heuristic device will be helpful to provide new perspectives on cases with complex causal structure, but it will always have to be complemented with and checked against other means of causal judgement.

8 Conclusion

I have started with an in-depth analysis of the braking case, arguing that it involves double prevention and causal redundancy. I have provided a simple model of the case, according to which both the rental company's failure and the driver's failure are causally on a par. A potential objection to this approach is that it builds in the symmetry between the rental company and the driver on the level of the causal model. It is then no surprise, one could object, that the rental company and the driver are causally on a par. In response to this objection, I have critically analyzed a prominent alternative way of framing the braking case which takes the driver's actions as preempting the rental company's failure. I have shown that such a construal of the braking case is implausible given the standard description of the case.

One might think that a definition of actual causation should not be model-dependent if it is to be employed in legal inquiry. Agents who disagree about judgements of actual causation are likely to disagree also about the best causal models. Looking at cases that

²²Thanks to an anonymous referee for raising this worry.

involve the non-use or misuse of missing or malfunctioning safety devices suggests that such disagreement can only be solved by looking at the causal details of the individual cases. But my discussion of the braking case and associated examples has illustrated that the modelling framework benefits causal inquiry as a heuristic tool: it can help us highlight those assumptions about causal structure that matter for our verdicts of actual causation. More specifically, causal models provide a framework (1) for looking for certain signatures of causal priority, (2) for seeking how apparent symmetries between factors can be disturbed (thereby helping to generate hypotheses for the search for relevant evidence), and (3) for comparing cases with supposedly similar structure.

References

- Abrams, Yuval. Omissive overdetermination: Why the act-omission distinction makes a difference for causal analysis. *University of Western Australia Law Review*, 49(57):57–86, 2022.
- Andreas, Holger and Günther, Mario. A Ramsey test analysis of causation for causal models. *The British Journal for the Philosophy of Science*, 72(2):587–615, 2021.
- Baldwin, Richard A. and Neufeld, Eric. The structural model interpretation of the ness test. In Tawfik, Ahmed Y. and Goodwin, Scott D., editors, *Advances in Artificial Intelligence*, pages 292–307, 2004.
- Beckers, Sander. The counterfactual NESS definition of causation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6210–6217, 2021.
- Beckers, Sander and Vennekens, Joost. A principled approach to defining actual causation. *Synthese*, 195(2018):835–862, 2018.
- Blanchard, Thomas and Schaffer, Jonathan. Cause without default. In Helen Beebe, Huw Price, Christopher Hitchcock, editor, *Making a Difference*, pages 175–214. Oxford University Press, Oxford, 2017.

- Bochman, Alexander. Actual causality in a logical setting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 1730–1736. International Joint Conferences on Artificial Intelligence Organization, 2018.
- Chockler, Hana and Halpern, Joseph. Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research*, 22:93–115, 2004.
- Chockler, Hana, Fenton, Norman, Keppens, Jeroen, and Lagnado, David A. Causal analysis for attributing responsibility in legal cases. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law, ICAIL '15*, pages 33–42, 2015.
- Collins, John. Preemptive prevention. *The Journal of Philosophy*, 97(4):223–234, 2000.
- Fischer, David A. Causation in fact in omission cases. *Utah Law Review*, pages 1335–1384, 1992.
- Fischer, David A. Insufficient causes. *Kentucky Law Review*, 94(2):277–317, 2006.
- Fischer, Enno. *Actual Causation*. PhD thesis, Leibniz University Hannover, 2021a.
- Fischer, Enno. Causation and the problem of disagreement. *Philosophy of Science*, 88(5): 773–783, 2021b.
- Fischer, Enno. Actual causation and the challenge of purpose. *Erkenntnis*, 2023.
- Fischer, Enno. Three concepts of actual causation. *The British Journal for the Philosophy of Science*, forthcoming.
- Fumerton, Richard and Kress, Ken. Causation and the law: Preemption, lawful sufficiency, and causal sufficiency. *Law and Contemporary Problems*, 64(4):83–105, 2001.
- Hall, Ned. Structural equations and causation. *Philosophical Studies*, 132:109–136, 2007.
- Halpern, Joseph Y. Defaults and normality in causal structures. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning, KR'08*, pages 198–208. AAAI Press, 2008.

- Halpern, Joseph Y. Cause, responsibility and blame: a structural-model approach. *Law, Probability and Risk*, 14(2):91–118, 2015.
- Halpern, Joseph Y. *Actual Causality*. MIT Press, Cambridge, MA, 2016.
- Halpern, Joseph Y. and Hitchcock, Christopher. Actual causation and the art of modeling. In Dechter, Hector, Geffner, Rina, and Halpern, Joseph Y., editors, *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*, pages 383–406. College Publications, London, 2010.
- Halpern, Joseph Y. and Hitchcock, Christopher. Graded causation and defaults. *The British Journal for the Philosophy of Science*, 66:413–457, 2015.
- Halpern, Joseph Y. and Pearl, Judea. Causes and explanations: A structural-model approach. Part I: Causes. *The British Journal for the Philosophy of Science*, 56(4):843–887, 2005.
- Hanley, Brian J. What caused the bhopal gas tragedy? the philosophical importance of causal and pragmatic details. *Philosophy of Science*, 88(4):616–637, 2022.
- Hart, H. L. A. and Honoré, A. M. *Causation in the Law*. Clarendon Press, Oxford, 1959.
- Hastie, Reid, Penrod, Steven D., and Pennington, Nancy. *Inside the Jury*. Harvard University Press, 1983.
- Hitchcock, Christopher. The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*, 98(6):273–299, 2001.
- Hitchcock, Christopher. What’s wrong with neuron diagrams? In Campbell, Joseph Keim, O’Rourke, Michael, and Silverstein, Harry S., editors, *Causation and Explanation*. The MIT Press, 2007.
- Hitchcock, Christopher and Knobe, Joshua. Cause and norm. *The Journal of Philosophy*, 106(11):587–612, 2009.
- Knobe, Joshua and Shapiro, Scott J. Proximate cause explained: An essay in experimental jurisprudence. *University of Chicago Law Review*, 88(1):165–236, 2021.

- Lagnado, David A. and Gerstenberg, T. Causation in legal and moral reasoning. In *Oxford Handbook of Causal Reasoning*, pages 565–602. Oxford University Press, Oxford, 2017.
- Lagnado, David A., Fenton, Norma, and Neil, Martin. Legal idioms: a framework for evidential reasoning. *Argument and Computation*, 4(1):46–63, 2013.
- Liefgreen, Alice and Lagnado, David A. Drawing conclusions: Representing and evaluating competing explanations. *Cognition*, 234:105382, 2023.
- Liepina, Ruta, Sartor, Giovanni, and Wyner, Adam. Causal models of legal cases. In Pagallo, Ugo, Palmirani, Monica, Casanovas, Pompeu, Sartor, Giovanni, and Villata, Serena, editors, *AI Approaches to the Complexity of Legal Systems*, pages 172–186, 2018.
- Mackie, John Leslie. *The Cement of the Universe: A Study in Causation*. Clarendon Press, Oxford, 1974.
- Malone, Wex. Ruminations on cause-in-fact. *Stanford Law Review*, 9(1):60–99, 1956.
- McDermott, Michael. Redundant causation. *The British Journal for the Philosophy of Science*, 46(4):523–544, 1995.
- Metz, Joseph. Preemptive omissions. *Erkenntnis*, 2022.
- Moore, Michael. *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford University Press, Oxford, 2009.
- Pearl, Judea. *Causality. Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, second edition, 2009.
- Puppe, Ingeborg. The concept of causation in the law. In Kahmen, Benedikt and Stepanians, Markus, editors, *Critical Essays on "Causation and Responsibility"*, pages 67–109. De Gruyter, Berlin, 2013.
- Stapleton, Jane. Choosing what we mean by 'causation' in the law. *Missouri Law Review*, 73(2):433–380, 2008.

- Stapleton, Jane. An 'extended but-for' test for the causal relation in the law of obligations. *Oxford Journal of Legal Studies*, 35(4):697–726, 2015.
- Statham, Georgie. Normative commitments, causal structure, and policy disagreement. *Synthese*, 197:1983–2003, 2020.
- Weslake, Brad. A partial theory of actual causation. *The British Journal for the Philosophy of Science*, forthcoming.
- Woodward, James. *Making Things Happen*. Oxford University Press, Oxford, 2003.
- Woodward, James. The problem of variable choice. *Synthese*, 193:1047–1072, 2016.
- Wright, Richard W. Causation in tort law. *California Law Review*, 73(6):1735–1828, 1985.
- Wright, Richard W. Once more into the bramble bush: Duty, causal contribution, and the extent of legal responsibility. *Vanderbilt Law Review*, 54(3):1–51, 2001.
- Wright, Richard W. The NESS account of natural causation: A response to criticisms. In Goldberg, R., editor, *Perspectives on Causation*, pages 265–322. Hart Publishing, Oxford, 2011.
- Wright, Richard W. Causation (contribution) and the 'no worse off' limitation on liability. *University of Western Australia Law Review*, 49(1):5–56, 2022.