

Hasteners and delayers: why rains don't cause fires

Caroline Torpe Touborg¹

Published online: 26 May 2017

© The Author(s) 2017. This article is an open access publication

Abstract We typically judge that hasteners are causes of what they hasten, while delayers are not causes of what they delay. These judgements, I suggest, are sensitive to an underlying metaphysical distinction. To see this, we need to pay attention to a relation that I call positive security-dependence, where an event E security-depends positively on an earlier event C just in case E could more easily have failed to occur if C had not occurred. I suggest that we judge that an event C is a cause of a later event E only if E security-depends positively on C. This explains our causal judgements in typical cases of hastening and delaying as well as in atypical cases, where we judge that hasteners are not causes of what they hasten or that delayers are causes of what they delay.

Keywords Causation · Hastening · Delaying

1 Introduction

There is a puzzling asymmetry in our causal judgements about hasteners and delayers: we typically (though not always) judge that hasteners are causes of what they hasten, while delayers are not causes of what they delay.¹ This asymmetry is seen, for example, in our judgements about the following two cases:

¹ This asymmetry was first noticed by Bennett (1987: 373), and received its first full-length treatment in Mackie's article 'Causing, delaying, and hastening: Do rains cause fires?' (1992). For further discussion, see in particular Sartorio (2006), and Yablo (2004).

✉ Caroline Torpe Touborg
ctt@st-andrews.ac.uk; caroline.touborg@gmail.com

¹ Department of Philosophy, University of St Andrews, 17-19 College Street, St Andrews, Fife KY16 9AL, Scotland, UK

Hastened forest fire: There is dry weather in April, followed by lightning in May and again in June. The lightning strike in May starts a forest fire. However, if the forest had not burned in May, it would have burned in June, because of the June lightning.

Delayed forest fire: There is heavy rain in April, followed by lightning in May and again in June. In May, the forest is too wet to burn. In June, however, the lightning strike starts a forest fire. If there had not been heavy rain in April, the situation would have been as in *Hastened forest fire*, and the forest would have burned already in May.²

In *Hastened forest fire*, the May lightning hastens the burning of the forest: if the May lightning had not occurred, the forest would have burned later—in June instead of May.³ Reflecting this, we judge that the May lightning is a cause of the forest fire. In *Delayed forest fire*, by contrast, the heavy rain in April delays the forest fire: with the heavy rain in April, the forest burns in June; if there had not been heavy rain in April, the forest would have burned already in May. Correspondingly, we judge that the April rain is *not* a cause of the forest fire; instead, the correct verdict here seems to be that the forest burns *in spite of* the heavy rain in April.

This asymmetry in our judgements about hasteners and delayers is puzzling because it seems that the metaphysical situation is symmetric (Schaffer 2005: 306): except for the fact that hasteners make the effect occur earlier, while delayers make the effect occur later, hasteners seem to be related to what they hasten in exactly the same way as delayers are related to what they delay. If the May lightning had not occurred, the forest would still have burned—it would simply have burned later. Similarly, if the April rain had not occurred, the forest would still have burned—it would simply have burned earlier. The only difference, it seems, is that hasteners make events happen *earlier*, while delayers make them happen *later*. And as Mackie asks (1992: 486): ‘Why should moving something to an earlier time count as causing it, while moving it to a later time does not?’

Based on this, it may seem that the best we can do is to give a pragmatic explanation of why we intuitively treat hasteners and delayers differently. Schaffer (2005: 306), for example, suggests the following pragmatic explanation:

Perhaps the hastener/delayer asymmetry can be explained away *pragmatically*. [...] Our intuitions are asymmetric, I propose, because: (i) we have an intuitive image of ourselves as agents intervening in the course of events, (ii) hasteners close opportunities for further intervention whereas delayers open them, and (iii) we are more concerned with closing opportunities because of the finality involved. *Finalizing* feels more salient than *postponing*.

² These cases were originally presented by Bennett (1987: 373).

³ Is the forest fire that would have occurred in June the same event as the forest fire that actually occurs in May? So far, I simply assume that it is. However, I provide arguments for this assumption in my discussion of fragility in Sect. 2.

This explanation appeals directly to the surface level characteristics of hasteners and delayers: that hasteners make events happen *earlier*, while delayers make events happen *later*. However, any explanation based on these surface level characteristics runs into trouble when confronted with atypical cases of hastening and delaying, where we intuitively judge that a hastener is *not* a cause of the event it hastens, or a delayer *is* a cause of the event it delays. Consider for example the two cases below:

The holiday: Suzy is planning to go to Greece in October. However, the teaching schedule for October is changed so that she is unable to go in October. She reschedules her trip and goes to Greece in September instead. (Cf. Mackie 1992: 485).

Fatal antidote: Billy is bitten by a snake and is promptly taken to hospital, where he is given the antidote. Twelve hours after the snakebite, he dies from a reaction to the antidote. If he had not been given the antidote, he would have died from the snake's poison already two hours after the snakebite. (Cf. Mackie 1992: 485).

In *The holiday*, the change in the October teaching schedule hastens Suzy's trip to Greece. Even so, it seems intuitively correct to say that the change in the October teaching schedule is *not* a cause of her trip—rather, we might say, she manages to go to Greece *in spite of* the change in the teaching schedule. In *Fatal antidote*, on the other hand, the antidote delays Billy's death. Even so, it seems intuitively correct to say that the antidote *is* a cause of his death.

A satisfactory explanation of the asymmetry in our judgements about hasteners and delayers must be able to account for these atypical cases as well. Thus, the asymmetry in our judgements cannot be explained simply as a result of the fact that hasteners make what they hasten happen *earlier*, while delayers make what they delay happen *later*.

In this paper, I am going to suggest a new explanation of the asymmetry in our judgements. In brief, my suggestion is that our causal judgements are sensitive to an underlying metaphysical distinction between three different ways in which an earlier event *C* may be related to a later event *E*. The distinction between these three relations may be drawn reductively, i.e. without appealing to causal notions. To see the distinction, we need to pay attention to the *modal* features of the later event *E*. More precisely, we need to pay attention to *E*'s *degree of security*. I give a precise characterisation of degree of security in Sect. 3. For now, however, the following rough characterisation is sufficient: an event has a low degree of security when it could easily have failed to occur, whereas it has a high degree of security when it could not easily have failed to occur. Based on how *C* influences *E*'s degree of security, we may then distinguish the following three relations:

- (1) *Positive security-dependence* is, roughly, the relation that holds between *C* and *E* when *E* would have been *less secure* if *C* had not occurred,
- (2) *Security-independence* is, roughly, the relation that holds between *C* and *E* when *E* would have been *just as secure* if *C* had not occurred, and

- (3) *Negative security-dependence* is, roughly, the relation that holds between *C* and *E* when *E* would have been *more secure* if *C* had not occurred.

As I will suggest (and explain in detail in the following sections), our causal judgements are sensitive to this metaphysical distinction: we judge that an event *C* is a cause of a later event *E* *only if* *E* security-depends positively on *C*. This, in turn, explains the asymmetry in our judgements about hasteners and delayers.

In typical cases, the surface-level distinction between hastening and delaying lines up with the above metaphysical distinction, so that a hastened event typically security-depends *positively* on the event that hastens it, while a delayed event typically does *not* security-depend positively on the event that delays it. In *Hastened forest fire*, for example, the forest fire security-depends positively on the May lightning: if the May lightning had not occurred, the forest fire would have been *less secure*. However, in exactly those atypical cases where we judge that a hastener is *not* a cause of the event it hastens, and that a delayer *is* a cause of the event it delays, we correspondingly find that the hastened event does *not* security-depend positively on the event that hastens it, while the delayed event *does* security-depend positively on the event that delays it. In *The holiday*, for example, we find that Suzy's trip to Greece security-depends negatively on the change in the October teaching schedule, while in *Fatal antidote*, we find that Billy's death security-depends positively on his being given the antidote. In this way, the deep metaphysical distinction between positive security-dependence, security-independence, and negative security-dependence can explain our judgements in both typical and atypical cases of hastening and delaying.

I will proceed as follows: in Sect. 2, I discuss the preliminary question of the fragility of events. In Sect. 3, I introduce the concept of *security*. In Sect. 4, I build on this to introduce the relations of positive security-dependence, security-independence, and negative security-dependence. In Sect. 5, I show that, in typical cases of hastening, the hastened event security-depends positively on the event that hastens it, whereas, in typical cases of delaying, the delayed event security-depends negatively on the event that delays it. In Sect. 6, I consider atypical cases of hastening and delaying. I begin by showing that, in the atypical cases presented above, the hastened event security-depends negatively on the event that hastens it, while the delayed event security-depends positively on the event that delays it. I then discuss a further kind of atypical case, where the hastened (or delayed) event is simply security-independent from the event that hastens (or delays) it. Finally, in Sect. 7, I discuss *why* the relation of positive security-dependence almost always holds between a hastened event and the event that hastens it, while the relation of negative security-dependence almost always holds between a delayed event and the event that delays it.

2 Hastening, delaying, and the question of fragility

In the following, I will assume that the causal *relata* are events.⁴ My aim in this section is to clarify my position on the question of the *fragility* of events.

⁴ This assumption is made in the vast majority of analyses of causation (see e.g. Paul and Hall 2013: 7).

The concept of fragility was originally introduced by Lewis (1986a: 196), who characterised it as follows:

Call an event *fragile* if, or to the extent that, it could not have occurred at a different time, or in a different manner. A fragile event has a rich essence; it has stringent conditions of occurrence.

Following Lewis, let us say that an event that does not have such stringent conditions of occurrence is modally *robust* (Lewis 1986a: 196).

Thus, the question of fragility is a question about which features of an event are essential to it, and which features are merely accidental. Consider, for example, the forest fire that in fact occurs in June (in *Delayed forest fire*). Could this forest fire—the very same event—have occurred earlier or later, or does it have its time of occurrence essentially? Could it have started from a different tree, and spread more slowly or rapidly, or is its manner of occurrence essential to it?

In the following, I will adopt Lewis's view on the question of fragility, since his view allows us to capture the nuances of our intuitive judgements about hasteners and delayers.⁵ On Lewis's view, there is, in any given region of space–time, a multitude of events that differ only in their fragility, i.e. that differ only in terms of what is essential to their occurrence, and what is merely accidental (see Lewis 1986a: 196–197; cf. Lewis 1986b). Consider, for example, the region of space–time where the June forest fire occurs. On Lewis's view, there is a multitude of events located in this region of space–time. Among these, there is a fairly fragile event—a forest fire that occurs essentially in June, essentially spreads from a particular old tree, etc. And there is a much more robust event—a forest fire that is essentially the burning of this particular forest, but which only accidentally occurs in June, only accidentally spreads from this particular old tree, etc.

This view nicely accommodates our intuitive judgements about hasteners and delayers: it allows us to capture the judgement that there is a possible reading of the claim that 'the April rain is a cause of the forest fire' on which it comes out true; and, at the same time, it allows us to capture the (much more natural) judgement that the heavy rain in April merely delays the forest fire, and does not cause it.

Let me begin by showing how Lewis's view accommodates the intuition that there is *some* sense in which the heavy rain in April is a cause of the June forest fire. Lewis's view accommodates this by allowing for fairly fragile events. In the present case, for example, Lewis's view allows for the existence of a fairly fragile forest fire, which has its manner and time of occurrence essentially: it is essential to it that it occurs in June, starts from that particular old tree, spreads in this particular way, etc. Call this fairly fragile event *fragile June fire*. The occurrence of *fragile June fire* depends counterfactually on the heavy rain in April: if there had not been heavy rain in April, *fragile June fire* would not have occurred at all. Instead, a numerically

⁵ My main reasons for adopting Lewis's proposal in the following is that it is simple, most readers may be expected to be familiar with it, and it can accommodate the data that I need to accommodate. It is important to note, however, that the data may also be accommodated by more sophisticated proposals about the causal relata, and what I say in the following may easily be applied within the framework of such proposals.

distinct forest fire would have occurred: a forest fire taking place already in May, originating (maybe) from a different tree, spreading in a different way, etc. If we take ourselves to be speaking of the *fragile June fire*, we should therefore say that the April rain is a cause of the forest fire. Indeed, this follows directly from the highly plausible principle that counterfactual dependence is sufficient for causation.

Usually, however, we are not interested in events with such stringent conditions of occurrence. Lewis (2004: 86) makes this point as follows:

we're usually quite happy to say that an event might have been slightly delayed, and that it might have differed somewhat in this or that of its contingent aspects. I recently postponed a seminar talk from October to December, doubtless making quite a lot of difference to the course of the discussion. But I postponed it instead of cancelling it because I wanted *that very event* to take place.

When we, for example, judge that the heavy rain in April *delays* the forest fire, we are not considering *fragile June fire*, which has its time and manner of occurrence essentially. Instead, to make sense of the idea that events can be hastened or delayed, we need to consider events that are much more robust—events whose precise time and manner of occurrence are not essential to them.

To capture the judgement that the April rain delays the forest fire, for example, we need to consider a fairly robust event: a forest fire that is essentially the burning of this particular forest, but whose precise time and manner of occurrence are not essential to it—it could have occurred in May or July instead of June, it could have started from a lightning strike hitting a different tree, it could have spread more slowly or rapidly, etc. Call this event *robust forest fire*. Considering this fairly robust event allows us to capture the judgement that the heavy rain in April delays the forest fire: if there had not been heavy rain in April, *robust forest fire*—that very event—would have occurred already in May; because of the heavy rain in April, it occurs later. Furthermore, it seems intuitively clear that the April rain is not a cause of *robust forest fire*.

In the following, I will focus on fairly robust events, such as *robust forest fire*. Firstly, as noted above, it is only appropriate to speak of hastening and delaying when we are considering such fairly robust events. And secondly and more importantly, the puzzling asymmetry between hasteners and delayers only comes into view when we consider such fairly robust events: when we consider fragile events that have their time of occurrence essentially, any prior event that influences their timing must be counted among their causes. It is only when we consider fairly robust events, which could have occurred earlier or later, that the asymmetry emerges. Here, the question arises: why do we judge, when we consider *robust forest fire*, that the hastening lightning strike in May is among its causes, while the delaying rain in April is not?

3 Security

I suggest that the key to understanding the asymmetry between hasteners and delayers is to look at a particular modal feature of events, namely *how easily they could have failed to occur*. In this section, I will introduce the concept of *security*, which provides a precise way of capturing this idea.

Before we can ask how easily an event could have failed to occur, we first need to specify which event we are talking about. Given the view presented in Sect. 2, there is, for each possible answer to the question of fragility, an event that fits exactly that answer. There is, for example, a fragile forest fire that has its time and manner of occurrence essentially; another more robust forest fire, which is essentially the burning of this particular forest, but could have happened at a somewhat different time and in a somewhat different manner, etc. To pick out a single event, it is therefore not enough to point to the space–time region in which the forest fire occurs. In addition, we need to specify whether we are talking about a fragile forest fire, or a more robust forest fire, etc. It is only once we have managed to specify a single event, corresponding to a particular answer to the question of fragility, that we can consider the further question of how easily this event could have failed to occur.

With this clarification in place, let me now illustrate the intuitive idea that some events could easily have failed to occur, while others could not. As an example of an event that could easily have failed to occur, consider a football player who takes a free kick in an important match, and succeeds in scoring a goal. More precisely, consider the fairly robust event of the goal, which is essentially the ball hitting the back of the net, but whose precise time and manner of occurrence are not essential to it. Let time t be the time at which the football player kicks the ball, and let time T be the time at which the ball hits the back of the net. And suppose that, had circumstances at time t been just ever so slightly different—if she had hit the ball slightly differently, or hesitated for a moment, or if the other players had moved just slightly differently—the ball would not have found the back of the net at all. In that case, we intuitively judge that—based on the circumstances at time t —the goal could easily have failed to happen. And, importantly, the reason why we reach this judgement is that, had things been just slightly different at time t , the goal would not have happened.

As an example of an event that could not easily have failed to happen, consider instead a football player who, during training, kicks the ball towards the middle of an undefended goal. More precisely, consider, as above, the fairly robust event of the goal. Once again, let t be the time at which the football player kicks the ball, and let T be the time at which the ball hits the back of the net. In this case, it seems reasonable to suppose that circumstances at time t would have to be *very* different in order for her not to score. Correspondingly, we intuitively judge—based on the circumstances at time t —that the goal could not easily have failed to happen. And, as before, our judgement is based on an assessment of the extent to which the goal depends on the precise circumstances at time t .

In the following, I will use possible worlds to make this idea more precise. Let me begin by introducing a particular notion of distance between worlds, namely the notion of *distance-at-a-time*. The distance-at-time- t between two worlds, w and w^* , is a measure of the dissimilarity between the state of w at t and the state of w^* at t .⁶ when the state of w at t is very similar to the state of w^* at t , then there is only a small distance-at- t between w and w^* , i.e., w and w^* are very close-at- t . When the state of w at t is very dissimilar from the state of w^* at t , then there is a large distance-at- t between w and w^* , i.e., w and w^* are very distant-at- t . I will say that a world w^* is *closest-at- t* to a world w just in case there is no world w^{**} such that the distance-at- t between w and w^{**} is strictly smaller than the distance-at- t between w and w^* (note that this allows for more than one world being closest-at- t to w).

In the following, I will only consider worlds governed by deterministic laws.⁷ And, in considering a case governed by a particular set of laws, I will only consider worlds governed by those same laws. Letting a not- E -world be a world that evolves (in accordance with the relevant set of laws) such that event E does not occur in it, we may now introduce the concept of *security* as follows:

Security: the security of an event E in a world w at a time t (where t may be any time prior to E 's time of occurrence, and w may be any world) is given by the distance-at- t between w and the not- E -world(s) that are closest-at- t to w .

If the not- E -world(s) that are closest-at- t to w are very distant-at- t , then E has a high degree of security at t in w . If the not- E -world(s) that are closest-at- t to w are very close-at- t , then E has a low degree of security at t in w .⁸

⁶ This seems to require us to pick out the same time—time t —in two different worlds, w_1 and w_2 . One might worry that this is not easy to do. However, there are ways around this difficulty. One option is the following: suppose that we are interested in a particular time—time t —in w_1 . Instead of supposing that there is a single correct way of picking out the same time in world w_2 , we allow that any time in w_2 may be picked out as time t . This yields a multitude of ways of picking out time t in world w_2 . For each particular way of picking out time t in w_2 , we may then determine the distance-at- t between w_1 and w_2 , relative to that way of picking out time t in w_2 . A world w_1 may then be close-at- t to w_2 relative to one way of picking out time t in w_2 , while it may not be close-at- t to w_2 relative to another way of picking out time t in w_2 . When I talk about ‘closeness-at- t ’ in the following, one may therefore take this as shorthand for ‘closeness-at- t , relative to a particular way of picking out time t ’.

⁷ There is, in fact, good reason to believe that the actual laws of nature are not deterministic. However, allowing for indeterministic laws would introduce further complications that I do not wish to consider here. For that reason, I follow the common practice of assuming determinism (see e.g. Paul and Hall 2013: 8).

⁸ With this precise definition of security, we are now in a position to see how the concept of security interacts with the concept of fragility: when we consider the occurrences in a particular space–time region, such as the space–time region in which the forest burns, these occurrences correspond, as detailed above, to a multitude of more or less fragile events. For a given earlier time t , we may now determine the degree of security at time t of each of these more or less fragile events. Doing so, we find that there is a systematic relationship between fragility and degree of security: roughly, increasing the fragility of an event (i.e. considering more details of its time and manner of occurrence to be essential to it) lowers its degree of security at the times before its occurrence. Considering the forest fire, for example, we find that even though things would have to be very different at time t in order for the robust forest fire not to occur, things would only have to be ever so slightly different in order for a particular fragile forest fire (occurring essentially in a particular manner at a particular time) not to occur.




Fig. 1 The graph illustrated here shows a typical evolution of an event's security over time: at times long before E 's occurrence, E has a very low degree of security, i.e. if things had been just slightly different at these times, E would not have occurred. Closer to E 's occurrence, however, E 's degree of security begins to grow, and just before E 's occurrence, E has a quite high degree of security, i.e. things would have to be very different at this time in order for E not to occur

Since we can evaluate an event E 's security at any time prior to E 's occurrence, we may draw a graph of how E 's security evolves over time. Such a graph may, for example, look like the one illustrated in Fig. 1. (This graph only contains comparative information about how an event's degree of security at one time relates to its degree of security at earlier or later times. For this reason, no unit is required on the y-axis. Indeed, it is an open question whether it is possible to find an absolute measure of an event's degree of security, and what the appropriate unit may be. For now, however, comparisons of degrees of security are all we need.)

The concept of security makes precise our intuitive judgements about how easily an event could have failed to occur, based on the circumstances at some earlier time: saying that an event E has a low degree of security at an earlier time t captures the idea that E could easily have failed to happen, based on the circumstances at t . And saying that an event E has a high degree of security at an earlier time t captures the idea that E could not easily have failed to happen, based on the circumstances at t .

As an illustration, let us once again consider the two football examples presented above. Consider first the case of the football player who takes a free kick in an important match. To assess the goal's degree of security at the time t when she takes the kick, we need to ask: how close-at- t are the closest-at- t possible worlds where she does not score a goal? Based on my description of the case, we find that the closest-at- t possible worlds where she does not score are *very* close-at- t : the state-at- t of the closest-at- t no-goal-worlds differs only ever so slightly from the state-at- t of the actual world. Thus, we find that the goal has a very low degree of security at t , corresponding to the intuitive judgement that the goal could easily have failed to happen, based on the circumstances at t .

By contrast, consider now the case of a football player who kicks the ball towards the middle of an undefended goal during training. In this case, we find, based on my description of the case, that the closest-at- t worlds where she does not score are very distant-at- t : the closest-at- t worlds where she does not score are worlds where she kicks the ball at a very different angle, or hesitates for a long time, or where the other players are moving *very* differently. Thus, we find that the goal has a very high degree of security at t , corresponding to the intuitive judgement that the goal could not easily have failed to happen, based on the circumstances at time t .

4 Security-dependence

With the concept of security, we may now go on to consider how an event C influences the security of a later event E , where we shall understand influence in counterfactual terms. In particular, considering the time t at which C occurs, we may ask: would E have been *less secure* at t , *just as secure* at t , or *more secure* at t , if C had not occurred? For any event C and any later event E , C will be related to E in exactly one of these three ways. I will call these three relations *positive security-dependence*, *security-independence* and *negative security-dependence*. They may be defined more precisely as follows:

Positive security-dependence: E security-dependes positively on C , where C occurs at time t ,⁹ if and only if

1. C and E both occur, and
2. if C had not occurred, then E would have been *less secure* at t .

Security-independence: E is security-independent from C , where C occurs at time t , if and only if

1. C and E both occur, and
2. if C had not occurred, then E would have been *just as secure* at t .

Negative security-dependence: E security-dependes negatively on C , where C occurs at time t , if and only if

1. C and E both occur, and
2. if C had not occurred, then E would have been *more secure* at t .

Considering an event C , occurring at time t , and a later event E , Fig. 2 illustrates a case where E security-dependes positively on C .

It is important to note that relations of security-dependence are doubly modal: the notion of security is itself a modal notion—to evaluate an event’s security in a world w at a time t , we need to consider the distance-at- t between w and other possible worlds. And to evaluate which of the three relations holds between an event C , which occurs at t , and a later event E , we then need to compare E ’s *actual* degree of security with the degree of security E would have had at t , if C had not occurred.

Thus, determining which of the three relations holds between an event C and a later event E requires us to evaluate counterfactuals of the form ‘if C had not occurred, then [...]’. My preferred method for making these evaluations was originally proposed by Maudlin (2007: 21–34), and Hall and Paul later tailored it specifically to the needs of counterfactual theories of causation (see Hall 2007b: 17–21; Paul and Hall 2013: 47–53). Paul and Hall (2013: 47–48) describe this method as follows:

Suppose event C occurs at t , and event E occurs later. To evaluate “if C had not occurred, then [...],” we construct a counterfactual state of the world at

⁹ In cases where C goes on for some time, the relevant time t is the time when C begins to occur. The same applies to the following definitions.




Fig. 2 The *solid line*—as in the previous graph—indicates how *E*'s degree of security evolves over time in the actual world. The *dotted line*, on the other hand, indicates how *E*'s degree of security evolves from time *t* onwards in the closest possible world where *C* does not occur. (The figure does not show how *E*'s degree of security evolves until time *t* in the closest not-*C*-world; the reason for this is that my preferred way of evaluating counterfactuals simply remains silent on what happens in the closest not-*C*-world until the relevant time *t*.) Since the *dotted line* lies below the *solid line* at time *t*, the diagram shows that *E* is less secure at *t* in the closest not-*C*-world than it is in the actual world, and thus *E* security-depends positively on *C*. Note that since *E* still occurs in the closest not-*C*-world, *E* does not depend counterfactually on *C* in this example

time *t* as much like the actual state at time *t* as possible, save for the fact that *C* does not occur. [...] We then evolve the resulting state forward in time, in accordance with the actual laws of nature.¹⁰

Using this method for evaluating counterfactuals, we may determine, whenever we have an event *C* and a later event *E*, whether *E* security-depends positively on *C*, is security-independent from *C*, or security-depends negatively on *C*.

The distinction between these three different relations is metaphysically significant. And I suggest that our causal judgements are sensitive to exactly this distinction:

- I. We judge that *C* is a cause of a later event *E*
only if *E* security-depends positively on *C*.
- II. We judge that *C* is (causally) *irrelevant* to a later event *E*
if and only if *E* is security-independent from *C*.
- III. We judge that *E* occurs *in spite of C*
if and only if *E* security-depends negatively on *C*.¹¹

Hopefully, these claims already seem intuitively plausible.¹² In the following, I will now show how accepting them allows us to accurately capture our intuitive judgements in both typical and atypical cases of hastening and delaying.

¹⁰ As Paul and Hall (2013: 49–53) point out, there is a further question of what exactly replaces *C* in the closest-at-*t* world(s) where *C* does not occur. For discussion, see, e.g., Schaffer (2005: 302–306), Hall (2007a: 125) and Hitchcock (2007: 506–507).

¹¹ Note that positive security-dependence is necessary, but not sufficient, for our judgement that *C* is a cause of *E*. By contrast, security-independence is both necessary and sufficient for the judgement that *C* is causally irrelevant to *E*, and negative security-dependence is both necessary and sufficient for the judgement that *E* occurs *in spite of C*.

¹² I provide further support for these claims in my paper 'Security-dependence and the intransitivity of causation' (ms).

5 Security-dependence in typical cases of hastening and delaying

In this section, I will show how, in typical cases of hastening, the hastened event security-depends *positively* on the event that hastens it, whereas, in typical cases of delaying, the delayed event security-depends *negatively* on the event that delays it. This difference, I suggest, explains the asymmetry in our judgements.

To show this, I will consider the representative examples presented by *Hastened forest fire* and *Delayed forest fire*. Consider first *Hastened forest fire*, which provides an example of a typical case of hastening, where we judge that the hastener is a cause of the event it hastens. My claim is that the forest fire—considered as a fairly robust event, in keeping with what I have said in Sect. 2 above—security-depends *positively* on the May lightning. To see this, we will need to consider the three possible worlds shown in Fig. 3.

More precisely, we need to consider the state of each of these three worlds at the time t when the May lightning occurs. To see that the forest fire security-depends positively on the May lightning, we may now proceed in the following three steps:

Step 1: assessing the forest fire's degree of security at t in @: relative to @, the closest-at- t world where the forest fire does not occur is w^* . Thus, the forest fire's degree of security at t in @ is given by the distance-at- t between @ and w^* .

Step 2: assessing the degree of security that the forest fire would have had at t , if the May lightning had not occurred: relative to @, the closest-at- t world where the May lightning does not occur is w . (This is confirmed by the intuitive judgement that, if the May lightning had not occurred, the forest would have burned in June: this judgement shows that we take the closest-at- t world where the May lightning does not occur to be a world in which the June lightning still occurs.) Relative to w , the closest-at- t world where the forest fire does not occur is w^* . Thus, the degree of security that forest fire would have had at t , if the May lightning had not occurred, is given by the distance-at- t between w and w^* .

Step 3: comparing the forest fire's degree of security at t in @ with the degree of security it would have had at t , if the May lightning had not occurred: intuitively, the distance-at- t between w and w^* is *smaller* than the distance-at- t between @ and w^* . Thus, the forest fire would have been *less secure* at t , if the May lightning had not occurred, i.e. *the forest fire security-depends positively on the May lightning*.




Fig. 3 The three worlds that are relevant for evaluating security-dependence in *Hastened forest fire*; I have used italics to emphasise differences from the actual world @. Note that I have put the actual world, @, on the far left, since this order of the worlds best reflects the comparative similarity between them




Fig. 4 The three worlds that are relevant for evaluating security-dependence in *Delayed forest fire*; I have used italics to emphasise differences from the actual world @. Note that @ is here placed in the middle; this is intended to reflect the fact that w differs from @ in one way, while w^* differs from @ in another way

Now consider *Delayed forest fire*, as an example of a typical case of delaying, where we judge that the delayer is *not* a cause of the event it delays. My claim is that the forest fire security-depends *negatively* on the heavy rain in April, which delays its occurrence. To see this, we will need to consider the three possible worlds shown in Fig. 4.

More precisely, we need to consider the state of each of these three worlds at the time t when the April rain begins. To see that the forest fire security-depends negatively on the April rain, we now proceed in the same three steps as above:

Step 1: assessing the forest fire's degree of security at t in @: relative to @, the closest-at- t world where the forest fire does not occur is w^* . Thus, the forest fire's degree of security at t in @ is given by the distance-at- t between @ and w^* .

Step 2: assessing the degree of security that the forest fire would have had at t , if the April rain had not occurred: relative to @, the closest-at- t world where the April rain does not occur is w . (This is confirmed by the intuitive judgement that if the April rain had not occurred, the forest would have burned already in May.) Relative to w , the closest-at- t world where the forest fire does not occur is w^* . Thus, the degree of security that the forest fire would have had at t , if the April rain had not occurred, is given by the distance-at- t between w and w^* .

Step 3: comparing the forest fire's degree of security at t in @ with the degree of security it would have had at t , if the April rain had not occurred: intuitively, the distance-at- t between w and w^* is larger than the distance-at- t between @ and w^* . Thus, the forest fire would have been *more secure* at t , if the April rain had not occurred, i.e. *the forest fire security-depends negatively on the April rain*.

These results show that there is a metaphysically significant difference between the way in which the May lightning is related to the forest fire in *Hastened forest fire*, and the way in which the April rain is related to the forest fire in *Delayed forest fire*: the forest fire in *Hastened forest fire* security-depends *positively* on the May lightning, while the forest fire in *Delayed forest fire* security-depends *negatively* on the April rain. This difference, I suggest, explains the asymmetry in our causal judgements:

When an event E security-depends *positively* on an earlier event C , we judge (provided other reasonable conditions for causation are satisfied) that C is a cause of E . Thus, the fact that the forest fire in *Hastened forest fire* security-depends

positively on the May lightning explains (provided reasonable conditions are satisfied) our judgement that the May lightning is among its causes. When, on the other hand, an event E security-depends *negatively* on an earlier event C , we judge that C is *not* a cause of E —rather, E occurs *in spite of* C . Thus, the fact that the forest fire in *Delayed forest fire* security-depends *negatively* on the April rain explains our judgement that the April rain is *not* a cause of the forest fire—and that, rather, the forest burns *in spite of* the April rain.

6 Atypical cases of hastening and delaying

In this section, I will consider atypical cases of hastening, where we judge that a hastener is *not* a cause of what it hastens, and atypical cases of delaying, where we judge that a delayer *is* a cause of what it delays. In these atypical cases, I will show that the hastened event does *not* security-depend positively on the event that hastens it, whereas the delayed event *does* security-depend positively on the event that delays it. This explains why our judgements about these atypical cases are reversed relative to our judgements about typical cases of hastening and delaying.

To show this, I will begin by considering the representative examples presented by *The holiday* and *Fatal antidote*. Consider first *The holiday*, presented in Sect. 1, which provides an example of an atypical case of hastening. In *The holiday*, Suzy is planning to go to Greece in October. However, a change in the October teaching schedule makes it impossible for her to go then. Although she would have preferred to go in October, she therefore changes her plans and goes to Greece in September instead. Thus, the change in the October teaching schedule hastens her trip. However, we intuitively judge that the change is not a cause of her trip—rather, she manages to go *in spite of* the change. Correspondingly, Suzy's trip to Greece in fact security-depends negatively on the change in the October teaching schedule. To see this, we will need to consider the three possible worlds shown in Fig. 5.

More precisely, we need to consider the state of each of these three possible worlds at the time t when the October teaching schedule changes. We may now proceed in the usual three steps:



Fig. 5 The three worlds that are relevant for evaluating security-dependence in *The holiday*; I have used italics to emphasise differences from the actual world @. Note that I have put the actual world, @, in the middle, since this best reflects the fact that w differs from @ in one way, while w^* differs from @ in a different way

Step 1: assessing the trip's degree of security at t in @: relative to @, the closest-at- t world where Suzy does not go to Greece is w^* . Thus, the trip's security at t in @ is given by the distance-at- t between @ and w^* .

Step 2: assessing the degree of security that the trip would have had at t , if the October teaching schedule had not changed: relative to @, the closest-at- t world where the teaching schedule does not change is w . (This is confirmed by the intuitive judgement that, if the teaching schedule had not changed, Suzy would have gone to Greece in October.) Relative to w , the closest-at- t world where Suzy does not go to Greece is w^* . Thus, the degree of security that the trip would have had at t , if the October teaching schedule had not changed, is given by the distance-at- t between w and w^* .

Step 3: comparing the trip's degree of security at t in @ with the degree of security it would have had at t , if the October teaching schedule had not changed: intuitively, the distance-at- t between w and w^* is larger than the distance-at- t between @ and w^* . Thus, the trip would have been *more secure* at t , if the teaching schedule had not changed, i.e. *the trip security-depends negatively on the change in the October teaching schedule*.

This result explains our judgement that the change in the October teaching schedule is *not* a cause of Suzy's trip. Furthermore, it fits perfectly with the nuances of our judgement: as I have suggested above, when an event E security-depends negatively on an earlier event C , we judge that E occurs *in spite of* C . This suggestion fits exactly with our intuitive judgement in the present case, where it seems appropriate to say that Suzy manages to go on her trip to Greece *in spite of* the change in the October teaching schedule.

Now consider the atypical case of delaying presented in *Fatal antidote*. In *Fatal antidote*, presented in Sect. 1, Billy is bitten by a snake and is immediately given an antidote to the snake's poison. He dies twelve hours later from a reaction to the antidote; if he had not been given the antidote, he would instead have died just two hours after the snake-bite. Here, we intuitively judge that the delayer—namely, Billy's being given the antidote—is a cause of the delayed event, i.e. Billy's death. Correspondingly, Billy's death security-depends *positively* on his being given the antidote. To see this, we will need to consider the three possible worlds shown in Fig. 6.

More precisely, we need to consider the state of each of these three worlds at the time t when Billy is given the antidote. We may again proceed in the usual three steps.




Fig. 6 The three worlds that are relevant for evaluating security-dependence in *Fatal antidote*; I have used italics to emphasise differences from the actual world @. Note that I have put the actual world, @, on the far left, since this order of the worlds best reflects the comparative similarity between them

Step 1: assessing the degree of security of Billy's death at t in @: relative to @, the closest-at- t world where Billy doesn't die is w^* . Thus, his death's degree of security at t in @ is given by the distance-at- t between @ and w^* .

Step 2: assessing the degree of security that Billy's death would have had at t , if he had not been given the antidote: relative to @, the closest-at- t world where Billy is not given the antidote is w . (This is confirmed by the intuitive judgement that Billy would have died already two hours after the snake-bite, if he had not been given the antidote.) Relative to w , the closest-at- t world where Billy doesn't die is w^* . Thus, the degree of security that Billy's death would have had, if he had not been given the antidote, is given by the distance-at- t between w and w^* .

Step 3: comparing the degree of security of Billy's death at t in @ with the degree of security that his death would have had at t , if he had not been given the antidote: intuitively, the distance-at- t between w and w^* is *smaller* than the distance-at- t between @ and w^* . Thus, Billy's death would have been *less secure* at t , if he had not been given the antidote, i.e. *Billy's death security-depends positively on his being given the antidote*.

This result—together with the satisfaction of reasonable conditions—allows us to judge that being given the antidote is a cause of Billy's death.

The holiday and *Fatal antidote* are examples of atypical cases where a hastened event security-depends *negatively* on the event that hastens it, while a delayed event security-depends *positively* on the event that delays it. However, as one might expect from the fact that we have distinguished three possible relationships—positive security-dependence, negative security-dependence, *and* security-independence—there is a further kind of atypical cases: cases where the event that is hastened (or delayed) is simply security-independent from the hastener (or delayer). The following is an example of this type of case:

Train tracks: A train is approaching a fork in the tracks. Suzy turns the lever in the signal box from left to right, so that the train travels down the right-hand track and arrives at its destination at 12 noon. If Suzy had not turned the lever, the train would instead have travelled down the slightly longer left-hand track, arriving five minutes past noon.¹³

Considering the fairly robust event of the train's arrival at some time around noon, we find that this event would have happened whether or not Suzy had turned the lever. Suzy's turning the lever simply hastens the train's arrival. As Mackie (1992: 496) notes,

¹³ Cf. Mackie (1992: 496). This case has been much discussed in the literature (see, for example, Hall 2000: 205; Hall 2007a: 118; Paul and Hall 2013: 232; Sartorio 2005: 74–75). The reason why this case has received so much attention is that it presents a counterexample to the transitivity of causation: intuitively, Suzy's turning the lever causes the train's journey along the right-hand track, and the train's journey along the right-hand track in turn causes its arrival, but our intuitive judgement is that Suzy's turning the lever is *not* a cause of the train's arrival. This suggests that the distinction between positive security-dependence, security-independence, and negative security-dependence may, in addition to explaining our judgements about hasteners and delayers, provide an explanation of what happens in the counterexamples to the transitivity of causation. I develop this idea in my paper 'Security-dependence and the intransitivity of causation' (ms).

one feels uncomfortable about saying that this [...] hastener—turning the lever—causes the train's arrival. And the same is true of the [...] delayer we get if we modify the example so that [Suzy's] action is to turn the lever from right to left.

Intuitively, it seems right to say that Suzy's turning the lever is simply *irrelevant* to the train's arrival. Correspondingly, we find that the train's arrival is security-independent from Suzy's turning the lever: considering the time t when Suzy turns the lever, and assuming that the two tracks are relevantly similar (apart from the fact that one is slightly longer than the other), we find that the train's arrival would have been *just as secure* at t , if Suzy had not turned the lever. This, I suggest, is the reason why we judge that Suzy's turning the lever is not a cause of the train's arrival, and that, instead, it is simply irrelevant to the train's arrival.

7 Explaining the asymmetry between typical cases of hastening and delaying

So far, we have seen how the asymmetry in our judgements about hasteners and delayers results from the fact that our causal judgements are sensitive to the deep metaphysical distinction between positive security-dependence, security-independence, and negative security-dependence. However, a question still remains: *why* does the surface level distinction between hastening and delaying so often line up with this deep metaphysical distinction? In other words: *why* do hastened events so often security-depend positively on the events that hasten them, and *why* do delayed events so often security-depend negatively on the events that delay them?

The key to answering this question, I suggest, is to see how hastening and delaying is connected with preemption. For hasteners, I suggest that:

- (i) Hasteners are typically preempting causes of what they hasten.¹⁴
- (ii) Effects always security-depend positively on their preempting causes.¹⁵

Similarly, for delayers I suggest that:

- (i) Delayers typically block a would-be preempting cause of what they delay.
- (ii) Effects always security-depend negatively on events that block their would-be preempting causes.

I cannot offer complete arguments for these two sets of claims. Rather, I will merely point out that they both fit perfectly with typical cases of hastening and delaying. Consider, for example, *Hastened forest fire*: the structure of *Hastened*

¹⁴ See e.g. Yablo (2004: 131). As has often been noted, the converse also holds: preempting causes are often (though not always) hasteners (see e.g. Yablo 2004: 132; Paul 1998).

¹⁵ I argue more fully for this claim in my paper 'Security-dependence and the intransitivity of causation' (ms).

forest fire is a typical preemption structure, where the May lightning is a preempting cause of the forest fire, while the June lightning is a preempted backup. And in Sect. 5 we have already seen that—in accordance with my claim that effects always security-depend positively on their preempting causes—the forest fire security-dependes positively on its preempting cause, i.e. on the May lightning.

Now consider *Delayed forest fire*. The crucial thing to notice here is that *Hastened forest fire* and *Delayed forest fire* are closely related: *Hastened forest fire* shows what would have happened in *Delayed forest fire* if the April rain had not occurred. Keeping this in mind, we see that *Delayed forest fire's* rain in April interferes with the preemption structure exhibited in *Hastened forest fire*: the April rain *blocks* a would-be preempting cause of the forest fire (namely, the May lightning), leaving the forest fire to be caused by what would otherwise have been a mere backup (namely, the June lightning). And in Sect. 5 we have already seen that the forest fire security-dependes negatively on the April rain that blocks the would-be preempting cause.

Both of the suggested explanations lead to further questions: why are hasteners often preempting causes? And why do delayers often block would-be preempting causes of what they delay? My tentative answer is that introducing a preempting cause is an obvious way to hasten an event, while blocking a would-be preempting cause is an obvious way to delay it. Consider, for example, Billy who is planning (as usual) to break a window. If Billy is taking his time, enjoying the sunshine before throwing his rock at the window, Suzy can hasten the breaking of the window by introducing a preempting cause—viz., by throwing a rock of her own. And if Sally stumbles on this situation and wants to delay the breaking of the window, she can do so by blocking the would-be preempting cause—viz., by running up to Suzy and preventing her from throwing, leaving the breaking of the window to the slower Billy.

In atypical cases, however, it is possible to hasten something by blocking a would-be preempting cause, or to delay something by introducing a preempting cause. This, indeed, is exactly what happens in our two atypical cases of hastening and delaying: *The holiday* and *Fatal antidote*.

When we look more closely at *The holiday*, we find that the change in the teaching schedule blocks a would-be preempting cause: in the counterfactual situation where there is no change in the teaching schedule, Suzy has the option of going to Greece in September *and* the option of going to Greece in October, and she prefers to go in October. In this counterfactual situation, then, Suzy's being free in October is a preempting cause of her going to Greece, while her being free in September is a mere preempted backup. The change in the teaching schedule thus hastens her trip by blocking a would-be preempting cause—namely, Suzy's being free to go on a later trip in October—so that the trip to Greece has to fall back on Suzy's being free in September.

When we look more closely at the structure of *Fatal antidote*, on the other hand, we find that it has an early preemption structure: the antidote cuts off the process initiated by the snake's poison; however, the antidote subsequently brings about Billy's death on its own, though more slowly. In this atypical case, we thus have a preempting cause that delays the effect of which it is a cause.

8 Conclusion

In typical cases, we judge that hasteners are causes of what they hasten, while delayers are not causes of what they delay. This asymmetry presents a puzzle: why do we treat hasteners and delayers differently in this way? Mackie (1992: 487) sums up the question nicely as follows:

the central question, applied to the example of the forest fire, can be stated as follows. There is a relation that holds between the rain and the occurrence of the fire that it delays. And there is a relation that holds between the lightning and the occurrence of the fire that it hastens. What is the difference between these relations that makes it natural to regard the lightning, but not the rain, as a cause of the fire?

My suggested answer to this question is that the fire security-depends *positively* on the lightning, whereas it security-depends *negatively* on the rain. More fully, my suggestion is that there is a deep metaphysical distinction between three ways in which an event *E* may be related to an earlier event *C*: *E* may security-depend positively on *C*, *E* may be security-independent from *C*, or *E* may security-depend negatively on *C*. This deep metaphysical distinction typically (but not always) lines up with the surface level distinction between hastening and delaying, and our sensitivity to this deep distinction explains our judgements in both typical and atypical cases of hastening and delaying.

Acknowledgements I would like to thank Katherine Hawley, Derek Ball, Sara Bernstein, Gabriel Uzquiano Cruz, Ned Hall, Matthew McGrath, L. A. Paul, and Christian Loew for fruitful discussions at various stages during the completion of this paper. In addition, I would like to thank colleagues and visiting students at Arché Philosophical Research Centre, University of St Andrews.

Compliance with ethical standards

Conflict of interest This research has been carried out as part of the author's Ph.D., which is funded by the Philosophical Quarterly and the Janet T Anderson Trust.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bennett, J. (1987). Event causation: The counterfactual analysis. In J. E. Tomberlin (Ed.), *Philosophical perspectives, 1: Metaphysics, 1987* (pp. 367–386). Atascadero, CA: Ridgeview Publishing Company.
- Hall, N. (2000). Causation and the price of transitivity. *Journal of Philosophy*, 97(4), 189–222.
- Hall, N. (2007a). Structural equations and causation. *Philosophical Studies*, 132(1), 109–136.
- Hall, N. (2007b). *Structural equations and causation*, extended version. <http://nrs.harvard.edu/urn-3:HUL.InstRepos:3710361>.

- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review*, 116(4), 495–532.
- Lewis, D. (Ed.). (1986a). Causation. In *Philosophical papers* (Vol. 2, pp. 159–213). Oxford: Oxford University Press.
- Lewis, D. (Ed.). (1986b). Events. In *Philosophical papers* (Vol. 2, pp. 241–269). Oxford: Oxford University Press.
- Lewis, D. (2004). Causation as influence. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 75–106). Cambridge, MA: MIT Press.
- Mackie, P. (1992). Causing, delaying, and hastening: Do rains cause fires? *Mind*, 101(403), 483–500.
- Maudlin, T. (2007). A modest proposal concerning laws, counterfactuals, and explanations. *The metaphysics within physics* (pp. 5–49). Oxford: Oxford University Press.
- Paul, L. A. (1998). Keeping track of the time: Emending the counterfactual analysis of causation. *Analysis*, 58(3), 191–198.
- Paul, L. A., & Hall, N. (2013). *Causation: A user's guide*. Oxford: Oxford University Press.
- Sartorio, C. (2005). Causes as difference-makers. *Philosophical Studies*, 123(1/2), 71–96.
- Sartorio, C. (2006). On causing something to happen in a certain way without causing it to happen. *Philosophical Studies*, 129(1), 119–136.
- Schaffer, J. (2005). Contrastive causation. *The Philosophical Review*, 114(3), 297–328.
- Yablo, S. (2004). Advertisement for a sketch of an outline of a prototheory of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 119–137). Cambridge, MA: MIT Press.