

## *Coincidence and Common Cause*

TAMAR LANDO  
Columbia University

### **Abstract**

According to the traditional view of the causal structure of a coincidence, the several parts of a coincidence are produced by independent causes. I argue that the traditional view is mistaken; even the several parts of a coincidence may have a common cause. This has important implications for how we think about the relationship between causation and causal explanation—and in particular, for why coincidences cannot be explained.

### **1. Introduction**

It is a truism that coincidences cannot be explained. Assuming that the sort of explanation at stake here is causal explanation, the truism informs us that coincidences do not admit of causal explanation. But it should be equally obvious that there are many things we *can* say about the causal antecedents of a coincidence. Suppose, to take an example from David Owens, that I coincide with my long lost friend at the train station. He is there to leave on a trip to Africa, and I am there to meet my mother. The cause of one part of the meeting lies in my intentions with respect to my mother; the cause of the other part lies in my friend's intentions with respect to his trip. Given the wealth of causal information we have about the meeting, in what sense can it not be causally explained?

There is a traditional view about the causal structure of coincidences, which is nicely articulated by Owens when he says that “a coincidence is an event which can be divided into components separately produced by independent causal factors.”<sup>1</sup> According to the traditionalist, the several parts of a coincidence do not have a common cause. In our example, the causal factors responsible for my being at the station are independent of the causal factors responsible for my friend's being at the station—I would go to the station even if he didn't intend to leave for Africa, and he would go to the station even if I didn't intend to meet my mother. This view lends itself naturally to an answer to our question about why coincidences cannot be explained. Very loosely, the causal factors that explain each part of the coincidence do not come together in such a way as to provide an explanation for the whole. As Owens puts it, “the factors invoked to explain the outcome do not *mesh with one another* in such a way as to provide an explanation for it—they do not unite to produce it.”<sup>2</sup>

The central claim I wish to defend in this paper is that the traditional view of the causal structure of a coincidence is mistaken: even the several parts of a coincidence may have a common cause. Or, to put things slightly differently, one and the same thing can bring about the separate parts of a coincidence. If this is right, it leaves it mysterious why coincidences in general cannot be explained. In particular, in those cases where there is a common cause, why can't we explain the coincidence by citing that very cause? Once we reject the traditional view, we must think anew about whether the truism we began with is in fact *true*—and if so, in what sense.

Before launching into the argument, let me flag two points. First, both I and the traditionalist are concerned with a notion of coincidence that applies to events located in space and time, or facts that involve such events. Coincidences are things that *happen*. (I will not be discussing mathematical coincidences, even if there may be some.)<sup>3</sup> Second, there is clearly more to a coincidence than anything one might say about causal structure or causal explanation. Hart and Honore, for example, maintain that a coincidence must be very unlikely by ordinary standards and for some reason significant or important.<sup>4</sup> The second of these conditions seems to me to look in the wrong place. If the only two people in a supermarket aisle are wearing the same pair of 1980's red Nike Air Ship high-tops, this may be a coincidence, but is not, presumably, important in any way. We might say instead that a coincidence involves a "match" of some appropriate kind between the constituent parts of the given event. In Owens' example, my friend and I are at *the very same* station, at *the very same* time. (Compare this with, for example, the combination of my hiking the Adirondack mountains right now and your reading this paper.) Saying just what a "match" of this sort amounts to is not an easy task, but whatever the exact formulation, the condition we want is broadly psychological in character—the joint occurrence of the several parts of a coincidence must be cognitively salient to us in some way. Although a condition of this general sort is integral to the familiar concept of coincidence, I wish to put all questions about its precise nature to one side. Happily, as I view it, the metaphysics of coincidence (questions about causation and explanation) can to a large extent be separated from the "psychology" of coincidence, and so in discussing the former, we can remain reasonably agnostic as to how things will eventually shake out with respect to the latter.

Thinking about coincidence is important for a number of reasons. If there are facts which *in principle* cannot be explained, what follows from this? Do the events which those facts involve have a particular kind of causal structure, as the traditional view posits—one that frustrates any attempt to causally explain? Or do those events lack causes altogether, as Owens argues? If not, what is the right way to distinguish between causation and causal explanation so that these conclusions do not follow?

## 2. The Traditional View

Let's begin by getting clearer on the traditional view. According to the traditionalist, the several components of a coincidence are brought about by independent causal factors; they do not share a common cause.


Support for this view can be found in, *e.g.*, Owens (1990), Owens (1992), Monod (1970), Horwich (1982), Lange (2010), and Sober (1984). Lange talks about the coincidence that the CIA agent was in the capital just when His Excellency dropped dead. He says, “Of course, this pair of events is coincidental because the two events have no common cause—at least, none that is important and of a relevant kind” (Lange (2010), p. 317). Sober discusses a situation in which several people are viewing the same movie, and have a visual impression of a boat being dragged up a mountain. “There is a correlation between my impression and yours. This is no mere coincidence, we say. There is a common cause—the image on the screen” (Sober (1984), p. 217). Horwich says that a coincidence is “an unlikely accidental correspondence between independent facts, which suggests strongly, *but in fact falsely, some causal relationship* between them” (Horwich (1982), p. 104). Monod says, “Mais dans d’autres situations, la notion de hasard prend une signification essentielle et non plus simplement opérationnelle. C’est le cas, par exemple, de ce que l’on peut appeler les ‘coïncidences absolues,’ c’est-à-dire celles qui résultent de l’intersection de deux chaînes causales totalement indépendantes l’une de l’autre” (Monod (1970), p. 128).<sup>5</sup>

Components of a coincidence must not, according to the traditionalist, have a common cause. A moment’s thought, however, reveals that events have long and complex causal histories. As Lewis once said, “We might imagine a world where causal histories are short and simple; but in the world as we know it, the only question is whether they are infinite or merely enormous.”<sup>6</sup> If we reach far enough back, events seemingly unrelated may well share *some* causal ancestor. Are we to understand the traditionalist as maintaining that the several components of a coincidence have no common cause whatsoever? This would be uncharitable; it would count the traditionalist as denying that many of the events we ordinarily think of as coincidences are coincidences. (If all events have the Big Bang as a common cause, that ought not show that on the traditional view there are no coincidences.)<sup>7</sup>

The traditionalist can be taken to mean that the several parts of a coincidence do not have a *salient* common cause, or that the salient causes of component events are independent. Here we distinguish salient causes—sometimes referred to as *the* cause of an event—from other “causal factors” or mere background conditions. To take an example from Hart and Honore, when a fire breaks out, we typically think of the presence of oxygen in the room as a background condition and the lighting of the match, or the short-circuit, as a salient cause. We do this even though the fire would not have broken out had there been no oxygen in the room. On a charitable reading, the traditionalist denies only that the several components of a coincidence have a common salient cause.

### 3. Counterexample to the Traditional View

But even understood in this way, there are, I think, straightforward counterexamples to the traditional view. Consider the following scenario.



**Figure 1.** (a) Nodes 3 and 4, the constituent events of the coincidence, proceed from independent causes; (b) The causal structure of *Pianos*. Nodes 1, 2, and 3 represent (respectively) the throwing of the ball, the striking of the first *A*-note, and the striking of the second *A*-note; the causal relations in the case are given by taking the transitive closure of the relation represented by directed edges.

*Pianos.* A boy is playing with a ball in the courtyard of an apartment complex. He throws the ball too high, and it bounces off of the balcony of one apartment, sails through the air, bounces onto the balcony of another apartment, and finally falls to the ground. On each of the two balconies sits a grand piano. As the ball lands on the first balcony it strikes a note on the first piano, and as the ball lands on the second balcony, it strikes a note on the second piano. On each of the two pianos, the note struck is the high *A*.

It is a coincidence that the same note is struck on both pianos. Nevertheless, the throw of the ball causes the first *A*-note to be struck, and the throw of the ball causes the second *A*-note to be struck. Thus the throw is a common cause of each of the coincidence's constituent events. If this is right, then the traditional view is mistaken: the constituent events of a coincidence need not be produced by independent causes. (Here and in the remainder of the paper I understand "common cause" as follows: *c* is a common cause of *e* and *f* if *c* is both a cause of *e* and a cause of *f*.<sup>8</sup>)

Let me briefly make explicit pictorially what I take to be the upshot of the example. The traditionalist claims that the several parts of a coincidence do not have a common cause. For her, arrows representing the causal ancestry of those constituent events must proceed from different sources (see Figure 1a). I claim that this need not be the case—these arrows can proceed from one and the same source. *Pianos* is a case in point; the causal structure of the case is represented in Figure 1b. There is a causal path from the throw of the ball, through the striking of the first *A*-note, to the striking of the second.

One might object that while the throw of the ball is a common cause of the striking of the *A* notes, those events come about in part also because of the positions of the pianos on each of the balconies. And while the throw is a common cause, the positioning of the first piano precisely *here* is independent of the positioning of the second piano precisely *there*. Thus it may look like the traditionalist has a line of defense. She can say that we judge the striking of the same note to be a

coincidence only insofar as we judge that the positioning of the pianos on each of the balconies are themselves salient causes, and those causes are independent of one another.

But this move does not stave off the counterexample for long. First, by elaborating the background circumstances of the case, we can pump the intuition that the positioning of the pianos is not a salient cause after all. Imagine that the pianos have been sitting on their respective balconies for the past 75 years, that they weigh several tons a piece, that nobody has so much as contemplated moving them in all that time, and so on. Now along comes the kid with his ball. Here it seems entirely natural to treat the throw of the ball as the salient cause and the positioning of the pianos as a background condition.<sup>9,10</sup>

Second, we can modify the counterexample slightly so that the positioning of the first and second pianos do have a common cause, and so are not independent after all. Suppose that instead of pianos on the balconies, there are two electric keyboards. Before arriving on the balconies, these heavy keyboards were being transported by plane. The pilot, realizing that he was dangerously low on fuel, pushed several of them overboard in an effort to reduce the cargo. Sailing through the air, the keyboards landed where they now stand on each of the balconies. The pilot's shove is a common cause of the first and second keyboards being positioned in the way they are. But still, the striking of the same note on both keyboards is a coincidence. Indeed, the fact that the keyboards got to the balconies in the way they did does not make it any less coincidental that when the boy loses control of his ball the very same note is struck. In this scenario, none of the salient causes of one of the constituent events is independent of the salient causes of the other constituent event.

#### 4. Responses to the Counterexample

There are two ways the traditionalist can defend her view against the counterexample. She can argue that (a) it is *not* a coincidence that the same note was struck on the pianos, or that (b) the constituents of the purported coincidence do not have a common cause. In response to the first of these defenses, let me elaborate slightly on the details of the case. The boy, we are to imagine, was not trying to hit matching notes. Moreover, he could have easily thrown his ball at a slightly different angle, or with slightly different force, the wind could have easily been blowing in a slightly different direction, and so on. In any one of these scenarios, different notes on the two pianos would have been struck. To deny, in these circumstances, that it was a coincidence that the same note was struck seems to me to violate strong pre-theoretic intuitions.<sup>11</sup>

It may be helpful to consider a second example, this one involving well-defined probabilities. Imagine that an experiment is conducted in which a random mechanism selects a ball out of a red urn containing 100 numbered balls. Next to the red urn are 100 blue urns lined up in a row, each of which also contains 100 numbered balls. When a ball is drawn from the red urn, the number  $n$  of the ball is recorded, and a random mechanism then selects a ball from Urn  $n$ . The reader may assume

that the entire process is triggered by the flip of a switch, so that there is just a single causal chain running from the flip of the switch through the drawing of the first ball, to the drawing of the second. Now imagine that when this process is set in motion, ball 4 emerges from the red urn, and, lo and behold, ball 4 emerges from Urn 4 as well. Think how natural it would be to remark, upon seeing the second ball emerge: “What a coincidence! The very same number came up in both urns.” We say this knowing full well that the random drawings from both urns participate in a single causal chain—*i.e.*, that the flip of the switch triggers the selection mechanism in the first urn which in turn triggers the selection mechanism in the second urn. So again, we have an example of a coincidence in which the causes of the several component events are not independent. (We’ll return to this second case for further discussion in §6.)<sup>12</sup>

Let’s turn to the second response to the counterexample. In denying that the striking of two *A*-notes has a common cause, the traditionalist must deny in particular that the throwing of the ball is a common cause. But if the throw is not a common cause, then the throw did not cause the first *A*-note to be struck, and the throw did not cause the second *A*-note to be struck. (What sense is there in maintaining that the throw caused one of these events but not the other?) These commitments give rise to two difficulties.

The first concerns the grounds on which we deny that the throw is a cause of each of the *A*-note strikings. Adopting some terminology from Lewis, let us say that an event *e* counterfactually depends on an event *c* if, had *c* not occurred, *e* would not have occurred. Lewis (1987a) famously took causation to be the ancestral of this relation.<sup>13</sup> Whatever one thinks of the counterfactual theory of causation, many philosophers follow Lewis in taking counterfactual dependence of *e* on *c* to be *sufficient* for “*c* caused *e*.” There are, we should note, some apparent counterexamples even to that claim. First, for someone with a non-Lewisian view of counterfactuals in which “backtracking” is allowed, there may be (1) cases in which a cause counterfactually depends on its effect; (2) cases in which two independent effects of a common cause counterfactually depend on one another; and (3) cases in which two events that are constitutively related but not causally related counterfactually depend on one another. Very roughly, Lewis dismisses (1) and (2) by banning backtracking, and he gets around (3) by stipulating that causation is a relation between events that are not constitutively related.

One may, of course, reject Lewis’ view of counterfactuals, and thus deny that the counterexamples to sufficiency are merely apparent. Even so, I think it is right to say that counterfactual dependence is at the very least a good heuristic for causation in the following sense. In cases that do *not* intuitively fall under (1), (2), or (3), the counterfactual dependence of *e* on *c* presents very good *prima facie* support for causal dependence.<sup>14,15</sup>

But consider now the following counterfactuals:

- (A) Had the ball not been thrown, note *A* on Piano 1 would not have been struck.
- (B) Had the ball not been thrown, note *A* on Piano 2 would not have been struck.

Both (*A*) and (*B*) are intuitively true.<sup>16</sup> Moreover, this case does not fall under (1), (2) or (3) above. The striking of the first *A* note did not cause the ball to be thrown; the striking of the first *A* note and the throwing of the ball are not independent effects of a common cause; and the striking of the first *A* note is not constitutively related to the throwing of the ball. Therefore if we accept the heuristic, we ought to allow that the ball caused each of the *A*-notes to be struck.

The second problem for this response to the counterexample involves what I will call “causal danglers.” In denying that the throw caused, *e.g.*, the first *A* to be struck, the traditionalist owes us some account of what *did* cause this event. But it seems that there is little she can say on this score. Once we deny that the throw of the ball caused the striking of the note, it looks like *nothing* did. (What better candidate than the throw?) This punches holes in the causal nexus. It saddles us with causal danglers; bonafide parts of the spatio-temporal world that are without salient causes.

How big of a problem is this? Some philosophers deny that all events have causes. Owens, for example, denies that the coincidental meeting at the train station between my friend and I has a cause. According to him, my friend’s intention to leave on a trip to Africa causes him to be at the station, and my intention to meet my mother causes me to be at the station. But nothing causes the coincidence itself; in particular, our joint intentions are not a cause of the meeting.<sup>17</sup> But the problems raised by our example run deeper. For on Owens’ view, each of the constituent events of the meeting at the train station is causally grounded, and only the meeting itself is not. In general, Owens allows that an event *e* can fail to have a cause, if each of the constituents of *e* is itself causally grounded. The traditionalist, however, must deny not only that the striking of the *same* note on both pianos has a cause, but that each of the constituents of that event—the striking of an *A* here and the striking of an *A* there—has a cause. (Something like Owens’ view can be defended on the grounds of metaphysical parsimony; no similar appeal to parsimony is possible here.)

The arguments presented here are not decisive—one could, for example, flatly deny that holes in the causal nexus are problematic. But they do, I think, put pressure on any attempt to save the traditional view.

### 5. Contrastive and Aspect Causation

I’ve argued that the traditional view of coincidence falls prey to the counterexample presented in §3. My discussion in the forgoing sections made an implicit assumption—namely, that causation is a binary relation between events. In fact, there are really two assumptions here: causation is a *binary* relation, and the relata of the causal relation are *events*. Some philosophers may be tempted to respond to the counterexample in §3 by denying one or the other of these assumptions. Would doing so rescue the traditional view of coincidence, or some sufficiently similar version of that view? I argue in this section that it would not. I first imagine a “contrastivist” response to the counterexample—one that denies the first



assumption—and then turn to an “aspect-theoretic” response—one that denies the second assumption.

### 5.1 Contrastive Causation

We'll focus here on Schaffer's contrastive view of causation. As Schaffer (2005) points out, his contrastivism is a descendant of both van Fraassen's contrastive account of causal explanation and Hitchcock's contrastive account of causation.<sup>18</sup> According to Schaffer, causation is not a binary, but a quaternary relation between events. The true form of a causal statement is: *c* rather than *C* caused *e* rather than *E*. Here *c* and *e* are individual events, and *C* and *E* are contrast classes, or (nonempty) sets of events. Thus, for example, a contrastivist may say that Jane's smoking cigarettes rather than drinking whiskey caused her to develop lung cancer rather than liver malfunction. To say simply that *c* caused *e* is in an important sense elliptical on the contrastivist view.

Of course, we do, as a matter of fact, say that one thing caused another, without making explicit reference to contrast classes. The contrastivist holds that where no such reference is made, contrast classes are supplied by the conversational context in which the assertion takes place. If in a context in which you and I are discussing the relative harms of smoking cigarettes and drinking whiskey, you say that smoking cigarettes caused Jane to develop lung cancer, the causal contrast is naturally taken to include drinking whiskey. Although the focus in Schaffer (2005) is on causation itself, Schaffer also takes causal explanation to be contrastive.<sup>19</sup> Thus Jane's smoking rather than drinking *explains* her getting lung cancer rather than liver malfunction. In general, the form of causal explanation is quaternary.

In proposing a solution to certain difficulties that arise in the theory of causation, Schaffer discusses how shifts in the description of the events *c* and *e* can introduce shifts in contrast class. The problem that Schaffer addresses is given in an example that he adapts from McDermott.<sup>20</sup> As Schaffer puts it, “One might accept that McEnroe's tension caused his serving awkwardly but deny that his tension caused his serving. One wants to say: he was going to serve anyway.” This poses a problem for any view of causation as an extensional relation between coarse-grained events. For if in fact McEnroe served awkwardly, then on a coarse-grained view of events, the event of his serving *just is* the event of his serving awkwardly. And given extensionality, the event under the first description stands in the same causal relations as the event under the second.

Schaffer argues that the difficulties can be solved by a turn to contrastivism. His solution is based on the idea that shifts in the description of the effectual event *e* trigger shifts in contrast class. In particular, when we describe the effectual event as “serving,” the implied contrast is standing still; and when we describe the effectual event as “serving awkwardly,” the implied contrast is serving smoothly. Making these contrasts explicit gives us a full account of the causal relations that obtain, and dissolves the paradox.

The underlying contrastive truth is: McEnroe's being tense rather than calm caused his serving awkwardly rather than smoothly, but it did not cause his serving rather than



standing still. The different descriptions of the effect (“McEnroe’s serving awkwardly” versus “McEnroe’s serving”) suggest different intended contrasts. His tension made a difference to how he served, not to whether he served.<sup>21</sup>

Now it might be thought that by taking up a contrastive account of causation, an appeal to shifts in contrast classes triggered by shifts in description can similarly be used to explain what goes on in *Pianos*. In particular, the contrastivist may be tempted to say that when we describe the effectual event as the ball landing in such and such locations in physical space—just *here* on the first balcony, and just *there* on the second—the implied contrast class consists of events in which the ball lands in different locations in physical space. And when we describe the effectual event as the ball striking two *A* notes, the implied contrast class consists of events in which the ball strikes two different notes. The underlying contrastive truth in *Pianos* is that relative to some salient causal contrast class *C*, the throw of the ball caused the ball to land in those locations in physical space rather than some other ones, but did not cause two *A* notes to be struck rather than two different notes.

Putting things a bit more precisely, let *c* be the event of the throw, and let *e* be the event of the ball striking two *A*s (equivalently, the event of the ball landing in the locations where it did). Moreover, let *E* and *E\** be respectively some collection of events in which the ball lands in different locations, and some collection of events in which the ball strikes two different notes. Then the claim is that *c* rather than *C* causes *e* rather than *E*, but did not cause *e* rather than *E\**. This would be a way of rescuing the traditionalist up to a point: the traditionalist could say that the striking of the *A* notes rather than two different notes do not have a common cause in the throw after all. (To be clear, Schaffer himself nowhere proposes this: the response to *Pianos* is imagined on behalf of a traditionalist looking to contrastivism to escape the counterexample.)

Now there is, I think, an obvious problem with this attempt to shore up the traditional view. It goes back to what events populate the classes *E* and *E\**. If *E* is simply *all* landings of the ball in locations other than the actual ones—or, at any rate, all such landings that could easily have come about given the physical constraints present—and *E\** is the collection of *all* such landings that produce two different notes, then *E\** is just a subset of *E*. For if the ball strikes two different notes, then at least one of the locations where it lands is different from the actual one. So it would be very odd for the contrastivist to say that *c* rather than *C* caused *e* rather than *E*, but did not cause *e* rather than *E\**. (Another way to put this: if the throw relative to *C* is good enough to “rule out” all alternatives to *e* in the bigger set *E*, then how could it not be good enough to rule out all alternatives to *e* in the smaller set *E\**?)

At this point the contrastivist may resort to constructing highly idiosyncratic contrast classes *C* and *E* in order to differentiate causal relations in just the right way. Suppose, for example, that *E* consists *only* of those landings of the ball in different locations from the actual ones which nevertheless produce two *A* notes. Let the causal contrast class *C* consist of throws of the ball that would produce such landings. Then the contrastivist can accurately say that *c* rather than *C* caused

$e$  rather than  $E$ , but  $c$  rather than  $C$  did not cause  $e$  rather than  $E^*$ . The problem with this is that the contrast classes involved on both the causal and effectual side are entirely artificial, and are precisely *not* the sort of contrast classes triggered by the natural language descriptions in the case. But the whole idea behind the imagined contrastivist move to save the traditionalist was to appeal to shifting contrast classes triggered by shifting natural language descriptions.

### 5.2 Aspect Causation

Perhaps a better strategy is to deny the second of our assumptions—that events are the causal relata. According to Paul, causation is a relation between *aspects*. “An aspect is a particular’s (a particular event’s or individual’s) having a property. Aspects are things that correspond one to one with thing-property pairs such that the property is had by the thing; so aspects are in an important sense part of the spatiotemporal world.”<sup>22</sup> Paul (2000) argues that only by understanding causation as a relation between aspects can we block counterexamples to the transitivity of causation—a principle she is eager to preserve.

Can passing from events to aspects save the traditionalist from the counterexample in §3? The aspect theorist may be tempted by the following thought. That the ball landed in the precise locations in physical space where it did and that two  $A$  notes were struck are distinct aspects of the same event. The angle and force of the throw are causally responsible for the ball landing in those locations in physical space, but are not causally responsible for the two  $A$  notes being struck. If that account of causal relations in the case is right, then the striking of the two  $A$  notes do not have a common cause in aspects of the throw after all.

Although something about this is tempting, there are two problems that mirror precisely the problems we encountered above in treating causation as a binary relation between events. First, and most importantly, it is not at all clear on what grounds we should deny that aspects of the throw are causally responsible for the two  $A$  notes being struck. Recalling our earlier arguments, if we take counterfactual dependence between aspects to be at least a good heuristic for aspect causation, then we seem to get the opposite verdict. With Paul, we can define counterfactual dependence of aspects as follows.<sup>23</sup>

**Counterfactual dependence of aspects.** For any two distinct, actual events or individuals  $c$  and  $e$ , and logically distinct properties  $p$  and  $q$ :  $e$ ’s having  $q$  is counterfactually dependent on  $c$ ’s having  $p$  if and only if, had  $c$  occurred (existed) without  $p$ , then  $e$  would not have occurred (existed) with  $q$ .

But now consider the counterfactuals:

- (C) Had the throw of the ball occurred without it being thrown at just that angle, with just that force, and so on, the striking of the first piano would not have been a striking of note  $A$ .
- (D) Had the throw of the ball occurred without it being thrown at just that angle, with just that force, and so on, the striking of the second piano would not have been a striking of note  $A$ .

Are these counterfactuals true? Certainly if the ball had been thrown at a slightly different angle, it *might* not have struck an *A* on the first or second pianos. Whether it *would* not have done so is, I think, underdetermined by the description of the case. So let us now add to that description that the ball could only hit two *A*'s if it were thrown at just the angle, and with just the force with which it was in fact thrown. Surely doing so does not weaken the intuition that the striking of the same note was a coincidence (and may, in fact, strengthen it). In this case, both (*C*) and (*D*) are true. So if counterfactual dependence between aspects is a good heuristic for causation, then the relevant aspects of the throw do cause each of the *A* notes to be struck.<sup>24</sup> (The above remarks apply only to someone who takes counterfactual dependence of aspects as a heuristic for aspect causation; Paul's view is more complex, requiring both what she calls "influence" and "lawful entailment."<sup>25</sup>)

The second problem for the aspect-theoretic solution is that it saddles us anew with causal danglers. If aspects of the throw are not causally responsible for the striking of the *A* notes, then what is? Again, it looks like *nothing* is; the striking of two *A*'s is causally gratuitous. But, according to Paul, aspects are "part of the spatiotemporal world." So in denying that something causes the striking of two *A* notes, we punch holes in the causal nexus. Again, we cannot simply dismiss the problem as one might in the example of the meeting at the train station. There, each component of the meeting (my being at the station and my friend's being at the station) had a cause, even if, according to some, the meeting itself did not. Here the aspect theorist denies not simply that the striking of the *same* note has a cause, but that each of the two *A* note strikings has a cause. And so the picture we get is that there are aspects—bonafide parts of the spatiotemporal world—that are causally gratuitous not because they are in some sense awkwardly conjunctive or gerrymandered. The simpler aspects that compose them are themselves causally gratuitous. How big of a problem is this? Perhaps causally gratuitous aspects are less problematic than causally gratuitous events, although I am not sure why this should be the case.

It should be clear that the objections raised in this section are not objections to an aspect-theoretic view of causation, but rather to the thought that one can save the traditional view of coincidence by turning to aspects rather than events as the causal relata. (Similar remarks apply to the previous section on contrastive causation.)

## 6. Causation versus Causal Explanation

Where does this leave us?

The traditional view is mistaken, I've argued; even the several parts of a coincidence can have a common cause. Passing to a contrastive or aspect theoretic view of causation does not get the traditionalist out of the water. But what, then, of the truism we began with—that coincidences cannot be explained? How can we account for this if the several parts of a coincidence may have a common cause? In the remainder of the paper I want to explore what I take to be an attractive response to that challenge.

Let's start by noting that at least two of the views considered above posit quite a tight connection between causation and (causal) explanation. For Schaffer, not

only is causation contrastive on both the causal and effectual sides, but explanation is likewise contrastive on both sides (see note 19). For Owens, who takes the more traditional view that causation is a binary relation, the reason we cannot explain a coincidence is that a coincidence literally has no cause. These are two examples of the ways in which facts about causal structure are supposed to filter up, as it were, to facts about causal explanation.<sup>26</sup>

But why buy into the identification of causation and causal explanation? Long ago, Davidson taught us to draw a sharp distinction between the two. He argued that causation is an extensional relation between coarse-grained events. To take one of his examples: if it was the drying she gave herself on the beach with a coarse towel that caused the splotches to appear on Flora's skin, then it was the drying she gave herself that did it.<sup>27</sup> Substituting a different, in this case truncated, description of the causal relatum does not change the truth value of a singular causal statement.

Causal explanation, on the other hand, is an intentional relation between facts of a certain kind, or events under a description. One and the same event differently described may explain or fail to explain a given fact; likewise, one and the same event differently described may be explained or fail to be explained by a given fact. Suppose, to take an example from Ruben (1990), that I strike a match and it catches fire. My striking of the match is the penultimate thing that ever happened to the match. We have here two different descriptions of the same event, but only one of those descriptions is explanatory. The fact that the match was struck explains why the match lit. The fact that the penultimate thing that ever happened to the match occurred does not explain why the match lit.<sup>28</sup> To explain why it was that the match lit, it is not enough to simply cite the cause of that event under any description whatsoever.

The same is true if we instead fix the description of the cause and vary the description of the effect. Consider Gavrilo Princip and the Archduke Franz Ferdinand in Sarajevo, in an example adapted from Mackie (1974) and Searle (1983). Princip's pulling of the trigger caused Ferdinand to die. And Ferdinand's death was arguably the event that started the First World War. It's quite natural to say that Princip pulling the trigger explains why Ferdinand died; but his pulling the trigger does not explain why the First World War began.<sup>29</sup>

These toy examples show that facts about what explains what do not simply *filter up* from facts about what caused what. But there is, I think, a broader point. In general, we should be quite wary of inferring facts about what can or cannot be explained from facts about what caused what. For to explain is not simply to cite a cause. I argued above that in *Pianos*, the throw of the ball caused each of the *A* notes to be struck. By separating causation from causal explanation, however, we can still maintain that there is no explanation for the fact that the same note was struck. This even though the striking of the same note just was, as a coarse-grained event, the striking of an *A* here and an *A* there, and those two events have a common cause in the throw. Or, to revisit the example from §4, while the flipping of the switch caused the two 4-balls to emerge from their respective urns, the fact that the switch was flipped does not explain the fact that it was the very same ball in both urns. This even though the emergence of the same ball just was, as a

coarse-grained event, the emergence of a 4-ball here and a 4-ball there, and these events have a common cause in the flip of the switch.

The general move to pry apart causation and explanation allows us to make sense of the idea that even where there is a common cause, explanation of certain facts may not be forthcoming. But of course more needs to be said about why those facts cannot be explained—or, more cautiously, cannot be explained in such a way as to remove the coincidental nature of what happened. So let me now stick my neck out a bit further. In general, it seems obvious to me that there is a difference between explaining on the one hand why *a* has the property *F* and separately why *b* has the property *F*, and on the other hand explaining why whatever properties *a* and *b* have, they stand in some particular relationship. Consider again *Pianos*. The fact in need of explaining is that the *same* note is struck. Even if we could show, separately, why the note struck on the first piano was an *A* and why the note struck on the second piano was an *A*, it seems to me we would not yet have explained why it is that the *same* note was struck on the two pianos. Why not? This last fact is a *relational* fact; the former are not. To explain the relational fact, it is not enough to show separately why each of the relata had to be what it was. We want a reason for the relationship itself, the fact that the notes *matched*. Likewise, in the case of the two urns. Even if we could explain separately why ball 4 emerged from the first urn and ball 4 emerged from the second urn (it is doubtful that we can), it seems to me we would not yet have explained why the *same* ball emerged from the two urns. To do so, we would have to exhibit some relationship between the numbers on the two balls—a relationship that is in this case incompatible with the fact that the two numbers are probabilistically independent.

Presumably providing an explanation of the right kind for the relevant relational facts—one that would remove the coincidental nature of what happened—would involve saying what one note taking the value it does *has to do with* the other taking the value that it does, or in the case of the urns, what the number that emerges from the first urn has to do with the number that emerges from the second. But of course in the examples as presented, an *A* on the first piano does not guarantee or make more probable an *A* on the second; a *B* on the first piano does not guarantee or make more probable a *B* on the second. (The boy, we are to imagine, could have easily thrown his ball in such a way as to hit an *A* here and a *B* there, or vice versa.) So we cannot give an account of the events that identifies a reason, not simply for each note to be what it is, but for the notes to be the same. Likewise in the urn example, a 4 in the first urn does not guarantee or make more probable a 4 in the second; a 5 in the first urn does not guarantee or make more probable a 5 in the second. The only thing a 4 in the first urn guarantees is that the next ball will be chosen from *Urn 4*, not that the *number* of that ball will itself be 4. So we cannot give an account of the events that identifies a reason why the balls drawn from the two urns were the same.

Some philosophers will perhaps feel uneasy about the claim that even if we could explain why it was an *A* on the first piano, and separately why it was an *A* on the second, we would not yet have explained why it was the very same note. What more is there to the notes being the same, they may say, than the first and second both being *A*'s? But careful: this, I believe, is another piece of the same faulty reasoning

we've been discussing. From the fact that *as events* there is nothing more to the striking of the same note than the striking of two *A*'s, it does not follow that there is nothing more to *explaining* why the notes were the same than explaining why each individually was an *A*. (We explain facts, and not events. The fact that the same note was struck is different from the fact that an *A* was struck here and an *A* was struck there. This can be seen by noting that the former does not entail the latter.)

But never mind. To the very insistent philosophers of this stripe: the point of these remarks is not so much to legislate what does and does not count as an explanation of *any* kind, but rather to zero in on what kinds of explanation our judgments about coincidence are sensitive to. One could put all of this a bit more cautiously by saying that there are both "thin" and "thick" explanations of relational facts. In explaining why the first note was an *A* and, separately, why the second note was an *A*, we give at best a thin explanation for the sameness. We explain separately two non-relational facts which together entail the relational fact. But we offer no reason for the relational fact itself, over and above whatever reasons there are for each of the non-relational facts to obtain. To say this is emphatically not to suggest that there always *are* reasons for relational facts over and above the reasons for the relevant non-relational facts. Indeed coincidences, on one way of understanding the truism, are relational facts for which there are no explanations of the "thick" kind.

### 7. Explaining a Purported Coincidence

The remarks in the previous section are schematic. A full defense of them would require a full-blown theory of explanation, a task that is beyond the scope of this paper. Instead, I want to explore in this section the kind of thing that *would* count as an explanation of the relevant relational facts, thereby bringing into focus a bit more fully what is missing in the case of genuine coincidences.

One natural place to begin is by imagining what it would take for us to reverse our judgments in *Pianos*—to judge, that is, that the striking of the same note on both pianos is not a coincidence after all. Suppose that instead of losing control of the ball, the boy in the courtyard of the apartment complex set out with the express purpose of hitting the same note on each piano. If he skillfully pulled this off, we would not say it was merely a coincidence that the same note was struck. After all, this is precisely what he set out to do.

But it would be easy, I think, to make too much of the role of intentions here. (The *intentional* is more properly opposed to the *accidental* than to the *coincidental*.<sup>30</sup>) So let us suppose, instead, that the kid did not intend to hit the pianos, but simply threw his ball too high. This time, however, imagine that there is a thick wall between where the kid is playing and the apartments. The ball must pass through a small hole (or "tunnel" if you will) in the wall if it is to reach the balconies of the two apartments. Because of the precise position of the hole with respect to the balconies, any ball that passes through the hole produces the very same note on both pianos. An *A* on the first piano is invariably followed by an *A* on the second, a *B* on the first piano is invariably followed by a *B* on the second, and so on. Imagine that when the boy loses control of his ball, it goes through the hole in the wall and

strikes two *A* notes. Here I think we would no longer say it was only a coincidence that the same note was struck. Given that the ball reaches the two pianos at all, the note struck on the second piano is bound to match that struck on the first.

In the modified version of *Pianos* we're considering, facts about the physical setup (wall, hole, and pianos) explain why it is that the same note is struck on the two pianos. They provide, to be clear, an explanation not of why the ball landed in this or that location in physical space, and not of why it was in particular an *A* struck on each of the pianos, but of a relational fact—the fact that the two notes matched. This is precisely the sort of explanation we lack, I argued, in the original *Pianos*.

Consider another example.

**Deck of Cards.** Two friends, Dylan and Augie, visit a casino together. They split up toward the beginning of the night, and make their way among the betting tables independently. At one of the tables, a card dealer asks Augie to select a single card from a full, covered deck. The card he pulls at random is the Queen of diamonds. Several hours later Dylan visits the same table. The card dealer asks Dylan to select a card, and the card he pulls is also, as it happens, the Queen of diamonds.

In *Deck of Cards*, it is only a coincidence that Augie and Dylan picked the same card. The cards could have easily been shuffled slightly differently; Augie could have easily pulled the card just below or just above the one that he in fact pulled while Dylan still chose the Queen of diamonds, and vice versa. In all of these scenarios, Augie and Dylan would have chosen different cards.

But now consider a situation in which we would deny that it was a coincidence that Augie and Dylan picked the same card. Suppose, for example, that because of the particular way Augie handles his cards, he makes a mark on the Queen of diamonds as he pulls it out of the deck, causing it to protrude slightly from the rest of the cards. When it comes time for Dylan to choose, his eye is drawn inadvertently to the protruding card. Augie could have easily (let us suppose) drawn the Jack of spades. But had Augie drawn the Jack of spades, he would have left a mark on that card, and Dylan would have drawn *it* several hours later. Under these circumstances, we would deny that it was merely a coincidence that Augie and Dylan chose the same card. Those same circumstances explain why Augie and Dylan picked the same card. They explain the relational fact that could not be explained in the case where each of them chose at random.

Consider one final example, a simplified version of the urn case considered above.

**Urns.** There are two urns, each containing 100 numbered balls, and in each of these urns there is a physical mechanism for mixing up the balls and shooting out one of them. First the “mix and shoot” mechanism is activated in the first urn, and a ball is shot out. The momentum of this ball coming down the chute causes the activation of the mix and shoot mechanism in the second urn, and a ball is shot out there too. We can imagine that the mix and shoot mechanism in the first urn is itself triggered by an On-Off switch, so that there is just a single causal chain running from the flipping of the switch to the shooting out of the first ball, and then the shooting out of the second



ball. Now suppose that when the switch is flipped, the balls that emerge from the two urns both have the number 4 printed on them.

It is only a coincidence that the same number comes up in both urns. Sure, the flipping of the switch caused the 4-ball to emerge from Urn 1, and also caused the 4-ball to emerge from Urn 2. But if we ask, “Why did the same number come up in both the first and the second urn,” there is nothing intelligent we can say on this score. A 4-ball in the first urn does not make a 4-ball in the second urn more likely; the number drawn in the second urn is probabilistically independent of the number drawn in the first. (Note that what makes this a coincidence is not simply that the combination of  $M$  balls each labeled, say, “4” has low probability. Assuming a uniform distribution, any particular sequence has the same low probability. But there are many, many other sequences which are not cognitively salient to us in the same way, and which, if they occurred, would not prompt us to remark, “What a coincidence!”<sup>31</sup>)

But consider now a situation in which we would deny that it was a mere coincidence that the same number came up in the two urns. Suppose that each of the balls in the two urns has a magnet attached to it in such a way that matching balls (balls with the same number) attract each other and repel all other balls. As the 4-ball emerges from the first urn and goes down the chute, it not only triggers the mix and shoot mechanism of the second urn, but brings to the fore the 4-ball in that urn. We can imagine a physical setup in which the magnet mechanism was perfected to such a degree that whatever ball emerges from the first urn is, with very high probability, matched by the same number emerging from the second urn. The circumstances just described are ones in which we would deny that it was a mere coincidence that the same number came up in both urns. Those circumstances provide an explanation for why the *same* ball emerges; they explain the relational fact that could not be explained in the original scenario, even though in both cases there is just a single causal chain running through the several events.

Note one interesting feature of these examples. In each of the modified scenarios, the circumstances described explain the relevant relational fact *without* explaining the non-relational circumstances that are so related. In the modified version of *Pianos*, for example, the position of the hole in the wall with respect to the keyboards explains why it was the *same* note on each of the two pianos, but does not explain why it was in particular an *A* that was struck on each. Likewise in the modified version of *Deck of Cards*, the mark Augie makes, the protrusion from the deck, and so on, explain why Dylan picked the same card as Augie, but do not explain why it was in particular the Queen of diamonds that they both picked. Finally in *Urn*s, the magnet mechanism explains why the same number came up in both urns, but does not explain why it was in particular a 4. Examples like this show that one can sometimes explain a relational fact without thereby explaining the particular occurrences which are so related.

Very often (but perhaps not always) those explanations take the form of exhibiting a probabilistic relationship between the several parts of the purported coincidence; of showing how, from some natural perspective, the occurrence of one

part of the coincidence makes the occurrence of the other(s) more probable, or even certain. In the modified *Pianos* for example, given the position of the hole in the wall, an *A* on the first piano is *bound* to be matched by an *A* on the second. In the modified *Urns*, the emergence of a number *k* ball in the first urn is, with very high probability, matched by the emergence of a number *k* ball in the second.<sup>32</sup> These probabilistic relationships—or the recognition of them—defeat the judgment that what happened was a mere coincidence. By contrast, the recognition that the several parts of a coincidence have a common cause does not defeat that judgment.

Coming full circle, our original problem was to say why, even when the components of a coincidence have a common cause, the coincidence cannot be explained. The answer I've given is that two component events can issue from a common cause without there being an explanation of the appropriate kind for the relational fact, and in particular without their being a probabilistic relationship between the several parts of the coincidence.<sup>33</sup> The flip of the switch is a cause of each of the balls emerging from their respective urns—they would not have emerged were it not for the flip. But there is nevertheless no probabilistic relationship between the *numbers* that emerge in the first and second urns that would explain the match. And again in *Pianos*, the throw is a common cause of the striking of each of the pianos—neither of the keys would have been struck had the kid not thrown his ball. But there is nevertheless no probabilistic relationship between the *note* struck on the first piano and the note struck on the second. An *A* on the first piano does not make more probable an *A* on the second. The relational fact that the same note was struck is not in any way explained, illuminated, or made intelligible by appeal to the common cause.

## 8. Conclusion

I've argued that, when it comes to coincidence, talk of causal structure is a red herring. The traditionalist's thesis that the several parts of a coincidence do not have a common cause issues the wrong verdicts in specific cases, and misleads us as to why coincidences cannot be explained. By separating causation and causal explanation, we can refrain from making implausible commitments about the causal structure of a coincidence, and from poking holes in the causal nexus. Although there are many things we can say about the causal antecedents of a coincidence, certain explanations, as the truism rightly informs us, are simply not to be had.

## Acknowledgements

I thank John Campbell, John Collins, Branden Fitelson, Greg Gandenberger, Philip Kitcher, John Morrison, Christopher Peacocke, Achille Varzi, and two anonymous referees for very helpful discussions and comments on early drafts of the paper.

## Notes

<sup>1</sup> Owens (1992), pg. 13.

<sup>2</sup> Owens (1992), pg. 17.

<sup>3</sup>The thesis that there *are* mathematical coincidences it itself controversial. For a defense of this thesis see, e.g., Lange (2010).

<sup>4</sup>Hart & Honore (1959), pg. 74.

<sup>5</sup>Translation: “But in other situations, the notion of chance has an essential and not simply functional meaning. This is true, for example, of that which we can call ‘absolute coincidences,’ which is to say those that result from the intersection of two causal chains completely independent from one another.”

<sup>6</sup>Lewis (1987b), pg. 214.

<sup>7</sup>As Lange says, “For example, the Big Bang is a common cause of the CIA agent’s presence and His Excellency’s death. Typically, however, the Big Bang is too remote to qualify as a relevant kind of common cause, and so it remains coincidental that the CIA agent was in the capital just when His Excellency dropped dead.” Lange (2010), p. 317.

<sup>8</sup>Thus no causal “branching” structure is implied. In particular, *c* may be a common cause of *e* and *f* even if there is a causal path from *c* to *f* that goes through *e*. It should be noted that one can also construct counterexamples to the traditional view which have a causal branching structure. Here is a modification of *Pianos* suggested to me by Greg Gandenberger: Instead of playing with a ball, the child in the courtyard is playing with a slingshot that launches two rocks simultaneously. The release of the slingshot causes the two rocks to hit the two high *A* notes on the pianos. There are three events: (1) the release of the slingshot, (2) the striking of the high *A* note on the first piano, and (3) the striking of the high *A* note on the second piano. There are direct causal paths from (1) to (2) and from (1) to (3). Moreover, no causal path from (1) to (3) goes through (2).

<sup>9</sup>If you aren’t moved by pianos, imagine instead that the boy is playing with his ball in front of two rock formations, each resembling a deer. When the boy loses control of his ball it first hits the tip of the nose of one deer, bounces off it, then hits the tip of the nose of the other deer before falling to the ground. The rock formations themselves have been there for millennia. Here again it seems very natural to treat the throw of the ball as a salient cause and the positions of the rock formations as only a background condition. I thank John Collins for suggesting this modification to the example.

<sup>10</sup>Yet another problem: even if the positioning of the pianos is taken to be a salient cause, there is still of course another salient cause that the striking of the *A* notes have in common—namely, the throw of the ball. So the traditionalist who adopts this tack would have to say that the striking of the same note is only partly a coincidence. But this seems wrong. The striking of the same note was *entirely* coincidental.

<sup>11</sup>Note that it is not simply the *low probability* of what happened (from some natural perspective) that makes the striking of the same note on the two pianos a coincidence. It may well be that the probability of striking some particular *E* on the first piano followed by some particular *C* on the second piano is just as small as striking the very same *A* on the two pianos. But we would not be inclined to say, upon witnessing that latter event, “What a coincidence! An *E* was struck here, and a *C* there.” The reason for this is that the striking of the *same* note is cognitively salient to us in a way that that other state of affairs is not. Not every low probability event is cognitively salient.

<sup>12</sup>I thank an anonymous referee for suggesting this example (here slightly modified). In fact, the referee claims that it is *not* a coincidence, but only a “low probability event,” that ball 4 is drawn from both urns. I think, to the contrary, that this is a prototypical example of the folk notion of coincidence, and could serve just as well as *Pianos* to make the point against the traditionalist. The judgment that it is a coincidence stems not only from the fact that this event has low probability, but also from the fact that there is a cognitively salient match between component events; had number 4 emerged in the first urn and number 6 emerged in the second, we would hardly be tempted to remark, “What a coincidence!”

<sup>13</sup>Two events, *c* and *e*, are related by the ancestral if there is a chain of events from *c* to *e* in which each node in the chain counterfactually depends on the previous node.

<sup>14</sup>Another way of saying this: If *e* does counterfactually depend on *c* in a case that does not fall under (1), (2), or (3), we would need very good reasons to *deny* that *c* is a cause of *e*.

<sup>15</sup>Note that the difficulties for the counterfactual theory of causation most widely discussed today—cases of late preemption and trumping—are objections to the *necessity* of Lewis’s conditions, not sufficiency. See, e.g., Collins (2007), Schaffer (2000), and Collins et al. (2004).

<sup>16</sup> We can stipulate as part of the example that there are no preempted causes waiting to take effect in the event of an unsuccessful throw.

<sup>17</sup> Indeed, according to Owens, no coincidence has a cause. The separate component events of a coincidence each has a cause, but the coincidence itself does not. In particular, the combination of the causes of component events is not a cause of the event they compose. See Owens (1992).

<sup>18</sup> As Schaffer explains, for Hitchcock causation is contrastive on the *c*-side. The true form of a causal statement is: *c* rather than *C* caused *e*. For van Fraassen, on the other hand, causal explanation is contrastive on the *e*-side. The true form of an explanatory statement is *c* causally explains *e* rather than *E*. Schaffer's view differs from both in that he thinks causation is contrastive on both ends; his view differs from van Fraassen in that he thinks causation itself—rather than just causal explanation—is contrastive. See van Fraassen (1980) and Hitchcock (1996). An earlier contrastive account of explanation was given by Dretske in Dretske (1970).

<sup>19</sup> See Schaffer (2005), p. 325, note 17, where Schaffer says that explanation is contrastive. More specifically, Schaffer believes [personal communication] that explanation is, like causation, contrastive on both the causal and effectual sides.

<sup>20</sup> See McDermott (1995), p. 540, and Schaffer (2005), p. 307.

<sup>21</sup> Schaffer (2005), pg. 308.

<sup>22</sup> Paul (2000), pg. 244.

<sup>23</sup> Paul (2000), p. 245.

<sup>24</sup> In evaluating the truth of counterfactuals (*C*) and (*D*), we hold fixed the positions of the pianos on each of the balconies. We do this because these pianos were positioned where they now stand well before the ball made its way through the air. It follows that one cannot object to the claim that “Had the ball been thrown at a slightly different angle, it would (or might) have hit two different notes” by pointing out that *had the pianos been differently positioned*, the ball might still have hit the two *A* notes.

<sup>25</sup> See Paul (2000), p. 251. It may seem as though in the present example, aspects of the throw influence what note is struck on the two pianos in the sense Paul requires but aspects of the throw do not lawfully entail facts about which note is struck. Perhaps once we supplement aspects of the throw with aspects of the positioning of the two pianos, we get lawful entailment of the relevant kind; this seems to me to depend in a critical way on what one takes to be admissible as a causal law. At any rate, I think this is the right place for someone like Paul—or rather, someone looking to save the traditional view of coincidence by turning to a view like Paul's—to push back on the view defended here.

<sup>26</sup> See note 17. Owens himself has a non-standard view of the causal relata. “A canonical causal statement employs the sentential connective “because”: the house burned at *t*(1) because the match was struck at *t*(0). It is widely agreed that “causally explains” is a sentential connective and I wish to equate causation with causal explanation. It is also said that causes are events. What gives this position its appeal is phrases like: the house was caused to burn at *t*(1) by the match's being struck at *t*(0). I take “the match's being struck” to be a nominalisation of the sentence “The match was struck” which appears in the canonical rendering of that statement. On this understanding, I shall often speak of causes as events.” Owens (1989), p. 49–50.

<sup>27</sup> See Davidson (1967), p. 698.

<sup>28</sup> The example is taken from Ruben (1990), p. 189.

<sup>29</sup> The example is adapted from Mackie (1980), p. 248–249, and Searle (1983), p. 98–100.

<sup>30</sup> From the fact that the kid did not intend to strike the same note, we cannot conclude that it was only a coincidence that the same note was struck. This is in part the point of introducing the modified version of *Pianos* below, where the ball must pass through a hole in the wall if it is to reach the balconies.

<sup>31</sup> If you don't think that two 4-balls is a coincidence, consider now increasing the number of urns and the number of balls in each urn so that *M* urns are lined up in a row, each containing *N* balls. The shooting out of the ball from one urn triggers the “mix and shoot” mechanism of the next urn, so that there is just a single causal chain. Suppose that for all *M* urns, the very same number ball is shot out. As we increase *M* and *N*, it's difficult to deny that the result is a coincidence (I would like to thank Philip Kitcher for very helpful discussion of this example.)

<sup>32</sup> It is important to keep in mind that the probabilities mentioned are calculated from some natural perspective. Take the original *Pianos*, where there is no wall between the boy and the balconies, and the striking of the same note is a genuine coincidence. Given the exact position of the boy's hand one

millisecond prior to letting go of the ball, the direction of the wind, and so on, one could perhaps calculate (given enough resources) exactly where the ball lands. Conditional on all of this information, one might say, the probability of hitting the same note on both pianos is extremely high, and therefore what happens is not a coincidence. But the fact is that this particular perspective—a millisecond before the throw, with all physical parameters known—is not the natural perspective we adopt when we (rightly) declare that what happened was a coincidence. The perspective we adopt is a more natural one, influenced by all of the nearby epistemic possibilities in which the ball is thrown at a slightly different angle, the wind is blowing in a slightly different direction, and so on. From *this* perspective, it is not the case that an *A* on the first piano makes an *A* on the second more likely, or that it is highly probable that the same note will be struck on both pianos.

<sup>33</sup> When we talk about the probabilistic relationships that obtain, or do not obtain, between the several parts of a coincidence, it is crucial that we think in terms of something finer than coarse-grained events. In the *Urn* case, for example, the emergence of *a* ball in the first urn is the same coarse-grained event as the emergence of the 4-ball (because, in fact, the ball that emerged was a 4-ball). And likewise, the emergence of *a* ball in the second urn is the same coarse-grained event as the emergence of the 4-ball. But there *is*, of course, a probabilistic relationship between the emergence of *a* ball in the first urn, and the emergence of *a* ball in the second. By contrast, there is no probabilistic relationship between the *number* that emerges in the first urn and the number that emerges in the second. Clearly what matters to our judgment is the latter fact. A fuller treatment of this topic would delve into this point more deeply.

## References

- Ned Hall, Laurie Ann Paul, and John Collins, editors. *Causation and Counterfactuals*. MIT Press, Cambridge, Mass, 2004.
- John Collins. Counterfactuals, causation, and preemption. In D. Jacquette, editor, *Philosophy of logic*, pages 1127–1144. Amsterdam: Elsevier, 2007.
- Donald Davidson. Causal relations. *Journal of Philosophy*, 64(21), 1967.
- Fred Dretske. Epistemic operators. *Journal of Philosophy*, 67(24), 1970.
- Herbert Lionel Adolphus Hart and Anthony Maurice Honoré. *Causation in the Law*. Oxford University Press, 1959.
- Christopher Read Hitchcock. The role of contrast in causal and explanatory claims. *Synthese*, 107(3):395–419, 1996.
- Paul Horwich. *Probability and Evidence*. Cambridge University Press, 1982.
- Marc Lange. What are mathematical coincidences (and why does it matter)? *Mind*, 2010.
- David Lewis. Causal explanation. In *Philosophical Papers*, volume II. Oxford University Press, 1987.
- David Lewis. Causation. In *Philosophical Papers*, volume II. Oxford University Press, 1987.
- John Leslie Mackie. *The Cement of the Universe*. Oxford: Oxford University Press, 1980.
- Michael McDermott. Redundant causation. *British Journal for the Philosophy of Science*, 40:523–544, 1995.
- Jacques Monod. *Le Hasard et la Necessite: Essai sur la philosophie naturelle de la biologie moderne*. Editions du Seuil, 1970.
- David Owens. Causes and coincidences. *Proceedings of the Aristotelian Society*, 90:49–64, 1989–1990.
- David Owens. *Causes and Coincidences*. Cambridge University Press, 1992.
- Laurie Paul. Aspect causation. *Journal of Philosophy*, 97(4):235–256, April 2000.
- David-Hillel Ruben. *Explaining Explanation*. Routledge, London, 1990.
- Jonathan Schaffer. Trumping preemption. *Journal of Philosophy*, 97(4):165–181, 2000.
- Jonathan Schaffer. Contrastive causation. *The Philosophical Review*, 114(3):297–328, July 2005.
- John Searle. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press, 1983.
- Elliot Sober. Common cause explanation. *Philosophy of Science*, 1984.
- Bas van Fraassen. *The Scientific Image*. Oxford University Press, 1980.