Chancy Causation and the Problem of Aggregate Events

Johann Frick (jdfrick@princeton.edu)

1. Introduction: Causation as "Causal Dependence"

In his classic article "Causation", David Lewis (1973a) gives an analysis of singular event-causation that is centered around the notion of "causal dependence". According to Lewis, an actual event *c* is a cause of an actual event *e* if and only if (i) *e* causally depends on *c*, or (ii) there is a chain of events d_1, d_2, ... d_n, such that d_1 causally depends on *c*, d_2 causally depends on d_1, ..., and *e* causally depends on d_n.[1] The concept of causal dependence is cashed out in terms of counterfactual conditionals: an actual event *e* causally depends on an actual event *c* if and only if, had *c* not occurred, *e* would not have occurred.[2]

Lewis's original account of causation in his 1973 paper takes place against the background assumption of determinism. But what if the assumption fails? It is likely that the fundamental dynamical laws of our world are not deterministic, but probabilistic. The laws of nature, plus the state of the world

[1] Clause (ii) of the analysis is required to deal with the problem of *early preemption*. (See Lewis 1986a, pp. 200-202). We have a case of early preemption when there are multiple would-be causes for an outcome, but one preempts the other would-be causal process from running to completion. To take a common example, after Menzies and Beebee (2019): Celine and Alfred are two assassins who have made a pact that one of them will shoot the dictator when he appears on his balcony at noon. At noon, Celine and Alfred both take aim at the dictator (events *c* and *a*, respectively). Celine pulls her trigger and fires a shot that kills the dictator; however, Alfred desists from pulling his trigger when he sees Celine fire her gun. In this case, Celine's taking aim intuitively is a cause of the death of the dictator (event *e*), while Alfred's taking aim is merely a preempted potential cause. However, the death of the dictator does not causally depend on Celine taking aim. Had Celine not taken aim, the dictator would have been killed by Alfred instead. Taking causation as the ancestral of counterfactual dependence, as Lewis does, solves the problem with early preemption; for there is a chain of causal dependencies running from Celine's taking aim to the death of the dictator, but not from Alfred's taking aim to the death of the dictator.

[2] Lewis defends a possible world semantics for counterfactuals. Such a semantics states truth conditions for counterfactuals in terms of similarity relations between possible worlds. According to this account, the counterfactual "had *c* not occurred, *e* would not have occurred" is true in the *actual* world if and only if there is some possible world in which *c* does not occur and *e* does not occur that is closer to the actual world than any possible world in which *c* does not occur and *e* does occur. See Lewis (1973b).

at a time, make later states of the world more, or less, objectively probable. Surely, even in such an indeterministic world, there is causation?

Lewis agrees. In his "Postscripts to 'Causation'" (1986b) he maintains that the fact that some event e had a probability $p < 1$ of occurring does not imply that e was not caused by some other event. In order to accommodate the possibility of "chancy" causation, Lewis proposes the following revision to his notion of causal dependence: Instead of couching causal dependence in terms of counterfactual conditionals about the *occurrence* of events, causal dependence is now put in terms of counterfactual conditionals about the *probability* of an event's occurring. Specifically, an actual event e causally depends on an actual event c if and only if, had c not occurred, e's probability of occurring would have been much less than its actual chance of occurring.[3]

In this paper, I take a critical look at the prospects for a causal dependence account of chancy causation. Philosophers have for some time suspected that Lewis's theory of chancy causation may confront a problem of false *positives*, in that there may be events that significantly raise the probability of an event e without, for all that, being a cause of e.[4] What has not been widely noted, by contrast, is that Lewis's account also faces a serious problem of false *negatives*.[5]

---

[3] Here, and throughout the paper, the terms "chance", "probability", and "likelihood" are used synonymously. The probabilities involved in Lewis's theory of chancy causation should be understood as objective, single-case chances, indexed to a time t immediately after the occurrence of c. For further discussion, see Lewis (1980).

[4] See, in particular, Hitchcock (2004). One of Hitchcock's cases is briefly discussed in Section 2 below.

[5] I here exclude, as posing *no* serious challenge to Lewis's theory of chancy causation, what may be called the "old" problem of false negatives. This problem was thought to arise in cases where the actual cause of an event appeared to *lower* the probability of the event's occurring. A canonical example, due to Suppes, is given by Hitchcock (2004, 404): "A golfer lines up to drive her ball, but her swing is off and she badly slices the ball, sending it on a trajectory well to the right of the hole. Her slice decreases the probability that it will land in the cup for a hole-in-one. By chance, however, the ball bounces off a tree trunk at just the right angle to send it on a trajectory back toward the cup. As it happened, her slice did cause the ball to go into the cup, even though the slice lowered the probability of this outcome." For cases of this type, I am entirely persuaded by the solution presented by Christopher Hitchcock in Hitchcock (1996). According to Hitchcock, many causal claims are inherently contrastive in nature; they are true or false *relative* to a specific alternative. But often, as in Suppes' example, there is no uniquely correct alternative that we must refer to for the purpose of counterfactual comparisons of probability. If the relevant alternative is one in which the golfer does not hit the ball at all, then her slicing the ball certainly *did* raise the probability of a hole-in-one. In this case, it seems entirely apposite to say that the golfer's slice *caused* the hole-in-one. By contrast, if the relevant alternative is one in which the golfer hits the

It is a direct corollary of Lewis's causal dependence account that, in certain contexts, an event *c* which only *slightly* raises *e*'s probability of occurring cannot be a cause of *e*. While not overly troubling at first blush, I will argue that this feature of Lewis's view has implications that are hard to defend.

The problems with Lewis's theory can be brought out particularly clearly by considering its implications for causal claims about *aggregate events*, a notion which I explicate in Section 3. The first of these problems is that Lewis's theory implies, implausibly, that an event *c* may cause an aggregate event *E* without being a cause of any of the "constituent events" $e_1, e_2, \ldots e_n$ that together *make up* the aggregate event. This is the Problem of Aggregate Effects, discussed in Section 4. Lewis's view also implies, again implausibly, that a group of events $c_1, c_2, \ldots c_n$ can together cause an effect *e*, without *any* of $c_1, c_2, \ldots c_n$ being a cause of *e*. This is the Problem of Aggregate Causes, discussed in Section 5. Finally, it implies that one and the same event *c* can fail to be a cause of *e*'s occurring yet also be a cause of *f*'s *not* occurring, even when *f* just *is* "the non-occurrence of *e*". This is the Problem of Causation and Prevention, discussed in Section 7. These three problems, I claim, seriously undermine the attractiveness of a causal dependence approach to chancy causation, lending support instead to an account, briefly sketched in Section 6, in terms of causal *contribution*.

Before we can begin reviewing the problem cases, however, we must first remind ourselves of Lewis's reasons for proposing his revised definition of causal dependence under indeterminism.

2. Chancy Causation and the "Restriction"

Why does Lewis see the need to revise his original 1973 notion of causal dependence to accommodate the phenomenon of chancy causation? There are certainly some cases of causation under indeterminism that his original analysis handles without trouble: Suppose that some event *e* would have had 0% chance of occurring, had *c* not occurred, whereas it has some objective probability *p* of occurring if *c* does occur. Suppose further that both *c* and *e* occur. In that case,

---

ball *squarely*, her slicing the ball lowered the probability of a hole-in-one. But here, Hitchcock points out, rather than saying that the golfer's slice "caused" the hole-in-one, it is more natural to speak of the ball landing in the cup *despite* the badly sliced shot. The old problem of false negatives derived its *prima facie* cogency by vacillating between these two alternatives. Once they are appropriately disambiguated, the problem disappears.

e causally depends on *c*, in the unreconstructed 1973 sense of the term, and hence *c* is a *cause* of *e*, even on Lewis's original analysis. The fact that *e* was never *certain* to occur is neither here nor there. What matters is that *e* was certain *not* to occur, had *c* not occurred.

Another class of cases, however, proves more troublesome. Here is Lewis:

> [T]here is a second case to be considered: *c* occurs, *e* has some chance *x* of occurring, and as it happens *e* does occur; if *c* had not occurred, *e* would still have had some chance *y* of occurring, but only a very slight chance since *y* would have been very much less than *x*. (Lewis 1986a, 176)

The following case may serve as an illustration:

> *Cancer:* Experimental evidence shows that upon exposure to the carcinogen diazonium metabolite, mice have an 88% chance of developing a stomach tumor. A mouse is exposed to diazonium metabolite, and does, in fact, develop a tumor. However, we also know that, prior to being exposed to diazonium metabolite, the mouse already had an independent 3% chance of developing a stomach tumor. Can we still say that exposure to diazonium metabolite was *a cause* of the mouse's developing the tumor?

Here, it is not clear that, had the mouse not been exposed to diazonium metabolite, it would not have developed a stomach tumor. (Indeed it *would* have, in 3% of cases). Granted, the mouse would have been much less *likely* to develop a tumor. But that is not enough to satisfy the counterfactual conditional that the original notion of causal dependence relies on. Under Lewis's original analysis, we therefore cannot say that administering diazonium metabolite was a cause of the mouse's developing a stomach tumor. This seems wrong.

There is one solution to this problem that Lewis explicitly rejects:

> I have said that if distinct events *c* and *e* both occur, and if the actual chance of *e* (at a time *t* immediately after *c*) is sufficiently greater than the counterfactual chance of *e* without *c*, that

> implies outright that *c* is a cause of *e*. Some philosophers find this counterintuitive. They would correct me thus:
> "No; if there would have been some residual chance of *e* even without *c*, then the raising of probability only makes it *probable* that in this case *c* is a cause of *e*. Suppose, for instance, that the actual chance of *e*, with *c*, was 88%; but that without *c,* there would still have been a 3% probability of *e*. Then most likely (probability 97%) this is a case in which *e* would not have happened without *c;* then *c* is indeed a cause of *e*. But this just might be (probability 3%) a case in which *e* would have happened anyway; then *c* is not a cause of *e*. We can't tell for sure which kind of case this is." (Lewis 1986a, 180)

This objection, Lewis writes, "presupposes that the case must be of one kind or the other: either *e* definitely *would* have occurred without *c*, or it definitely would *not* have occurred." (ibid). If this were, in fact, the correct analysis, then only in the former case would *c* be a cause of *e*, and Lewis's original notion of causal dependence would continue to serve.[6]

However, as Lewis rightly notes, this objection "solves" the problem of indeterministic causation only by assuming it away:

> I suspect that my opponent is someone who has not wholeheartedly accepted my stipulation of the case in question. (…) [H]e remains at least somewhat inclined to think that the case involves no genuine chance, but a kind of counterfeit chance that is compatible with determinism. (…) If it *is* a case of counterfeit chance, then his objection is well taken. For then *e* is after all predetermined one way or the other, both in actuality and in the counterfactual situation without *c*; but predetermined partly by details of prior historical fact that are far too minute to be discovered in advance. (…) That is all very well, but then his objection is off target. I was not speaking of a case of counterfeit

[6] It is worth noting that this is essentially the approach that the American tort law takes to cases like *Cancer*. For two representative discussions, see Robinson (1985) and Porat and Stein (2003). For a rare critique of the law's assumption that the underlying probabilities in "indeterministic" cases are merely epistemic and not objective, see Parascondola (1996).

> chance, I insist, but of a different case: probabilistic causation of a genuine chance event. (Lewis 1986a, 183)

Short of denying that there exist objectively chancy events in nature, Lewis thinks we are left with what might be called a "pooling conception" of chancy causation. Various causes c_1 to c_n increase the probability of an event e's occurring by contributing to a "probability pool". Once the probability of e is determined, the dice are cast, so to speak, and the event either occurs or it does not. The individual causes make no *additional* contribution to the outcome; they bring e about only *via* their contribution to the probability pool.

Against the backdrop of this pooling conception, Lewis proposes to solve the difficulty raised by *Cancer* and similar cases by relaxing his notion of causal dependence:

> We cannot quite say that without the cause, the effect would not have occurred; but we can say that without the cause, the effect would have been *very much less probable* than it actually was. In this case also, I think we should say that e depends causally on c, and that c is a cause of e. (Lewis 1986a, 176; emphasis mine).

Lewis is explicit that "much less probable" is to be understood in *relative*, not in *absolute* terms:

> "Much less" means less by a large factor – not by a large difference. If x [e's chance of occurring, if c occurs] is already small, the difference of y [e's chance of occurring, had y not occurred] and x could not be large. It is x that sets the standard for how small the chance of e must be without c. We could have one case in which the absence of a cause would lower the chance of an effect from 100% to 10%; another in which the lowering would be from 10% to 1%; yet another in which the lowering would be from 1% to 0.1%; … and all would count equally as cases of chancy causal dependence. So it will not do to simplify our counterfactual and say that without c, the chance of e would be low *simpliciter*. (Lewis 1986a, 177)

This, then, is the settled account of chancy causation in Lewis (1986a): An actual event c is a cause of an actual event e if and only if (i) e causally depends on c, or (ii) there is a chain of events d_1, d_2, ... d_n, such that such that d_1 causally depends on c, d_2 causally depends on d_1, ..., and e causally depends on d_n, where e causally depends on c if and only if, had c not occurred, e's probability of occurring would have been *smaller by a large factor* than its actual chance of occurring.

It has been understood for some time that Lewis's account of chancy causation confronts a worry about false positives. As Christopher Hitchcock (2004) and others have pointed out, there are instances of causation involving multiple probability raising events that do not seem to conform to Lewis's pooling conception.[7] These are cases in which each of a number of events c_1 to c_n significantly raise the probability of some effect e occurring, but intuitively some of c_1 to c_n are *not* a cause of e. Consider the following example from Hitchcock:

> *Atoms*: "A box contains two carbon-14 atoms, and a very sensitive detector is placed near the box. The presence of each atom increases the probability that a decay event will be detected during some given time interval. In fact, a decay event is detected, and one of the atoms in the box is now a nitrogen atom." (Hitchcock 2004, 415)

The presence of either carbon atom increases the probability of a decay event being detected, but intuitively only the presence of one of the atoms is a cause of the decay event being detected. That is, there is a causal chain running from the presence of one of the atoms to the decay event being detected, which is absent in the case of the other atom. If Hitchcock is right, Lewis's theory of chancy causation, which would count the presence of both atoms as a cause of the decay event being detected, thus confronts a problem of false positives.

Rather than totally invalidating Lewis's theory, however, I believe that Hitchcock's counter-example should lead us to restrict its scope. As Hitchcock himself notes, there are other cases where the pooling model appears entirely appropriate:

[7] See also Glynn (2011).

> *Lottery*: Desmond and James decide to purchase two lottery tickets at a price of $1 each, and to split their winnings, if they are lucky. Instead of each individually purchasing and owning a ticket, they walk to the store together, slap their dollars down simultaneously, and purchase two tickets, each registered in both their names. In fact, one of those tickets wins, and the two men share a million dollars (after Hitchcock, 2004; my paraphrase)

Both Desmond and James's causal input (contributing a dollar towards buying lottery tickets) raised the probability of their winning a million dollars. But, although only one of the *tickets* won, it is false to say that it was only Desmond's dollar or only James's dollar that was a cause of their winning. Rather, *both* men's dollars were causally efficacious, by providing the funds with which the winning ticket was purchased. As an analysis of cases like *Lottery*, Lewis's theory of chancy causation thus appears wholly appropriate. And although I shall not here propose, on Lewis's behalf, a way of restricting the scope of his theory to cases like *Lottery* and excluding problem cases like *Atoms*, this does not strike me as an insuperable problem. As long as it will be admitted, therefore, that the problem cases of my own subsequent discussion are more like *Lottery* than like *Atoms*, and would thus pose difficulties even for a version of Lewis's theory that had been amended to meet Hitchcock's challenge, my critique still has a target.

In what follows, I shall present a series of problem cases which demonstrate that, in addition to the problem of false positives identified by Hitchcock, Lewis's account of chancy causation also confronts a problem of false *negatives*. His theory fails to classify as causes some events which, intuitively, *are* causes.

It is an immediate corollary of Lewis's account of chancy causation that (i) if event *e* would have been only *somewhat*, but not *very much* less likely (i.e. not less likely by a large factor), had event *c* not occurred, and (ii) there is no chain of intermediate events d_1, d_2, ... d_n such that, had d_n not occurred, *e* would have been very much less likely to occur, had d_{n-1} not occurred, d_n would have been very much less likely to occur, and so on back to *c*, then *c* cannot be a cause of *e*. Call this feature of Lewis's view the "Restriction". In cases not involving preemption, the Restriction serves, in effect, to filter from the pool of candidate causes all but those events that make a given effect very much more likely to

occur.[8] For instance, if prior to *c*, *e*'s likelihood of occurring was already very high, the Restriction implies that *c* cannot be a cause of *e*. For *c* can no longer *greatly* raise the probability of *e's* occurring.

It is tempting to view the Restriction as just a minor detail of Lewis's account of chancy causation. But this would be a mistake. It is only thanks to the Restriction that Lewis can purport to offer a *unified* account of causation across both deterministic and indeterministic contexts, in terms of the notion of causal *dependence*. We saw that, in order to handle cases like *Cancer*, Lewis must relax his notion of causal dependence. In *Cancer*, we cannot quite say that, had *c* not occurred, *e* would not have occurred. But we can say something close enough: Had *c* not occurred, *e*'s probability of occurring would have been *very much lower*. That, Lewis thinks, is sufficient to say that *e* causally depends on *c*, and that *c* is a cause of *e.* By contrast, to claim that *e* causally "depended" on *c* if, had *c* not occurred, *e* would have nonetheless been *very likely* to occur, or if *e*'s likelihood of occurring was only *insignificantly* raised by *c*, would stretch the notion of causal dependence beyond breaking point. The plausibility of Lewis's causal dependence approach, applied to the sphere of chancy causation, therefore stands and falls with the plausibility of the Restriction.

Unfortunately for Lewis, there are a number of situations in which the Restriction generates implications that are hard to accept. We can see these problems particularly clearly by focusing on causal statements involving a special class of events – what I call "aggregate events".

3. What Are Aggregate Events?

Let us call an *aggregate event* any event *E* that has other events e_1, e_2, … e_n as its constitutive parts, and consists entirely in these events. Call e_1, e_2, … e_n the *constituent events* of *E*.[9]

---

[8] None of the cases that will concern us in the following involve instances of preemption, so it is clause (i) of the Restriction that will concern us. I take this clarification as read in what follows.

[9] Henceforth, in contexts where there are both aggregate events and their constituent events, let capital letters denote aggregate events and lower-case letters constituent events. My definition does not rule out that the constituent events e_1, e_2, … e_n which compose an aggregate event *E* can themselves be aggregate events. Denoting a given event by a capital or a lower-case letter indicates whether, in this context, the event is being considered *qua* aggregate or *qua* constituent event.

Consider the following examples of aggregate events and their constituents:

(a) The string quartet's playing a C major chord at t_1, which consists of the cellist playing a C, the violist playing an E, the second violinist playing a G, and the first violinist playing a C, all at t_1.
(b) The ongoing cholera epidemic in Yemen, which consists of approximately 540,000 Yemenis having become infected with cholera to date.
(c) The scoring of at least two goals in Sunday's match between FC Barcelona and Real Betis de Sevilla, which consisted of the own-goal scored by Alin Tosca of Real Betis in the 36th minute and the goal scored by Sergi Roberto of Barcelona in the 39th minute.

Aggregate events are not simply *mereological sums* of events. For one thing, not all mereological sums of events need *themselves* be events. They may fail to meet whatever constraints our theory of events imposes on this ontological category, for instance that they occur in a spatiotemporal region that is not too widely scattered.[10] For another, while mereological sums have their parts essentially[11], it is plausible that at least some aggregate events do not have (all) their parts essentially. Thus, while the aggregate event of at least two goals being scored in Sunday's match between Barcelona and Real Betis actually occurred in the manner described in example (c), it is reasonable to think that a different *version* of the numerically identical event could have been constituted in a different manner, say by Lionel Messi of Barcelona scoring two goals in the 50th and 73rd minute. This thought is even more compelling in example (b): it is implausible that the numerical identity of an epidemic should depend on exactly which people become infected, nor indeed on their precise number.[12] However, while I find this claim about the accidental mereology of some aggregate events to be plausible, nothing in my subsequent argument will depend on it being correct.

[10] See Lewis (1986b, 243).

[11] See Johnston (2007, 666).

[12] The view that events need not have all their parts essentially is defended in Smith (1983).

In general, since aggregate events are themselves events, we can ask many of the same questions about aggregate events that metaphysicians have debated regarding events *simpliciter*: How fine-grainedly are they to be individuated with regard to the time and manner of their occurring? How modally robust are they? Does the arrangement of their parts matter? Etc.[13] In what follows, I will remain agnostic about these questions. The problems that Lewis's account of chancy causation has with aggregate events do not depend on any particular theory of events being correct.

Although, to date, the concept of aggregate event has not received much detailed attention from philosophers, it is nonetheless an important adequacy condition for a theory of event causation that it give a satisfactory treatment of aggregate events. The notion of aggregate event is central to causal reasoning in many of the social and natural sciences. Consider the kinds of causal claims frequently advanced by macro-economists or economic historians: "The central bank's loosening of interest rates led to increased investment activity and a fall in unemployment." "The Great Depression was caused by irrational exuberance in the stock market, followed by a shortfall in aggregate demand." Or take a typical causal claim from the field of epidemiology: "In 2012, air pollution was a cause of 7 million deaths worldwide." These statements all traffic in causal claims about aggregate events, either as causes or effects or both. In general, when a lack of data, or the incompleteness or computational complexity of our causal models, make it impossible to establish clear causal links at the fine-grained level of individual (constituent) events, we can nonetheless often make informative causal statements at the aggregate level, using techniques like regression analysis, etc.[14] Thus, to give another example, climate scientists typically shy away from attributing particular storms, such as Hurricane Harvey, to global climate change. What they do assert is that climate change is leading to a greater *incidence* of such extreme weather events.[15]

---

[13] For some influential answers to these questions, see Davidson (1980); Kim (1973); Kim (1980); Lewis (1986b); Bennett (1998); and Johnston (2007).

[14] See Morgan and Winship (2016).

[15] This does not mean, of course, that climate scientists do not believe that, for some extreme weather events, global climate change was a cause. In fact, climate scientists are currently working to develop causal models that attribute *individual* climate events to global climate change. For a progress report, see Diffenbaugh et al. (2017). The source of climate scientist's

Given its centrality to causal reasoning in the social and natural sciences, a philosophical account of causation is considerably less attractive if it fails to give a plausible account of causal statements involving aggregate events.

4. The Problem of Aggregate Effects

What, then, are the problems that aggregate events pose for Lewis's account of chancy causation? The first concerns cases where aggregate events are *effects*.

With regard to the causation of aggregate events, the following principle is intuitively compelling:

> **The Principle of Aggregate Effects (PAE):** If c is a cause of an aggregate event E that consists of constituent events $e_1, e_2, \ldots e_n$, c must be a cause of at least of some of $e_1, e_2, \ldots e_n$.

For example, if an outbreak of cholera (c) is a cause of the aggregate effect of 2000 extra deaths in a population (E), this must be because cholera is a cause of the deaths of at least 2000 individuals in that population (e_1 to e_{2000}). These constituent events, after all, are what the aggregate event consists in. How could something be a cause of an aggregate event without, *ipso facto*, being a cause of (some of) that which *makes up* the aggregate event?

The trouble for Lewis is that in certain instances the Restriction makes his account of chancy causation incompatible with the PAE. Consider the following case:

> *901,000 Explosions*: 1 million bombs are each connected to a randomizing detonator that will either trigger, or fail to trigger, the bomb at noon. Initially, the randomizing detonator on each bomb is set such that each bomb has an objective 0.9 chance of exploding at noon.[16] At 9am, Amy pushes a button that turns the

current reluctance to attribute specific extreme weather events to climate change is epistemic, not metaphysical.

[16] In this and all future cases, assume that the objective probabilities are independent of one another.

> dial on each bomb's detonator up by a tiny amount, to 0.901. At noon, 901,000 bombs go off.

It ought to be uncontroversial that, by pushing the button, Amy is a cause of the aggregate event of there being at least 901,000 explosions at noon. So much, at least, is agreed by Lewis's account: Amy's pushing the button is a cause of at least 901,000 bombs exploding, since it raised the probability of this aggregate event by a large factor. Given that Amy pushed the button, the chance of there being at least 901,000 explosions was slightly greater than 0.5. By contrast, had Amy not pushed the button, this probability would have been less than 0.0005. Indeed, the probability would have been greater than 0.99 that there would be no more than 900,700 explosions.

But ask yourself "of which *particular* explosion was Amy's action a cause?", and it seems that Lewis must answer: "of *none*". For prior to Amy's pushing the button, every bomb already had a 0.9 chance of exploding, and Amy's action raised that probability by only a tiny amount, to 0.901. Amy's pushing the button therefore didn't make it significantly more likely that any *particular* bomb would explode. And hence the Restriction implies that it couldn't have been a cause of any particular explosion. In other words: Amy's pushing the button is a cause of an aggregate event (there being at least 901,000 explosions at noon), without being a cause of *any* of the constituent events that make up the aggregate event. This contradicts the PAE.

Forced to choose between the Restriction and the PAE, we should jettison the former and retain the latter. It is hard to make sense of how c could be a cause of an aggregate event E without being a cause of *any* of the events which together *constitute* E. How *else*, we might wonder, could c be a cause of E?

Conversely, we might ask: if, in the *901,000 Explosions* case, the *only* cause of each individual bomb's exploding was that its detonator was antecedently set to 0.9, then how come there were roughly 1,000 more explosions than we would have expected, had Amy had not intervened? Given that the two scenarios are said to be indistinguishable on the level of "micro causation", it is mysterious what further factor could explain the difference in macro outcomes.[17]

[17] In addition to these principled metaphysical reasons for wishing to preserve the PAE, there are also pragmatic considerations relating to moral or legal responsibility that militate against giving up this principle. Suppose that each of the explosions in *901,000 Explosions* killed one person, and that the families of the victims attempted to sue Amy for her role in the carnage. If Lewis's account of chancy causation came to inform legal decision making, the prosecution might have a

5. The Problem of Aggregate Causes

But it gets worse for Lewis. Aggregate events cause trouble for his account of chancy causation not just when they are effects but also when they are causes. Consider the following case:

> *999 Bombers*: A single bomb is connected to a randomizing detonator that will either trigger, or fail to trigger, the bomb at noon. Initially, the randomizing detonator is set to 0%. At 9am, simultaneously, 999 technicians each push a button. The effect of each technician's button-pushing is to raise the likelihood that the bomb will go off by 0.001. The *cumulative* effect of their actions is thus to raise the probability that the bomb will go off to 0.999. At noon the bomb detonates.

It is intuitively compelling that, *collectively*, the actions of the 999 technicians are a cause of the explosion. In other words, it is compelling that the aggregate event *C*, consisting of the actions of each of the 999 technicians at 9am, is a cause of *e,* the explosion at noon. After all, if not *C*, what else could account for the explosion? [18] It will be a strike against a theory of causation if it cannot accommodate this intuitive datum.

Can Lewis's account accommodate it? Is it the case that, had *C* not occurred, *e*'s likelihood of occurring would have been very much smaller? The answer is not straightforward. It depends on what exactly would have to be the case for *C* not to occur, which in turn is a function of *C*'s conditions of cross-world identity.

---

surprisingly hard time achieving the severe verdict that surely appears appropriate in this case: It is standardly assumed that, outside contexts of strict or vicarious liability, in order to be legally accountable for a criminal or tortious wrong *W*, an agent must have some degree of causal responsibility for *W*. "Now, granted", we can imagine Amy's defense attorney arguing in a Lewisian spirit, "Amy's action may have caused the *level* of deaths in *901,000 Explosions* to be higher than it otherwise would have been. But why should this matter, given that Amy's action didn't actually kill *anyone*, nor was even a *cause* of anyone's death?" If Lewis were right, Amy's action at most constituted a victimless crime. That seems wrong.

[18] This question is especially pressing for Lewis since, as we shall see, his theory implies that none of the *constituent* events of *C* is a cause of *e*.

Suppose, for instance, that the *number* of button-pushings is essential to the aggregate event *C*. This would make *C* fairly modally fragile. Had there been a different number of button-pushings at noon, *C* would not have occurred. A numerically different, though similar aggregate event would have occurred instead. But, in that case, *e* does not causally depend on *C*. For the nearest possible world in which *C* does not occur, i.e. the one that departs least from actuality, is not one in which none or only a few technicians push their buttons. Rather, it is one in which 998 or 1000 technicians push their buttons. This is not a world in which the likelihood of there being an explosion would have been smaller *by a large factor*. Instead, it would either have been only a little bit smaller, or indeed *higher*.

Of course, a case could also be made for a different way of individuating *C* – one on which *e does* counterfactually depend on *C*. Perhaps *C* occurs in all worlds in which *any* of the technicians push their buttons and is only *accidentally* constituted by 999 button-pushings in the actual world. (We mustn't confuse the description with which we pick out an event with features of those events that are essential to its occurrence.) Alternatively, we could retain the modally fragile way of individuating *C*, but specify that the possible world we should consider in determining whether *e* causally depends on *C* is not one in which an event very similar to *C* occurs, but rather one in which no event like *C* occurs. Lewis himself appears to have been sympathetic to this second kind of solution. In "Causation as Influence" he writes:

> What it means to suppose counterfactually that *C* does not occur depends on where we draw the line between *C*'s not occurring and *C*'s occurring differently in time and manner.
>
> That makes a problem. What is the closest way to actuality for *C* not to occur? It is for *C* to be replaced by a very similar event, one that is almost but not quite *C*, one that is just barely over the border between versions of *C* itself and its nearest alternatives. But if *C* is taken to be fairly fragile, then, if almost-*C* occurred instead of *C*, very likely the effects of almost-*C* would be almost the same as the effects of *C*. So our causal counterfactual will not mean what we thought it meant, and it may well not have the truth value we thought it had.
>
> When asked to suppose counterfactually that *C* does not occur, we do not really look for the very closest possible world where

> *C*'s conditions of occurrence are not quite satisfied. Rather, we imagine that *C* is completely and cleanly excised from history, leaving behind no fragment or approximation of itself.[19]

On either of these proposals it would be true that, in the possible world in which *C* does not occur (and no event *like C* occurs), *e*'s chance of occurring would have been very much lower.

However, while this may help Lewis to accommodate the intuitive datum that the aggregate event *C* is a cause of *e*, it pushes him into the arms of a second problem. Whatever the truth about *C*'s essential properties may be, in the actual world *C* is constituted by c_1, c_2, … c_{999}. And the following claim about aggregate events as causes is very plausible:

> **The Principle of Aggregate Causes (PAC):** If an aggregate event *C*, consisting of constituent events c_1, c_2, … c_n, is a cause of event *e*, then some of c_1, c_2, … c_n must be a cause of *e*.

The PAC strikes me as almost as compelling as the PAE. If c_1, c_2, … c_n *together* are a cause of *e*, in the form of the aggregate event *C*, then surely some of c_1, c_2, … c_n must have a claim to being regarded as causes of *e*.[20] How could an aggregate event be a cause of *e* if *none* of the constituent events that entirely make up the aggregate event are causes of *e*?[21]

---

[19] Lewis (2000, 190). See also Lewis (1973a, 210-211) for an earlier statement of the same point.

[20] In the present case, since all constituent events stand in the same relationship to *C*, it would be arbitrary to identify some as causes of *e*, but not others. Thus, *all* the constituent events in this case have a claim to be regarded as causes of *e*.

[21] It might be objected that, despite its intuitive plausibility, we shouldn't be all that surprised by violations of the Principle of Aggregate Causes. After all, the objection continues, analogous principles for other types of aggregate entities are false: The fact that an army (an aggregate entity composed of individual soldiers) surrounds the fortress does not entail that any individual soldier surrounds the fortress.

This objection fails, because the analogy is inexact. "The army surrounds the fortress" is most naturally read as "The army, *by itself*, *fully* surrounds the fortress." The corresponding claim about aggregate events would be that an aggregate event *C* is *the cause* of some effect *e* (suggesting that *e* has no other causes). If "is the cause of *e*" were the relevant relational predicate, the PAC would obviously be false. But this is not what the PAC says, nor is it the notion that Lewis himself is concerned with. He explicitly states that what he aims to analyze is the concept of being *a* cause, not *the* cause, of an event:

Once again, however, this intuitive principle comes into conflict with the Restriction. Consider what would have been the case if any given technician had failed to press his button, i.e. if any of the constituent events that *C* consists in had failed to occur: the probability that the bomb explodes would have been only very slightly lower – 0.998 instead of 0.999. For the closest possible world in which a given constituent event c_i does not occur is one in which the others *do* occur. Hence, according to Lewis, it isn't true of the action of any of the 999 technicians that it was a cause of the bomb's exploding.

Cases like *999 Bombers* thus confront Lewis's account of chancy causation with a dilemma: Either the aggregate event *C* is individuated in such a way that *C* is not a cause of *e*, given Lewis's account of causation as causal dependence. This would avoid violating the PAC. But it is a highly undesirable result nonetheless, for two reasons: For one thing, it is intuitively compelling that *C is* a cause of *e*. For another, if *C* is not a cause of *e*, and neither are any of c_1 to $c_{999,}$ according to Lewis, then it seems like *e* must count as an *uncaused* event, by Lewis's lights. This is not an acceptable result.

Alternatively, we accommodate the intuition that *C* is a cause of *e*, via one of the methods surveyed above. But then Lewis's account of causation implies that the Principle of Aggregate Causes is violated, since, according to Lewis, none of the constituent events that make up *C* are a cause of *e*.

I note in passing that the problems that Lewis's theory of causation has with the Principle of Aggregate Causes are not limited to instances of chancy causation. Structurally parallel problems also arise under determinism, namely in cases of symmetrical *overdetermination*. Figure 1 provides an illustration:[22]

> We sometimes single out one among all the causes of some event and call it "the cause", as if there were no others. (…) I have nothing to say about these principles of invidious discrimination. I am concerned with the prior question of what it is to be one of the causes (unselectively speaking)." (Lewis 1986a, 558-59)

The natural analogue of "is a cause of *e*" is not "surrounds *f*" but "helps to surround *f*". Plugging this back into the example, a truer analogue of the PAC would therefore read: "If an army helps to surround a fortress [perhaps as part of a larger siege by multiple armies], then it must be that some of the soldiers who compose the army help to surround the fortress." And this claim is, in fact, perfectly true. I am indebted to [Name Redacted] for discussion of these points, as well as to Kaiserman (2016).

[22] Based on an example in Lewis (1986a, Appendix E).

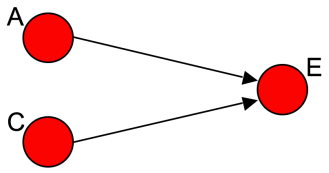

Figure 1

Let A, C, and E in the diagram represent three neurons, and let *a, c,* and *e* represent the events that correspond to the firing of these neurons. Neuron E will fire if stimulated by either neuron A or C, i.e. if either *a* or *c* occur, or if both *a* and *c* occur. If neither A or C fire, E will not fire. To sidestep questions of event-individuation, we stipulate that E will fire in an *identical* manner, whether both A and C fire or only one of these neurons fires. It is therefore beyond doubt that we are talking about the numerically identical event *e* in all scenarios where E fires.[23]

In our example, both *a* and *c* occur, and *e* occurs. Are either *a* or *c* causes of *e*? Lewis's counterfactual analysis implies that they are not: *e* does not causally depend on either. If nothing more could be said, *e* would have to count as an uncaused event – an unacceptable result. This corresponds to the first horn of the dilemma presented above.

Lewis is alive to this problem. He attempts to solve it by claiming that, while neither *a* or *c* are causes of *e*, the *aggregate* event *C* composed of *a* and *c* is:

> I should dispel one worry: that if we ever decline to count redundant causes as genuine causes, then we will be left with gaps in our causal histories – no cause at all, at the time when the redundant causes occur, for a redundantly caused event. That is not a problem. For consider the larger event composed of the two redundant causes. (…). Whether or not the redundant causes themselves are genuine causes, this larger event will be there to cause the effect. For without it – if it were completely

[23] By contrast, if neuron E fired significantly differently (more vigorously, say) if it was stimulated by both A and C, it could be argued that the case in which all three of *a, c,* and *e* occur would be one of *joint causation* of *e* by *a* and *c*, not a genuine instance of causal *overdetermination.* I am indebted to Hall and Paul (2013) for this observation.

> absent, with neither of its parts still present, and not replaced by some barely different event – the effect would not occur. For *ex hypothesi* the effect would not occur if both redundant causes were absent, and to suppose away both of them is just the same as to suppose away the larger event that is composed of them. (Lewis 1986a, 212)

This is structurally parallel to the proposal we surveyed in *999 Bombers*, and encounters the same problems that formed the second horn of the dilemma. For one thing, Lewis must rely on the controversial claim that the relevant counterfactual to look at, when considering what happens when the aggregate event *C* does not occur, is one in which neither *a* nor *c* occurs. It isn't altogether clear why we should grant this, given that the closest alternative possible world to the one in which *C* occurs arguably isn't one in which *neither a* nor *c* occur, but one in which only one of these events does not occur. For another, we again have a violation of the Principle of Aggregate Causes: *C* is a cause of *e*, but neither *a* or *c*, which together make up *C*, are said to be causes of *e.* This is implausible.

6. From Causal Dependence to Causal Contribution

The Problem of Aggregate Effects and the Problem of Aggregate Causes represent serious stumbling blocks for Lewis's theory of chancy causation. Moreover, the two problems are logically independent of one another. So even if, for instance, you aren't completely persuaded by the Problem of Aggregate Causes, because you don't find the PAC as compelling as I've made it out to be, you would still have the other problem to contend with, and vice versa.

Fortunately, a straightforward solution suggests itself that solves both problems in one fell swoop, namely to jettison the Restriction. It is the Restriction which prevents us from saying that Amy's pushing the button is not just a cause of the aggregate effect of there being at least 901,000 explosions, but also of the individual explosions that make up this effect; and it is the Restriction which prevents us from saying that each individual button-pushing by a technician is a cause of the explosion in *999 Bombers*, not just the aggregate event of 999 technicians pressing a button. And it is our inability to say these

things, which, in turn, leads to violations of the PAE and the PAC. Once we drop the Restriction, these problems go away.

But while jettisoning the Restriction would represent only a small emendation to the *letter* of Lewis's account, it is an emendation which, as we noted in Section 2, Lewis *cannot make* without giving up his ambition of providing an account of causation, for deterministic and indeterministic contexts, in terms of a unified notion of causal *dependence*. To say that, in *901,000 Explosions*, each individual explosion causally "depended" for its occurrence on Amy's pushing the button would stretch this notion beyond recognition. After all, had Amy not pushed the button, the explosion would still have had a 0.9 chance of occurring, and Amy raised the likelihood of explosion only very insignificantly. To speak of causal "dependence" in such a case is not just to *relax* the simple counterfactual analysis of dependence that Lewis appeals to in deterministic cases; it is to give the term "causal dependence" an altogether different meaning. Hence, while on one level giving up the Restriction represents only a minor modification to Lewis's account of chancy causation, it would very significantly alter its analytic thrust.

A better label for the new paradigm is that of causation as *causal contribution*. 'Causal dependence' is a binary concept. Either *e* causally depends on *c* or it does not. It makes no sense to say that *e* causally depended on *c* "a little bit", or "a lot". By contrast, 'causal contribution' is a *scalar* notion. There can be greater or smaller causal contributions to an event's occurring. Consequently, the fact that *c* raised *e's* probability of occurring by only a small degree is no obstacle to considering *c* a cause of *e*, under a conception of causation as causal contribution. I submit that in cases of indeterministic processes where the pooling conception is appropriate[24], an actual event *c* causally contributes to the occurrence of an actual event just in case, had *c* not occurred, *e* would have been less likely to occur. The degree to which *e* would have been less likely to occur is a measure of the size of *c*'s causal contribution.

A theory of chancy causation based not on causal dependence but on causal contribution would then look something like this: An actual event *c* is a cause of an actual event *e* if and only if (i) *c* causally contributes to *e*, or (ii) there is a chain of events d_1, d_2, ... d_n, such that *c* causally contributes to d_1, d_1 causally contributes to d_2, ..., and d_n causally contributes to *e*. This simple account is no

[24] See the discussion in Section 2, above.

doubt in need of much refinement.[25] But for our purposes it shall suffice. The theory gives plausible verdicts in all the cases we have so far considered.

We are not yet home free, however. [Name redacted] has presented me with an ingenious case that appears to challenge the PAC even on a causal contribution model of chancy causation. If [Name redacted]'s objection goes through, then the theoretical gains of switching from a causal dependence to a causal contribution model of chancy causation are more modest than I had hoped, since the PAC would ultimately have to be abandoned on either account. Here, then, is [Name redacted]'s case[26]:

> *The Two Assassins*: Unbeknownst to each other, two assassins have been assigned the same target. Both assassins attempt to dispose of their victim by poisoning her. The first assassin administers poison 1 (c_1) and the other poison 2 (c_2). Both poisons are highly lethal. Indeed, had the victim been exposed to either type of poison in isolation, she would have had a 0.99 objective chance of dying (whereas we assume that without either poison being administered, the victim had no chance of dying). As it happens, however, poison 1 and poison 2 are highly effective antidotes to one other. Having been administered *both* poisons, the victim has only a 0.10 chance of dying. Unfortunately, even small chances sometimes come to pass, and in this case the victim ends up dying (e).[27]

It should be uncontroversial that, as some level, the administration of both poisons was involved in the victim's death. After all, without either poison being administered, the victim had no chance of dying, and c_1 and c_2 are symmetrically related to e, so that if one is a cause of e, so is the other. However, even a causal contribution model of chancy causation seems to imply that neither c_1 nor c_2 was a cause of the victim's death: after all, it is true of either

[25] For a more worked-out account along the same lines, see Glynn (2011).

[26] Hall and Paul (2013, 65-66) refer to a version of this problem as the "Problem of Ambivalent Causes", but do not propose a solution.

[27] Again, let us bracket distractions due to the individuation of events by assuming that the time and manner of the victim's death was exactly like it would have been, had only one of the poisonings taken place.

poison that, had it not been administered, the victim would have been *more*, not *less* likely to die. Are we then forced to the absurd conclusion that the victim's death was *un*caused? Not quite. We could still resort to the (by now familiar) manoeuver of maintaining that, while neither *individual* poisoning was a cause of the victim's death, the *aggregate event C* of both poisons being administered *was* a cause of death. But then, it seems, we would have a new violation of the PAC, and it was in part a desire to avoid violations of the PAC that motivated a move to the causal contribution model in the first place.

Fortunately, the threat posed by *The Two Assassins* is more apparent than real. Take a closer look at [Name redacted]'s case and in particular ask yourself: *Why* exactly is it that when both poisons are administered together, this does not further raise, but rather lowers the victim's likelihood of dying? The answer, I believe, must be connected to the way in which the two toxins *interact* with one another. We can imagine, for example, that when poison 1 and poison 2 are present in the same body, they chemically react to give rise to a new chemical compound. While still mildly toxic, this new compound is far less dangerous to a person than either poison 1 or poison 2 in isolation. If this is roughly right, then on closer inspection *The Two Assassins* turns out not to be a case in which the victim's death was the *direct* causal result of the aggregate event *C*, constituted by c_1 and c_2. Rather, as Figure 2 illustrates, it is a case in which both c_1 and c_2 causally contribute to bringing about an intermediate event *b*, which in turn causes *e*. David Lewis (1986) refers to such events as "Bunzl events", in honor of Martin Bunzl's discussion in Bunzl (1979).

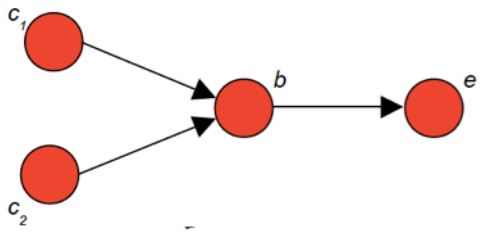

Figure 2

Viewed in this way, [Name redacted]'s case ceases to be a challenge to the compatibility of the causal contribution account and the PAC. The Bunzl event *b* is undoubtedly a cause of *e*, at least on a non-back-tracking reading of

counterfactual conditionals. (At the time when *b* occurs, it is true that, had *b* not occurred, *e* would have been very much less likely to occur). And it is also undeniable that c_1 and c_2 are causes of *b* (*eo ipso*, *C* is a cause of *b*). But then, by the transitivity of the causal relation, c_1 and c_2 are, after all, causes of *e*, although, had either c_1 or c_2 not occurred, *e* would have been more likely to occur.

It is important to understand why a proponent of a causal dependence theory of chancy causation could not similarly appeal to Bunzl events to argue that his account is, after all, compatible with the PAC in cases like *999 Bombers*. Unlike in *Two Assassins*, there is nothing to suggest that in *999 Bombers* there is an intermediate Bunzl event that stands in the causal chain between the individual button-pushings c_1 to c_{999} and the explosion. In particular, the aggregate event *C* is *not* a Bunzl event. It would be a mistake to say, in *999 Bombers*, that c_1 to c_{999} cause *e* by causing *C*. The relationship between c_1 to c_{999} and *C* is not a causal, but a *constitutive* one. It is true that, if a certain number of c_1 to c_{999} did not occur, *C* would not occur. (How many exactly depends on how modally robust we take *C* to be). But the reasons for this are not causal but *logical*. This concludes my response to [Name redacted]'s challenge.

I shall end my paper by mentioning a third class of cases in which Lewis's causal dependence account of chancy causation encounters serious trouble. Once again, these cases underscore the superiority of the causal contribution account, which avoids the implausible implications of Lewis's view.

7. The Problem of Causation and Prevention

In "Causation as Influence", his last article on the topic of causation, Lewis maintains that amongst the effects of an event may be the *absence* of another, "as when a vaccination prevents one from catching a disease."(Lewis 2000, 195). Though Lewis does not go into any details, it is likely that he would have been sympathetic to the following probabilistic account of the *prevention* of events, constructed on analogy with his causal dependence account of chancy causation:

> *c* causes *e* not to occur if and only if *c* occurs and *e* does not occur and *c* lowered *e's* probability of occurring by a large factor.

With this in mind, consider one final case:

Vaccination: Suppose that at some future time *t*, each out of 100 patients can either be healthy or sick. Initially each patient's likelihood of being healthy at *t* is 0.9. By giving each person a vaccine, we can lower their likelihood of being sick at *t* to 0.01, thus raising their likelihood of being healthy to 0.99. We administer the vaccine, and at *t*, exactly 99 patients are healthy.

Figure 3 summarizes the relevant features of this case:

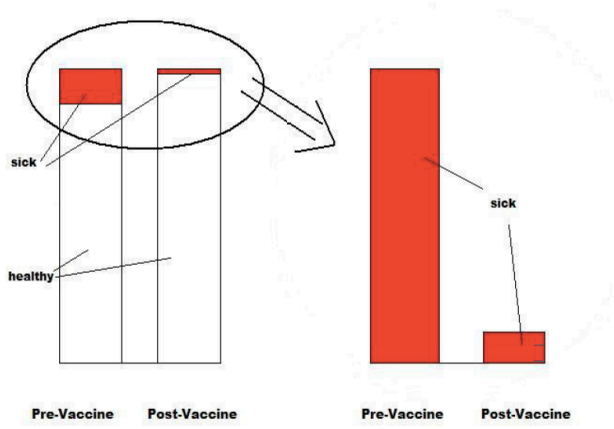

Figure 3

This case is *doubly* problematic for a Lewis-style theory of prevention. Suppose a proponent of the causal dependence account of chancy causation and prevention was asked the following three questions about this case:

(1) Was administering the vaccine to the patients a cause of the aggregate event of at least 99 persons being healthy at *t*?
(2) For any given patient, was administering the vaccine to this patient a cause of her being healthy at *t*?
(3) For any given patient, was administering the vaccine to this patient a cause of her not being sick at *t*?

The first problem we are already familiar with: Lewis's theory, implausibly, answers "yes" to question (1) but "no" to (2). The answer to (1) is affirmative, because the likelihood of at least 99 patients being healthy at *t* increases by a

large factor due to the vaccination. (Given the vaccination, the likelihood of at least 99 patients being healthy at *t* was slightly less than 0.74. By contrast, without the vaccination, this outcome was astonishingly unlikely – roughly 0.0003.). By contrast, the answer to (2) is negative, because for any given patient, her probability of being healthy at *t* increased by only a factor of 0.1. This is the familiar Problem of Aggregate Effects, discussed in Section 4. We have already explored in detail why this combination of answers is implausible.

But there is a second, novel problem. While Lewis's theory answers "no" to question (2), it answers yes to (3). For the vaccine does lower each patient's probability of being sick *by a factor of 10*, from 0.1 to 0.01. Lewis's account thus implies that while the vaccine is not a cause of anyone's being healthy at *t*, it *is* a cause of 99 patients' not being sick at *t*! But this is a nonsensical result. Being sick is simply the logical flipside of being healthy, so that something that is a cause of a patient's not being sick *ipso facto* is a cause of her being healthy.

The causal contribution model avoids this implausible conclusion. Since administering the vaccine to a patient both makes them more likely to be healthy and less likely to be sick at *t*, the model implies that, for the 99 patients who are healthy at *t,* administering the vaccine was a cause, both of their being healthy and of their not being sick. This, surely, is the correct result.

8. Conclusion

In this paper, I have argued against Lewis's causal dependence account of chancy causation. We have seen that by denying that events which only slightly raise the probability of an effect *e* can be amongst the causes, the causal dependence account lands Lewis in serious difficulties. A causal contribution model of chancy causation, on which an actual event *c* is a cause of *e* just in case *c* causally contributes to bringing *e* about by making *e*'s occurrence more likely, avoids the problems I have canvassed. It is possible that such an account faces drawbacks on other fronts. But at least along the dimensions that I have surveyed in this paper, the causal contribution model emerges as the clear victor.