

CAUSING AND CONTRIBUTING

by

Marc E. Johansen

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2013

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Marc Johansen, titled *Causing and Contributing* and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

Carolina Sartorio

Date: 11/26/2013

L. A. Paul

Date: 11/26/2013

Terry Horgan

Date: 11/26/2013

Michael McKenna

Date: 11/26/2013

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director: Carolina Sartorio

Date: 11/26/2013

Dissertation Director: L. A. Paul

Date: 11/26/2013

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: Marc Johansen

DEDICATION

To Kate

TABLE OF CONTENTS

ABSTRACT	6
CHAPTER 1 Introduction	8
CHAPTER 2 Trouble with intimacy	14
2.1 Introduction	14
2.2 The intimacy thesis	16
2.3 The problem of distal effects	19
2.4 More intimacy	23
2.5 Against <i>Co-dependence</i>	27
2.6 Overdetermination	29
2.7 Explaining failed additivity	33
CHAPTER 3 Causal contribution and causal exclusion	37
3.1 Introduction	37
3.2 The causal exclusion problem	39
3.3 Contributions as lawful constraints	45
3.4 Causal contributions and causal additivity	48
3.5 Solving the exclusion problem	52
3.6 Failed additivity in the circumstances	55
3.7 The present view in context	62
CHAPTER 4 Regularity as a form of constraint	68
4.1 Introduction	68
4.2 The problem of epiphenomena and other piggybacking problems	70
4.3 Contribution	75
4.4 Minimal sufficiency two ways	78
4.5 Adapting the regularity theory	82
4.6 Piggybacking problems	84
4.6.1 The piling-on problem	85
4.6.2 Replacing and appending	87
4.6.3 The gerrymandering problem	88
4.6.4 The problem of epiphenomena	89
4.7 Limitations of CR	91
4.8 Contributions reconsidered	93
REFERENCES	98

ABSTRACT

I develop a solution to the causal exclusion problem and explore its implications for a broader metaphysics of causation. In a series of three papers, I show that standard approaches to the exclusion problem are inadequate, that the solution to the problem lies not in causation but in the overlooked phenomena of contribution, and that contribution grounds a new kind of theory of causation.

In “Trouble with intimacy,” I show that the standard solution to the exclusion problem is inadequate. The problem is traditionally framed in terms of causal overdetermination. It charges that (i) if mental events are not identical to their physical realizers, they systematically overdetermine their common effects and (ii) such effects are not overdetermined. Critics often deny (i), claiming that mental events and their realizers are too intimately related to be overdetermining causes. I develop a class of cases that undermine this response. These cases show two things. First, mental events and their realizers overdetermine at least some of their effects. Second, overdetermination is not essential to the exclusion problem.

In “Causal contribution and causal exclusion,” I develop a solution to the exclusion problem. The exclusion problem is a symptom of our failure to attend not just to causation, but to the conceptually more basic notion of contribution — the influence that an event has on future states of the world independently of other events. I develop an account of contribution as a constraint on what world states may obtain in an event’s wake. The solution to the exclusion problem lies in the relation between the contributions of mental events and those of their realizers.

In “Regularity as a form of constraint,” I present the groundwork for a new type of regularity theory of causation. Traditional regularity theories have been

much too liberal: they entail a wealth of causal relationships that do not exist. We can correct this by grounding regularity/entailment relations in contributions. Traditional regularity theories fail because they identify causation with entailment by a non-redundant sufficient condition. This new breed of regularity theory succeeds by identifying causation with entailment via a minimally restrictive contribution.

CHAPTER 1

Introduction

Suppose that we add a drop of blue dye and a drop of yellow dye to a vial of water. The water turns green. Cases like this reveal an incongruity in the way that we've come to think about causation and causal influence. What an event causes depends on its environment. Had we not added yellow dye to the water, the addition of blue dye would not have caused the water to turn green. Yet the influence of an event seems, in at least some respect, not to depend on its environment. The addition of blue dye influences the water's color in fundamentally the same way whether the yellow dye is present or not. Call this influence the event's *contribution*.

Contemporary work in causation has largely overlooked contribution. Philosophical accounts of causation have traditionally divided into three camps. There are those views descended from Lewis (1973) that analyze causation in terms of counterfactual dependence relations. On the simplest version of such a view, an event c causes an event e just in case e counterfactually depends on c : if c had not occurred, e would not have occurred. There are those views that follow Hume and more recently J. L. Mackie (1965; 1974) in analyzing causation in terms of lawful regularities or entailments. Such views claim, roughly, that an event c causes an event e just in case the occurrence of c entails the occurrence of e given the laws and actual background conditions. And finally, there are those views that identify actual causation with the transfer of energy or some other physical connection between events (e.g., Dowe (2000), Fair (1979), Salmon (1998)).

Contribution is none of these things. In so far as ordinary events are concerned, whether one event stands in a dependence relation to another event, or lawfully entails that event, or transfers energy to that event all depend on factors in the environment. But the contribution that an event makes does not. As a result, while causation has received an increasing amount of attention in recent years, contribution has been ignored.

This neglect has allowed a number of otherwise tractable problems to fester in the metaphysics of causation. This dissertation aims to correct that. I develop an account of contribution and apply it in the service of two issues: the causal exclusion problem in the philosophy of mind and the reassessment the oft-maligned regularity theory of causation.

The exclusion problem stems from the tension between (a) the commonly held view that mental events are not identical to the neural events that realize them and (b) the fact that mental events and their realizers do not interact with one another in the manner we've come to expect of different causes. As a rule, unless two causes of the same effect are part of a single causal chain, the effect they cause together will have different properties than what either would have caused on its own. A window struck by two golf balls, for instance, will break apart in a different way than it would if struck by either one of the golf balls alone. Each impact makes a difference to what the other causes. Not so for mental events and their realizers. The cumulative effect of a mental event and its realizer is in no way different from what the realizer itself is sufficient to cause. Should you reach for a cup of coffee, for example, you do not reach more quickly or grasp the cup more firmly because your reaching was caused by both your desire for coffee *and* the physical realizer of that desire. In this respect, the two events act like one cause, not two.

The exclusion problem is commonly held up as an objection to theories of mind that claim that mental events and properties are not identical to physical events and properties. Kim (1993b; 1998; 2005) has argued forcefully that the exclusion problem presents a dilemma for such views: either accept that mental events reduce to physical events or admit that mental events are epiphenomenal, that they are causally inert. Those who wish to escape Kim's dilemma need to find a way out of the causal exclusion problem. They must explain why, if mental events and their realizers are in fact different events, they seem not to act like it.

In "Trouble with intimacy," I argue that solving the exclusion problem is more difficult than is commonly appreciated. Philosophers have often claimed that mental events and their realizers interact differently than most other events because, unlike most events, mental events and their realizers are intimately related. The impacts of the two golf balls against the window, for instance, are metaphysically independent of one another in a way that mental events and their realizers are not. And so the thought goes that mental events and their realizers do not interact with one another in the same manner as the impacts of the two golf balls because the intimate relation between them precludes it. I demonstrate that this proposed solution to the exclusion problem is inadequate. I develop a class of cases in which the interactions of intimately related events are in critical respects no different from the interactions of events that are not intimately related. In such cases, the realizers of mental events make a difference to what those mental events cause.

I argue that there are two important lessons to take from these cases. First, it's simply false that intimate relations preclude mental events and their realizers from interacting in the same fashion as events that are not intimately related. Second, intimate relations could not have provided the right type of distinction necessary

to solve the exclusion problem. Solving the exclusion problem in a manner that avoids Kim's dilemma requires an account of why *some* interactions of mental events and their realizers are unlike those of other events and some are not. No reasonably natural relation between a mental event and its realizer can do that on its own.

In "Causal contribution and causal exclusion," I argue that instead of turning to the close relation of mental events and their realizers to solve the exclusion problem we should look to their contributions. Contributions, I claim, are a kind of constraint that an event imposes on what the world may be like in the future. Given the natural laws, the mere occurrence of an event limits what future states of affairs may obtain. States that are eligible to obtain in the wake of an event must be not only consistent with the laws, but also with the occurrence of that event. I argue that if a pair of events are each to make a difference to what the other causes, their contributions must cross-cut one another: the contribution of each event must preclude the obtaining of some state of affairs that the other permits. The contributions of most events cross-cut one another in this way. The contributions of mental events and their realizers do not. While the contribution of the realizer precludes the obtaining of states that the contribution of the mental event permits, the contribution of the mental event does not preclude the obtaining of any state permitted by the contribution of its realizer.

This relationship between the contributions of mental events and their realizers provides a solution to the exclusion problem. Since the contribution of a mental event does not preclude the obtaining of any state that is not also precluded from obtaining by the contribution of its realizer, mental events cannot make a difference to what their realizers cause. However, since the contribution of the realizer does preclude the obtaining of some states that the contribution of the

mental event permits, the realizer may make a difference to what the mental event causes. Whether or not it does make such a difference will depend on the effect in question and what states are precluded from obtaining.

In “Regularity as a form of constraint,” I argue that regularity theories of causation in the tradition of Hume and J. L. Mackie warrant renewed consideration in light of the account of contribution that I develop in the previous chapter. Mackie’s regularity theory holds that causes are parts of minimally sufficient conditions for their effects. In essence, Mackie claims that an event c causes an event e just in case c is part of a condition C such that, given the laws, (i) C is sufficient for e and (ii) no proper part of C is also sufficient for e . This condition is famously too liberal; it routinely attributes causal relations where there are none to be found. This problem has led many philosophers to reject regularity theories like Mackie’s as a dead end.

My account of contribution provides a solution to the problem. Mackie’s minimalism is ultimately a minimalism about syntax. For Mackie, a condition is minimally sufficient for an effect only if it does not include any redundant parts; remove any of its parts and the condition is no longer sufficient for that effect. I argue that we can repair the regularity theory by trading Mackie’s syntactically minimal sufficient conditions for sufficient conditions that are minimal with respect to the contributions of their constituent events. A set of events is sufficient for some outcome if the events’ contributions collectively constrain future states of the world such that they assure that the outcome must occur. Those events are minimally sufficient for that outcome just in case their contributions collectively constrain future states of the world enough to assure that this outcome occurs, without doing so any more than is necessary to assure that it occurs. So understood, an event causes an effect only if it belongs to a set of events whose

contributions, collectively, are minimally sufficient for that effect.

CHAPTER 2

Trouble with intimacy

2.1 Introduction

Nonreductionists have a complicated relationship with intimacy. A nonreductive ontology must claim that some entities stand in intimate relations of ontic dependence. Intimately related entities are not identical. But neither are they wholly distinct; their existences, or occurrences, or instantiations are (metaphysically) necessarily correlated. It is this intimacy claim that distinguishes nonreductive ontologies from their reductive and pluralistic counterparts. Contra reductionism, some entities A and B are not so closely related that they are identical. Contra pluralism, A and B are too closely related to be wholly distinct from one another either. The connection between them is not just a contingent matter of law.

Nonreductive ontologies of mental events are no exception. Nonreductionism about mental events (henceforth ‘nonreductionism’) holds that while mental events are not identical to any physical events, they are intimately related to certain physical events. For example, suppose that I come to believe that my coffee has gone cold. Nonreductionism entails that my present holding of that belief (a mental event) is not itself a physical event, but is intimately related to a physical event of some sort that realizes my belief. The motivation for saying as much is familiar. We want to allow that mental events are multiply-realizable without embracing full-on psychophysical dualism.

Intimacy, however, is often put to a further end. It’s not uncommon to hear

friends of nonreductionism claim that intimate relations are causally significant. Intimacy, they say, limits the kinds of causal interactions that a pair of events may engage in. My having the belief that the coffee is cold can interact with my having the belief that cold coffee is undrinkable in ways that it cannot interact with an intimately related physical event that realizes it.

Nonreductionists' intentions in this regard are not altogether pure. They have (typically, though perhaps not exclusively) claimed this role for intimate relations in response to the causal exclusion problem. The problem stems from the fact that if nonreductionism is correct that mental events are not identical to their physical realizers, ordinary instances of mental causation would seem to amount to causal overdetermination: cases in which a single effect has multiple sufficient causes, such as the shattering of a fragile window by the simultaneous impact of two hailstones or an execution by firing squad in which every shot is fatal. For we typically wish to say that the effects of mental events are also caused by physical events in the brain.¹ To the nonreductionist, such talk is not just benign double-counting; it is an admission that the effects of mental causation are caused twice over, once by a mental event and once by a different physical event. Proponents of the exclusion problem find this result objectionable. Mental events and their realizers, they claim, do not systematically overdetermine their effects.

My concern in this paper is with one such strategy for resisting the causal

¹This claim is typically motivated by the physical causal closure principle, the claim that if a physical event has a cause at a given time, it has a sufficient physical cause at that time. The closure principle is frequently defended on the grounds that it is a minimal requirement for physicalism (Kim, 1998, 2005) or that it is a guiding assumption of best scientific practice (Papineau, 2001, 2002). For present purposes however, nothing rides on accepting physical causal closure. It will be enough to note that, as a matter of fact, our best theories of causation frequently seem to entail that an effect is caused both by a mental event and its realizer.

exclusion problem. Nonreductionists often claim to escape the exclusion problem on the grounds that mental events and their realizers, though different events, cannot overdetermine their common effects.² Overdetermination of the sort found in textbook cases requires events that are metaphysically independent of one another in a way that mental events and their realizers are not. So construed, the exclusion problem fails against the following principle:

Intimacy For any events C and C and any effect E , C and C overdetermine E only if C and C are not intimately related.

Intimacy is an attractive thesis. It promises to solve the exclusion problem on the cheap; accepting it does not carry significant metaphysical commitments beyond that of a broadly nonreductive ontology of mental events. Unfortunately, *Intimacy* does not succeed. The trouble is two-fold. In the first, *Intimacy* is simply false. Worse, however, *Intimacy* would not resolve the exclusion problem if it were true.

2.2 The intimacy thesis

If *Intimacy* is correct, intimate relations limit the kinds of interactions that their relata may have. My decision to get a fresh cup of coffee and the physical event that realizes that decision both cause me to get up from my desk. The impact of the baseball that Billy hit and the impact of the baseball that Suzy hit both cause the fragile window to shatter. *Intimacy* entails that these are different kinds of interactions, that my decision and its realizer are by their very nature unable to engage in the kind of interaction that occurs when the baseballs impact the window. So what does that difference come to?

²E.g., Bennett (2003; 2008), Pereboom (2002; 2011), Shoemaker (2001), Yablo (1992).

Let's first get one bad answer out of the way. If *Intimacy* is going to meaningfully engage the exclusion problem, the difference between these interactions can't just be the fact that the causes in the former case are intimately related and the causes in the latter case are not. This is, to be sure, a difference between these interactions, but it is not in and of itself a difference of the right sort. The exclusion problem arises in part because it seems, to many, simply mistaken to say that the ordinary effects of mental events are brought about in a manner analogous to effects in textbook cases of causal overdetermination — that my standing up from my desk is overdetermined in the same way that the window shattering is overdetermined by the impacts of the two baseballs. It would be small comfort to learn that the interactions of mental events and their realizers are exactly like the interactions of paradigmatically overdetermining causes *save for the relation between them*. Rather, the proponent of *Intimacy* needs to claim a causal difference between the interactions of mental events and their realizers on the one hand, and those that occur in textbook cases of causal overdetermination on the other. Intimate relations between events may be able to make such a difference, but they cannot themselves be that difference.

Thankfully, nonreductionists have a better answer at the ready. Friends of *Intimacy* typically wish to deny that intimately related events are redundant causes.³ Redundant causation is characterized by the presence of a backup causal process. An effect is caused redundantly when, in addition to one of its actual

³The point is not always explicitly set in terms of redundancy. For example, Yablo (1992) denies that mental events and their realizers “causally compete.” Others, like Bennett (2003; 2008), claim that intimately related events do not satisfy a necessary condition for overdetermination built on Lewis's (1973) counterfactual account of redundant causation. I take such talk to amount to a de facto denial that intimately related events can be redundant causes. For a similar reading of such positions, see Ney (2007).

sufficient causes, there is a further potential sufficient cause which would have brought about the effect had the actual sufficient cause failed to do so.

Overdetermination is a species of redundant causation in which both a sufficient cause and its backup bring about their common, redundantly caused effect.

The guiding idea here is that while mental events and their realizers may be sufficient causes of their common effects, they do not overdetermine those effects because they are not redundant causes. Intimately related events are too closely related to one another — too much like a single event — for either to serve as a causal backup for the other. To quote the old chestnut, the mental is nothing over and above the physical. So construed, ordinary cases of mental causation share a common causal structure with textbook cases of causal overdetermination. In both cases, a pair of events are each independently sufficient for their common effect: each is sufficient for that effect but they are not merely sufficient for it jointly, nor is either sufficient for it distally in virtue of first causing the other event proximally. The difference between them comes down to the fact that in cases of overdetermination the sufficient causes are redundant, while in ordinary cases of mental causation they are not. The intimate relations between mental events and their realizers preclude them causing their effects redundantly.

In the next section, I'll argue that this picture of mental causation is incorrect. Intimate relations preclude neither causal overdetermination nor causal redundancy. Mental events and their realizers can overdetermine their common effects and, when they do so, it is in part because they are redundant causes of those effects. Seeing this, however, requires that we focus on the right cases.

2.3 The problem of distal effects

To *Intimacy*'s benefit, both its friends and critics have taken a somewhat myopic focus on two sorts of cases. On the one hand are textbook cases of overdetermination. On the other are the cases of ordinary mental causation at issue in overdetermination arguments. *Intimacy* builds on a neat distinction between these two sorts of cases. The latter feature intimately related causes. The former do not. So by holding that intimately related causes cannot overdetermine their common effects, it might seem that we can neatly solve the exclusion problem facing nonreductive accounts of mental events.

But this is much too quick. *Intimacy* offers a necessary condition on overdetermination: a pair of causes overdetermine their common effect only if they are not intimately related. If that's correct, intimate relations must not only fail to obtain between overdetermining causes in the textbook cases, they must fail to obtain between overdetermining causes generally. The question we need to ask then is whether there are any further cases of overdetermination that are not cases of textbook overdetermination. I believe there are. Though the textbook cases of overdetermination featured causes that were not intimately related, overdetermination of the same sort can arise even when causes lack the ontological elbow room afforded them in textbooks.

Set overdetermination to one side for a moment. Notice that in so far as mental and physical events are different events, we have good reason to believe that they play different causal roles. Remember, the reason most people came to believe that mental events were not identical to any physical events was that the mental was multiply realizable. We wanted to say, for instance, the same mental event could occur in virtue of the occurrence of different physical realizers. The differences between those realizers can be slight. One may involve the firing of an

extra neuron or two, another could have involved several molecules being positions just slightly differently. These limited physical differences would not make any kind of a mental difference. But they would make a difference to the physical realizer and contribute to its causal profile. As such, the realizer may have causal capacities that are specific to it — capacities that aren't had by other realizers and aren't shared by the mental event it realizes.

With that in mind, consider the following case:

The elderly assassin

Assassin has been tasked with killing Victim. Though an expert sharp-shooter, Assassin is getting rather long in the tooth and has developed severe arthritis in the trigger finger. This arthritis has begun to impair his precision. As a precaution, Assassin's employers have planted a chip in his brain designed to detect just the physical event that actually realizes his decision to fire; if triggered, the chip will fire a second gun. Assassin decides to fire (D) and pulls the trigger (G_1). The physical event realizing this decision occurs (R), activating the chip (C) that triggers the second gunshot (G_2). As it happens, both bullets strike Victim's heart, killing him (V).

Call this case '*Elderly*'. The neuron diagram below illustrates *Elderly*'s basic causal structure. Shaded circles represent events that occur. Arrows represent causation. Dashed lines indicate intimate relations.

Elderly is a counterexample to *Intimacy*. It shows that intimately related causes can overdetermine their common effects. Assassin's decision to fire on Victim (D) and its intimately related realizer (R) overdetermine Victim's death (V) in virtue of causing the first gunshot (G_1) and second gunshot (G_2), respectively. Minus a commitment to *Intimacy* there's nothing terribly interesting

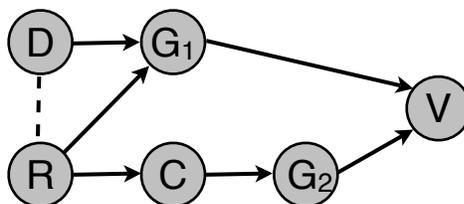


Figure 2.1: The elderly assassin

about their doing so. Both D and R are independently sufficient causes of V . D is sufficient for V independently of R in virtue of causing the first gunshot (G_1). R is sufficient for V independently of D in virtue of activating the chip (C) and consequently causing the second gunshot (G_2). Since D and R are neither identical nor only jointly sufficient for any of these effects, they satisfy the traditional account of overdetermination. More important given the current dialectic, however, *Elderly* satisfies the traditional account for the same reasons as textbook cases of overdetermination. Even though they are intimately related, the decision and realizer bring about different proximal effects (G_1 and G_2 respectively) in virtue of being subsumed differently under the laws. Each initiates a demonstrably separate causal chain and that chain runs to completion in Victim's death. They overdetermine Victim's death.

(Of course, this is not the whole of the causal story in *Elderly*. As D 's realizer, R also causes G_1 . So R in fact causes V via two causal chains: $R-C-G_2-V$ and $R-G_1-V$. But this is beside the point. Whatever we wish to say about this additional causal chain and its relation to either Assassin's decision D causing V via G_1 or the realizer R causing V via C is a further matter from whether D and R in fact overdetermine V via those causal intermediaries. I'll return to $R-G_1-V$ below, but we can set it aside for now.)

We can sharpen these observations into an argument. We saw previously that

Intimacy is best read as a rejection of the claim that intimately related causes ever bring about their effects redundantly. As we noted there, overdetermination is not the only form of redundant causation. More often than not, redundant causation takes the form of preemption: cases in which an effect is brought about by one causal process, but in which that effect would have been brought about by a second, backup causal process had the former process failed.⁴ If *Intimacy* is correct that intimately related causes cannot be redundant, then they must not only be precluded from overdetermining their common effects, they must also be precluded from preempting one another from causing those effects. Should either of a pair of intimately related events fail to bring about their would-be common effect, the other cannot act as a backup and bring it about instead.

And that's simply not true. We can easily imagine variations on *Elderly* in which the realizer causes Victim's death but the decision does not. Perhaps the first gun jams, or a selfless passerby leaps into the path of the first bullet, or maybe the second gun just fires more quickly and delivers the fatal blow before the first bullet can reach its target. In each of these scenarios, Victim's death is caused by the second gunshot and not the first, and so by the realizer of Assassin's decision and not the decision itself. If intimately related causes truly could not be redundant causes, that would be impossible.

We now have a blueprint for a general class of counterexamples to *Intimacy*. *Elderly* illustrates that intimately related causes can overdetermine a common

⁴I use 'preemption' broadly here as a catch-all for cases of redundant causation in which one potential cause of an effect fails to bring that effect about and another succeeds. (See Paul and Hall (2013) for another explicit example of this use.) 'Preemption' may also be used more narrowly to denote only those cases where the potential causes fails to bring about the effect *because* of the actual cause of that effect (e.g., Lewis (1973; 1986)). It denotes a type of redundant causation on either use.

distal effect if they bring it about in virtue of causing different proximal effects. Unless a pair of events must cause all and only the same effects then, they are able to act as overdetermining causes.

The moral of this case is that intimate relations between causes are not sufficient to distinguish causal overdetermination from the causal interactions that occur in ordinary mental causation. Some intimately related causes overdetermine their common effects. Others may not. But should some intimately related causes fail to be overdetermining causes, it is not solely on account of their intimacy.

2.4 More intimacy

Nonreductionists attracted to *Intimacy* can retrench. Overdetermination of the sort that arises in *Elderly* is not directly at issue in the causal exclusion problem. *Elderly* poses a challenge to *Intimacy* in part because it legitimately appears to be a case of causal overdetermination. Nonreductionists can (and should) allow that mental events and their realizers overdetermine effects like Victim's death. What they wish to deny is that overdetermination occurs in ordinary cases of mental causation. To do that, they need to distinguish the interactions of mental events and their realizers in those ordinary cases of mental causation from those like *Elderly*.

There is admittedly something fishy about *Elderly* and the blueprint that it provides for counterexamples to *Intimacy*. *Elderly* features a pair of intimately related causes that overdetermine an explicitly distal effect. To bring about that effect, each overdetermining cause relies on the help of a string of causal intermediates and a corresponding set of background conditions. But they do not rely on the same string of intermediates. Nor do they rely on the same background conditions.

Typical cases of mental causation just aren't like that. When I decide to reach for the cup of coffee on my desk, my decision and the neural event that realizes it both cause me to lift the cup off the desk. But they don't do so by way of a baroque physiology in which each kicks off a wholly distinct causal chain, meeting only in the event of the cup's lifting off the desk. Instead they both cause the cup to lift by causing the muscle contractions in my arm, and they both cause those muscle contractions by causing a calcium cascade through the relevant muscle fibers, and so on as we follow the path of causation upstream. There is a single, common set of intermediate effects that link both the decision and its realizer to their distal effect. Background conditions would seem to fair similarly. The background of events and properties that the realizer relies on to cause the lifting are, if not the same as those on which the decision relies, at least closely related. Nothing in the background plays a role similar to Assassin's chip in *Elderly* — a background condition on the sufficiency of one overdetermining cause, but entirely irrelevant to the causal work of the other.

Nonreductionists might use these differences to refine *Intimacy*. It's not unreasonable to suspect that intimately related causes might overdetermine some effects in cases like *Elderly* because, while *they* are closely related, the supporting factors on which their efficacy depends are not. This suggests a natural way to qualify *Intimacy*: intimate relations between causes *do* preclude overdetermination, but only so long as they extend in some manner to the supporting factors that are involved as well.

A thought along these lines seems to lie at the heart of Karen Bennett's treatment of the exclusion problem.⁵

⁵It is worth noting that while Bennett argues that mental events and their realizers do not overdetermine their common effects, the machinery behind her argument allows for intimately related events that overdetermine some of their common effects and not others. Bennett, and

“[T]here is an interesting connection between the relation that holds between mental events. . . and the physical ones that ‘underlie’ them, and the relation that holds between the physical cause and the effect. . . . [T]he conditions in which the physical event. . . occurs with the mental one are the same as the conditions in which the physical event or property manages to bring about the effect.” (Bennett, 2003, p. 489)

Bennett’s suggestion is that the intimate relations between mental events and their realizers preclude them from overdetermining their certain effects *because* those relations assure that mental events are relevant to their realizer’s can cause those effects. The realizer’s sufficiency for some effects requires that certain background conditions are present. The occurrence of the mental event “constitutively involves” the occurrence of those conditions. As a result, the physical cause depends on the mental for its efficacy in a straightforward, counterfactual sense — had the mental event not occurred, the physical event would not have caused its effect. Bennett takes this dependence to show that intimately related mental and physical causes do not overdetermine their common effects.⁶

The connection that Bennett describes consists of a series of intimate relations. They obtain between the mental event and its realizer, but also between the mental event and at least some of the other physical events that help the realizer bring about certain effects. For friends of *Intimacy*, this is more grist for the mill. A single intimate relation between a pair of causes does not preclude

perhaps other nonreductionists who appear to endorse *Intimacy*, might be better read as endorsing a position in the vicinity of the view that I refer to below as ‘*Co-dependence*’.

⁶Bennett is careful to deny that the mental cause is itself a background condition or causal intermediate with which the physical cause brings about its effects.

overdetermination. But a series of such relations plausibly just might. Suppose we extended *Intimacy* along these lines. Say that two sets of causes are *pairwise intimately related* just in case each member of one set is either identical with or intimately related to some member of the other set and vice versa. It's then just a small step from *Intimacy* to the following position:

Co-dependence For any causes C and C^* and their effect E , let B be the set of background conditions and causal intermediates that C relies on to cause E and B^* be the set of background conditions and causal intermediates that C^* relies on to cause E . C and C^* overdetermine E only if $\{C\} \cup B$ and $\{C^*\} \cup B^*$ are not pairwise intimately related.

Like *Intimacy*, *Co-dependence* holds that intimate relations between causes preclude causal overdetermination. But to do so, they need to obtain more pervasively than *Intimacy* requires. They need to obtain both between a pair of causes, and also between the supporting factors that those causes rely on to bring about the effect in question.

Co-dependence is just one way to adapt the idea that intimate relations preclude causal overdetermination in such a way that it accommodates cases like *Elderly*. There may be others according to which less is required to preclude overdetermination. (It might turn out, for instance, that we can soften the view to allow for non-intimately related causal intermediates or only require that intimate relations obtain between some background conditions and not others.) What there aren't, however, are positions of this sort that require more. We've run out of causal relata that might be intimately related. If *Co-dependence* is false, then so is the thought that intimate relations preclude causal overdetermination.

2.5 Against *Co-dependence*

Co-dependence may seem to be an improvement on *Intimacy*. But consider the following scenario:

ET: Black is a government neuroscientist conducting secret research on captured alien life. She's discovered that the aliens' brains exhibit an interesting redundancy: each hemisphere of the brain has its own independent realizer for the feeling of fear and those realizers differ from one another in kind.⁷ How their brains respond to fear depends both on how the fear is realized and which hormones are present in the brain at the time. If the fear is realized in the left hemisphere or hormone X is present, a weak alarm signal is sent to the motor system. If the fear is realized in the right hemisphere and hormone Y is present, a strong alarm signal is sent to the motor system. One day while Black is monitoring her subjects, an alien becomes afraid. The feeling of fear (*F*) is realized by a neural event in the right hemisphere (*R*) and hormone Y is present (*Y*), triggering an alarm signal (*S*) to be sent to the alien's motor system.

Unlike *Elderly*, the causes at the center of *ET* are proximal to their common effect and both they and their relevant background conditions are pairwise intimately related. Given the presence of the hormone (*Y*), the alien's fear (*F*) and its realizer (*R*) both cause the alarm signal (*S*). Suppose in fact that they cause it directly. Even so, *F* and *R* overdetermine *S*.

⁷That is, these mental events are multiply realized within individuals, rather than merely between individuals. To borrow a term from Horgan (2001), such events are *strongly multiply realized*.

To see why, notice that *ET* and *Elderly* share the same basic causal structure. *Elderly* featured a pair of intimately related events, both of which were independently sufficient causes of their common effect, but only one of which was an independently sufficient cause of that effect's properties. We find the same pattern here. Given Y , both F and R are independently sufficient for S . But they are not both independently sufficient for S having the property *being a strong signal*. Whether fear triggers a strong alarm signal depends on how it is realized. Fear realized in the left hemisphere, even in the presence of hormone Y , does not trigger strong alarm signals. Strong signals are triggered only in response to fears that are realized in the right hemisphere. So given Y , R is sufficient for S being a strong signaling. F is not.

As in *Elderly*, we can imagine variations on this case in which its causal structure breaks down into a form of preemption. For example, suppose that Black had introduced a chemical into the alien's brain that inhibits the production of alarm signals such that conditions that would have triggered a weak signal produce nothing and conditions that would have triggered a strong signal produce only a weak signal. F , R , Y , and S would all still have occurred. But F would not have caused S . The inhibitor prevents F from causing S , leaving R to act as the backup cause of S . This is a straightforward case of preemption, and as such an instance of redundant causation.

Co-dependence is therefore false. In *ET*, both F and R are independently sufficient causes of S given Y , F, Y and R, Y are pairwise intimately related, and F and R cause S directly, rather than by way of an intermediary. If *Co-dependence* were correct then, F and R would be unable to redundantly cause S , when in fact they can do just that.

2.6 Overdetermination

Intimacy and *Co-dependence* are false. Intimate relations do not stand in the way of mental events and their realizers being overdetermining causes. Nonreductionists therefore must look elsewhere for a solution to the exclusion problem. The trouble with looking to *Intimacy* or *Co-dependence* as a means of escaping the causal exclusion problem, however, runs deeper than this. Nonreductionism would be saddled with the exclusion problem even if *Intimacy* or *Co-dependence* were true.

To appreciate the underlying issue here, we need to consider why overdetermination poses a threat to mental causation in the first place. Why not just accept that, if nonreductionism is correct, ordinary cases of mental causation are instances of overdetermination, no different from deaths by firing squad? The answer, I think, is that ordinary cases of mental causation *are* different from deaths by firing squad and other textbook cases of causal overdetermination. If they are instances of overdetermination, it is overdetermination of a different and more puzzling sort than we find in textbook cases.

It will be helpful to distinguish several different types of causal interactions. Take a simple case to start.

Suppose that a pair of vandals, Billy and Suzy, spray paint a green mark on the wall. Billy presses down on the nozzle of a yellow spray can, Suzy presses down on the nozzle of a blue spray can, and where the two jets of paint meet they leave a green mark.

Distinguish (if only for illustration) the appearance of the green mark on the wall from the appearance of the mark simpliciter on the wall. The former is green essentially. The latter is not.

Billy's press and Suzy's press jointly cause the green mark. Though both cause

the green mark to appear, neither press alone is sufficient for it. Billy's press is sufficient for a yellow mark. Suzy's press is sufficient for a blue mark. To cause the green mark they must act in concert with one another. Contrast this with the relation that Billy's press and Suzy's press stand in to the mark simpliciter. Each press alone is sufficient for its occurrence. Given the laws and the background conditions that obtain (e.g., paint is in the can, the nozzle is pointed in the right direction), either press is enough to assure that the mark will appear on the wall. It may not have been green, but it would have occurred all the same.

This latter interaction is a textbook case of causal overdetermination. Overdetermination in cases like this is an artifact of how we individuate events. It requires that we recognize events like the mark simpliciter that have some of their properties accidentally; the mark might have been blue or yellow, but it would still be the numerically same event. Abandon such a theory of events for one that counts any difference in properties as a difference in event identity and Billy's press and Suzy's press are no longer individually sufficient for the appearance of the mark that actually occurs. Overdetermination would disappear, replaced by joint causation.

We can imagine overdetermination of a more thorough-going sort. Overdetermination in textbook cases like the vandals' spraying was impure in the sense that it was a product of the grain with which we individuate events. In contrast, overdetermination of this new sort would be pure; it would occur no matter how fine-grained and fragile we take events to be. Suppose that Billy and Suzy throw rocks at a fragile window and their rocks strike the window simultaneously, shattering it. For this to be a case of pure overdetermination, we need to imagine that the impact of each rock against the window is independently sufficient for the window to shatter in precisely the way it did when it was struck

by both rocks. But be careful. This is not to imagine that Billy's rock and Suzy's rock subtly impede one another such that they each strike the window less forcefully when thrown together than were they thrown separately, or that, really, one rock struck the window first and caused the shattering before the other could do so. No. That way lies closet joint causation or preemption. To count as pure overdetermination, each throw individually must cause the window to shatter precisely as it does when both vandals throw their rocks.

Pure overdetermination is strange. For one, the case we've imagined runs counter to the physical laws. When the window shatters under the impact of both rocks, the way it shatters — its full suite of properties — reflects the amount of force with which it was struck. To suppose that each throw was sufficient to cause the window to shatter in that way is to suppose that each rock alone strikes the window with the combined force of both rocks. Under the laws this cannot occur. A window struck by two rocks, each applying force f , should shatter in a manner that reflects its being struck with a force of $2f$. But by hypothesis this is not the case in pure overdetermination.

More generally, pure causal overdetermination stands apart from both impure causal overdetermination and joint causation in that it involves a failure of additivity.⁸ The causal interactions that occur in cases of joint causation and impure overdetermination are *additive*: acting together, multiple events cause a common effect to be propertied in a manner that they would not have caused individually. In contrast, pure overdetermination involves a striking failure of additivity. Each overdetermining cause is independently sufficient not only for their common effect, but for its full suite of properties.⁹

⁸I take this point from Paul (2007).

⁹In framing the issue this way, I'm ignoring certain extrinsic properties such as those that relate to an event's causal history (e.g., *is caused by two events*).

Additive causes play a difference-making role that non-additive causes do not. When the two vandals sprayed paint on the wall, each spray of paint caused the mark that appeared on the wall to have a certain set of properties. But in addition to that, each spray also made a difference to what properties the *other* spray caused the mark to have. Without the spray of blue paint, the spray of yellow paint would still have caused a mark to appear on the wall, but it would not have caused the mark to have all of the properties that actually possessed. In interactions that are not additive, one or more of the interacting events do not make that difference with respect to the other(s). They may cause a common effect to have certain properties, but they make no difference at all to what properties the other interacting events cause their effect to have. To put the point a slightly different way, a pair of additive causes make a difference to whether, when, where, or how their common effect occurs and they make that difference with respect to one another. If not for the fact that both causes brought about that effect, it either would not have occurred, or it would have occurred at a different time, or in a different place, or in a different manner. This type of difference-making is absent in causal interactions where additivity fails.

Pure causal overdetermination is not just an idle possibility if nonreductionism is true. If mental events and their realizers overdetermine their common effects, the overdetermination is pure. When I reach for the cup of coffee on my desk, my action is caused both by a mental event (my intention to act) and by a physical event (the realizer of that intention). Given the natural laws and the relevant background conditions, each event is sufficient for my action to occur in just the way it did: I do not reach more quickly or grasp the cup more firmly because my doing so is caused by a mental event *and* its realizer. The mental and physical events do not act jointly, nor does one event cause the other. This would be of no

consequence if the mental and physical events were identical. But for nonreductionism these events are different. And so we seem to have overdetermination of the pure sort.

The problem that overdetermination poses for mental causation then is at least this: if mental events and their realizers overdetermine their common effects, they do so purely. Should pure overdetermination occur, it would mean that there are two quite different *kinds* of interactions that events may engage in: additive interactions and non-additive interactions. That difference would call out for an explanation. We would need to explain why additivity fails in some cases of causal overdetermination and not in others — cases that ostensibly share a common causal structure. Absent that explanation, the failure of additivity engendered by pure causal overdetermination would make causation mysterious in a way that it need not be.

2.7 Explaining failed additivity

Intimacy and *Co-dependence* could not have substantively addressed the problems raised by pure overdetermination. They fall short on two counts. The distinction that they draw between cases in which overdetermination occurs and those in which it does not is in some respects too fine to do the explanatory work that the nonreductionist asks of it and in other respects too coarse.

First, note that the problem that pure overdetermination presents for nonreductionism does not concern overdetermination or even redundant causation per se. Pure overdetermination is problematic at least in part because it entails a failure of additivity that demands explanation. The problem lies in that failed additivity however, not in the overdetermination itself. Denying that mental events and their realizers are overdetermining causes therefore does not free the

nonreductionist from the need to account for that failed additivity any more than denying the antecedent of a conditional serves as a rejection of the consequent. Whatever the correct ontology of mental events may be, mental events and their realizers are not additive causes. That much is just a brute fact. So long as mental events and their realizers are different events with common effects then, we need to explain why they are not additive causes. Whether they are overdetermining causes or merely independently sufficient causes of their common effects is beside the point.

To that end, the claim that mental events and their realizers are not redundant causes is at best irrelevant and at worst a liability. Non-redundant causes of the same effect are often additive. Joint causes of an effect are not redundant and joint causation is just additive causation writ large. When two causes bring about an effect jointly (as, for instance, when the vandals' spraying causes the green mark) they do not each just make a difference to what properties one another cause that effect to have, they make a difference to what effect one another cause outright. Nonreductionists who deny that mental events and their realizers are redundant causes have therefore merely shifted the bump under the rug. The pressing question simply changes from "Why does additivity fail in some cases of causal overdetermination and not others?" to "Why does additivity fail in some cases of non-redundant causation and not others?" Even if mental events and their realizers were not redundant causes then, their failed additivity would be no less mysterious.

Second, *Intimacy* and *Co-dependence* also draw too coarse a distinction between cases to adequately account for the failed additivity that would arise in mental causation under nonreductionism. Consider *Elderly*. Assassin's decision and its realizer had two common effects in that case: the first gunshot and

Victim's death. Neither of these interactions is additive. Note however that additivity fails in a different way in the former interaction than in the latter. With respect to the gunshot, additivity fails in full. The decision and its realizer were independently sufficient causes of both the gunshot and its properties; neither event made a difference to what properties the other event caused their effect to have. With respect to Victim's death, however, additivity fails only in one direction. In causing Victim's death, the decision made no difference to what properties the realizer caused the death to have. The realizer, after all, was itself a sufficient cause of both the death and all of its properties. The decision was not. By causing the death, the realizer *did* make a difference to what properties the decision caused the death to have.

Part of the nonreductionist's task in accounting for the failed additivity that would occur in cases of mental causation is to account for this difference. She needs to explain why Assassin's decision and its realizer stand in a different relation to the first gunshot than they do to Victim's death. The intimate relation between them cannot do that. At best, the fact that the decision and realizer are intimately related distinguishes their interactions from those of events that are not intimately related. It cannot distinguish the interaction through which the decision and realizer cause one effect from the interaction through which they cause a different effect.

Moreover, note that just as we saw with overdetermination, we can't look to supporting factors such as background conditions and causal intermediates to adequately distinguish these interactions. The relationship that Assassin's decision and its realizer stand in to Victim's death in *Elderly* would seem to parallel the relationship that the alien's fear and its realizer stand in to the alarm signal in *ET*. In *ET*, a mental event and its realizer caused their common effect in virtue of the

same set of background conditions and without first causing an intermediate effect. None the less, the fear and its realizer did not cause the alarm signal additively. In causing the alarm signal, the realizer made a difference to what properties the fear caused the alarm signal to have. (The fear itself was sufficient to cause a *weak* alarm signal; the signal that it actually caused was strong.) The fear made no such difference with respect to its realizer. The realizer, together with the presence of the hormone, was sufficient for both the signal and its properties. Like the interaction through which Assassin's decision and its realizer cause Victim's death, the interaction through which the alien's fear and its realizer cause the alarm signal involves a one-way failure of additivity. Whatever might account for the difference in additivity between those interactions on the one hand and, on the other, Assassin's decision and realizer causing the first gunshot, it cannot just be a matter of the intimate relations between events.

CHAPTER 3

Causal contribution and causal exclusion

3.1 Introduction

Suppose that Billy and Suzy splash paint on the wall. Billy throws blue paint. Suzy throws yellow paint. A green spot forms on the wall where the blue paint and yellow paint run together. Billy's throw and Suzy's throw are joint causes of the green spot. Though each throw causes the appearance of the green spot, it does so only in virtue of the other throw. Had either vandal not thrown, the other vandal's throw would not have caused the appearance of a green spot.

Cases like this draw out an incongruity in the way we've come to think about causation. On the one hand, it is a simple reminder of the importance of background conditions to causation. What an event causes depends not only on its nature and the laws, but also on the background of events (and perhaps omissions) against which it occurs. Billy's throw might have occurred in the very same manner and at the very same time as it actually did and, had the circumstances in which it occurred been different, it may have had different effects. Had Suzy not thrown, Billy's throw would have caused the appearance of a blue spot. Had Suzy thrown white paint rather than yellow, Billy's throw would have caused a spot in a lighter shade of blue. Causation in this respect is subject to the vagaries of the environment. It involves an element of luck.

At the same time, our routine causal judgments often seem to be guided by the thought that causation, in some respect, is immune to this type of luck. Granted, the effects of Billy's throw are contingent on background conditions such as the

occurrence of Suzy's throw. But intuitively there is something that Billy's throw *does*, some influence that it exerts, that is not contingent on its background conditions. Suzy could have held her throw or used a different color of paint. She could have done anything else at all. The influence of Billy's throw would have been, in some respect, the same. Call this influence the contribution of Billy's throw. An event's contribution is the influence that it has on the world given only its nature and the laws.

Contemporary work on causation has lost track of contribution. Our best theories of causation do not address it. Contribution is not the causal relation. The causal relation that Billy's throw stands in to the green spot depends on Suzy's throw. The contribution of Billy's throw does not. By the same lights, contribution is not counterfactual dependence, or the lawful entailment of an event, or any of the more sophisticated variations on these relations that have been put forward as the basis for a theory of causation. Contribution is a recognizably causal notion. But it has gone missing from the current debate on causation. This neglect has come at a cost. By ignoring contribution, we've allowed otherwise tractable problems involving causation to fester.

My concern in this paper lies with one such problem: the causal exclusion problem for mental events. The exclusion problem is often posed as a challenge to the efficacy and irreducibility of mental events — believings, desirings, experiencings, and the like. Many of us wish to say that mental events are both causally efficacious and distinct from any underlying physical events in the brain. Should I reach for the coffee on my desk, it would seem that my reaching is caused both by my desire for coffee and by a neural event that realizes that desire. But if that's the case, the efficacy of the mental is superfluous. The cumulative effect of my desire and its realizer is in no way different from what the realizer itself is

sufficient to cause. In this respect, the two events act like a single cause. The presence of the second event adds nothing to the effect. At its heart, the causal exclusion problem is the problem of explaining how this is possible if mental events and their realizers are in fact different causes.

I will argue that contribution provides a new solution to the exclusion problem. Mental events and their realizers interact with one another in a different manner than events like Billy's and Suzy's throws because of a prior difference in their contributions. These contributions allow a mental event and its realizer to be both different and efficacious and yet act in the relevant respect like a single cause. The argument will proceed in six parts. Section one provides a brief account of the exclusion problem for mental events and the puzzle about causation that lends it its force. Section two then develops an account of contribution and sections three through five draw on that account to resolve the exclusion problem. Finally, I conclude in section six by comparing this new solution to the exclusion problem with a number of more familiar proposals. This solution has a notable advantage over its rivals. It does ready justice to the claim that the mental is nothing over and above the physical.

3.2 The causal exclusion problem

The exclusion problem has developed a bad rap. It is often dismissed by critics as obscure or, worse, a pseudo-problem. As they see it, the exclusion problem simply has not been shown to be a legitimate source of concern (e.g., Sider (2003)).

Before we can properly address the problem then, I believe that it is worthwhile to lay out the exclusion problem in a bit more detail than it is sometimes afforded and make the case that it warrants our attention.

Consider an rather ordinary case of mental causation. Suppose that you stand

and applaud after seeing a play. As a deliberate action, your applause is caused by various mental events such as your decision to applaud. As a physical action, your applause is also caused by certain physical events in your brain, including (it would seem) the physical event that realizes your decision to applaud. Now if your decision to applaud and its realizer are in fact one and the same event, this is innocent double counting. Though we refer to the decision and its realizer under different names or descriptions, they are one cause, not two. But if they are not the same event, your clapping has two different, but parallel, causal histories.¹ It is overdetermined in the sense that it has two or more sufficient causes that are not themselves causally related. Given their background conditions and the laws, the decision and realizer are both sufficient causes of the clapping.

The exclusion problem is that it is not clear how overdetermination of this sort can in fact occur. If either the decision or its realizer is a sufficient cause of your applause, the other is in some respect superfluous. Proponents of the exclusion problem sometimes seem to worry that superfluous causes are inherently problematic. For instance, a recurring theme in Kim's work on the exclusion problem is that the efficacy of physical events would deprive mental events of the opportunity to do any causal work:

[T]o acknowledge that [an event] p has a physical cause, p , at t is to invite the question: Given that p has a physical cause p , what causal work is left for [a mental event] m to contribute? The physical event

¹The arguments against the reducibility of mental events are well-traveled. Mental events cannot be type-identical with their their realizers because mental events are multiply-realizable. Similarly, mental events cannot be token-identical with their realizers because they have different modal properties. Your desire to applaud — that very event — could have been realized by a different physical event. But if your desire could have occurred in the absence of its actual realizer, it cannot be identical to that actual realizer.

therefore threatens to exclude, and preempt, the mental cause. This is the problem of causal exclusion. (Kim, 1998, p. 37)

The thought seems to be that if the physical realizer of a mental event is a sufficient cause of an effect, there is “no causal work left over” for the mental event to do for that effect and so, for lack of causal work, the mental event cannot cause that effect. If that is indeed true, your decision and its realizer cannot possibly overdetermine your applause. Should either event be a sufficient cause of your applause, the other would be ineligible, excluded.

What such concerns amount to is notoriously unclear. (Talk of causal work, for example, is at best a metaphor.) I myself am not moved by them and won't explore them here. But whatever the merits of these complaints, causal overdetermination by an event and its realizer poses a prior problem. For whether or not such overdetermination is objectionable in its own right, it is in the very least quite different from other causal interactions. Causal overdetermination is just one type of causal interaction. Say that a causal interaction occurs just in case two or more events that are not themselves causally related cause the same effect. If mental events and their realizers overdetermine their common effects, their interactions are quite unlike those of most other events.² Recognizing causal overdetermination by mental events and their realizers therefore carries a significant explanatory burden. If we're going to allow that mental events and

²To be sure, the problem may not be unique to mental events and their realizers. As many others have noted (e.g., Block (2003)), fundamentally the same problem can arise in other cases where reductions are controversial. For example, if biological events are not identical to the chemical events that underlie them, biological and chemical events would seem to overdetermine their common effects in the same manner as mental events and their realizers. In this respect, the exclusion problem itself is general problem that arises from the rejection of various reductions. Mental events and their physical realizers simply present a paradigmatic example.

their realizers overdetermine their effects, we need to explain why their interactions are fundamentally different from those of other events.

To illustrate, compare the case of your applause with a case that involves a superficially similar causal interaction. Suppose that a pair of golf balls crash into a window at the same time, causing it to shatter. Like the case of your applause, a pair of events that are causally unrelated bring about a common effect. In the former case, your decision and its realizer cause you to clap. In this latter case, the impact of each golf ball against the window pane causes the window to shatter. In fact, if we suppose that each impact is sufficient to cause the window to shatter, both these cases are examples of causal overdetermination.

Despite this similarity, the decision and realizer stand in a very different relationship to your clapping than the impacts of the two golf balls stand in to the window shattering. Notice that even if we suppose that the impact of each ball is sufficient to cause the window to shatter, neither impact itself is sufficient to cause the shattering to have all of its actual properties. The window shatters in a slightly different manner under the impact of two golf balls than it would have shattered under the impact of just one. It breaks into different shards. It is left with a larger hole. To borrow a term from L. A. Paul (2007), the causal interaction of the two impacts is *additive*. Acting together, the impacts cause their effect to be propertied in a way that neither event would have caused individually. In contrast, the case of your applause involves a failure of additivity. Both your decision and its realizer are sufficient causes of your clapping and all of its properties.³ You do not clap more loudly, or quickly, or with added enthusiasm because your clapping was caused by *both* the decision and its realizer rather than just one or the other.

³In framing the issue this way, I'm ignoring certain extrinsic properties of events, such as those that relate to an event's causal history (e.g., *is caused by two events*).

Additive causes play a difference-making role that non-additive causes do not. If two golf balls strike the window and shatter it, each impact causes the shattering to have a certain set of properties. But in addition to that, each impact also makes a difference to what properties the *other* impact causes the shattering to have. Without the impact of the first ball, the impact of the second ball might still have caused the shattering, but it would not have caused the shattering to have all of the properties that it actually did. In interactions that are not additive, one or more of the interacting events do not make that difference with respect to one another. They may cause a common effect to have certain properties, but they make no difference at all to what properties the other interacting events cause their effect to have. To put the point a slightly different way, a pair of additive causes make a difference to whether, when, where, or how their common effect occurs and they make that difference with respect to one another. If not for the fact that both causes brought about that effect, it either would not have occurred, or it would have occurred at a different time, or in a different place, or in a different manner. This type of difference-making is absent in causal interactions where additivity fails.

Such failed additivity is puzzling. Consider just how different the case of your applause is from the window shattering example. Suppose that we re-imagine the window shattering such that it involves a failure of additivity. We need to imagine that one ball shatters the window, and the other ball shatters the window, and that the window shatters in precisely the same way that it would have shattered had it only been struck by a single ball. And be careful. This is not to imagine that the two balls subtly impede one another so that they each individually strike the window with half as much force or that one ball strikes the window first and the other sails harmlessly through the already broken pane. No. To do that is to

imagine an additive interaction or no interaction at all. Rather, we need to imagine that each ball strikes the window in such a way that it is fully responsible for the shattering's properties, from the time that it occurs down to the shape and momentum of each shard of glass that forms as the window breaks apart.

As I see it, this puzzle is the heart of the causal exclusion problem. Should mental events and their realizers overdetermine their effects, there are two quite different kinds of causal interactions that events engage in: additive and non-additive. By adopting a theory of mind on which there is such a difference, we take on the burden of accounting for it. We need to explain why causal additivity systematically fails in the interactions of mental events and their realizers when it does not fail in causal interactions generally. Absent that explanation, we make causation mysterious in a way that it need not be. We might, after all, have adopted a theory of mind that does not engender a systematic failure of causal additivity. We might have identified mental events with physical events or denied that they have the effects we typically take them to have. Either route leads us away from mental-physical causal overdetermination and the failed additivity that ensues. Embracing that overdetermination means that we have some explaining to do.

In what follows, I'll argue that contributions provide that explanation. I believe that the reason that additivity fails in some interactions but not others lies in the contributions of the interacting causes. Additive causation requires causes that make the right sorts of contributions. While the events that commonly concern us make such contributions, mental events and their realizers do not.

3.3 Contributions as lawful constraints

So what is a contribution? The answer I want to explore here is that contributions are constraints that events impose on how the state of the world will evolve. Think of the world's evolution from one state to the next as a path through a state space. Given the laws, an event influences the course of that path. That influence is its contribution.

One brief caveat. How we should understand contribution depends on whether the laws are deterministic or indeterministic. As is common in discussions of causation, I'll assume that determinism is true. That is, I assume that the laws and complete state of the world at one time together entail the complete state of the world at every other time. This assumption may very well prove false, of course. Even so, I believe that it is fruitful. An account of contributions under indeterminism can take the present account as a starting point. Regardless, understanding the contributions of individual events under determinism is valuable in its own right.

The question then is what influence an event might have on its own, independently of the other events in its environment. To start, take a case in which there can be no question of how to distinguish the influence of a single event from that of its environment. Consider the complete state of the world at this very moment. The obtaining of this state is an event. Under determinism, that event imposes a well-defined constraint on how the world will evolve in its wake. Together with the laws, it determines the succession of states that will follow it. It dictates which events will occur and which will not. It lays out a complete path through the state space.

Most events that concern us do not have such influence. Births, deaths, handshakes, and the like (however world-changing they may be) do not fix the

world on a single, unwavering course. It is consistent with the laws for an ordinary, medium-sized event to be immediately followed by the obtaining of any number of different states. For example, Billy's throwing the blue paint at the wall at time t_1 may be followed at t_2 by the obtaining of a world state that includes a blue spot on the wall, or a green spot on the wall, or a light-blue spot on the wall. Each of these states — and many more to be sure — is consistent with the laws and Billy's throwing at t_1 . The same goes for events. Billy's throw does not itself assure that any event subsequently occurs.

So what influence can a single, ordinary event claim? Notice that while an event like Billy's throw at t_1 may not assure that any one state obtains at t_2 , it *may* assure that certain states do not obtain. Many states are just not compossible with Billy's throw under the actual laws. For if the laws are deterministic as we've assumed, the obtaining of the state of the world at one time lawfully entails the obtaining of the state of the world at every other time. Absent a violation of the laws then, those world states that could have obtained at t_2 had Billy not thrown at t_1 could *only* have obtained at t_2 had Billy not thrown at t_1 . Billy's throw makes it nomically impossible that these states obtain at that time.

Billy's throw at t_1 imposes a constraint on what the world can be like at t_2 . The state of the world that obtains at t_2 must be a nomically possible state of the world at t_2 *given Billy's throw at t_1* . Take the range of possible world states that are themselves internally consistent with the laws. Each of these states is, in a broad sense, a nomically possible state of the world at t_2 . Far fewer are nomically possible at t_2 given Billy's throw at t_1 . Billy's throw imposes a constraint on what the world can be like at t_2 in that it limits the number of states in this range that are eligible to obtain at t_2 .

In this respect, the influence of ordinary events like Billy's throw is

fundamentally no different from that of an obtaining world state. Both constrain the range of states that may obtain in their wake. They permit the world to evolve towards some states and prohibit it from evolving towards others. Where they differ is the severity of that constraint. The obtaining of a complete world state limits the range of states that may subsequently obtain so thoroughly that only a single state remains eligible to obtain at any given time. It lawfully entails that these states obtain. In contrast, the constraints imposed by ordinary events such as Billy's throw are open-ended. They limit the range of states that may subsequently obtain, but those limits leave multiple states eligible to obtain at any one time. This is no doubt a significant difference between these constraints. But it is a difference of degree, not kind.

I propose that this constraint is an event's contribution. Imposing that constraint is what an event *does*, on its own, simply given the laws. Going forward, it will be helpful to have some formalism. We can understand an event's contribution as the set of nomically possible worlds in which that event occurs. Each nomically possible world marks one way in which the state of the world may evolve under the laws. An event's contribution is a subset of this broader range of worlds. It carves out the range of ways in which the state of the world may lawfully evolve given the event's occurrence. In this way, contributions do indeed serve as a kind of constraint. If it is nomically possible for a world state to obtain in the wake of a given event, the contribution of that event must include a world in which the state obtains in the wake of that event. If it is not nomically possible for that state to obtain in the wake of an event, its contribution will not include such a world.

To illustrate, return again to the paint-throwing vandals. When Billy throws his paint at the wall, his throw makes a contribution. That contribution is the set of nomically possible worlds in which the throw occurs. That set includes the

actual world, of course, but also worlds that break with actuality in large and small ways. For example, it includes worlds where Suzy throws different paint or aims at a different target. It also includes worlds where the surrounding environment is quite different: worlds in which Billy throws in the middle of a desert or into an oncoming hurricane. How much variety is found in the contribution depends in part on what event we take it to be the contribution of. For instance, if we suppose that the contributing event is particularly fine-grained and modally fragile, the contribution may only include worlds in which the throw occurs precisely when, where, and how it actually does. Alternatively, if we grant that the throw could have occurred a bit earlier, or later, or that it still could have occurred if Billy had stood a foot to the left and thrown a little harder, then such worlds must also belong to its contribution.

3.4 Causal contributions and causal additivity

What does contribution have to do with causal additivity? Recall that additive causes are difference-makers. A pair of events are additive causes of some effect only in so far as they each make a difference to what the other causes. The first event in the pair must make some difference to whether, when, where, or how the second event causes their common effect to occur and the second event must make a similar difference with respect to the first event. Difference-making as such is a matter of contrasts. For one event to make a difference to what another event causes, it must do more than merely make it the case that this other event causes its effect to occur when, where, and how it does. The former event must make it the case that the latter event causes its effect to occur just as it does, *rather than* at a different time, or place, or in a different manner, or as opposed to not causing the effect at all. It must prevent at least one of these alternate possibilities from

being actualized.

Take our vandals. Billy's throwing blue paint and Suzy's throwing yellow paint cause the appearance of a green spot. But they also prevent one another from causing the spot to be a different color. Had Billy not thrown or had his throw been off the mark, Suzy's throw would have caused the appearance of a yellow spot. Similarly, had Suzy's throw not occurred, Billy's throw would have caused the appearance of a blue spot. Neither of these possibilities are actualized. Billy's throwing blue paint prevents the spot from being yellow. Suzy's throwing yellow paint prevents the spot from being blue. It is in part because of this that each throw makes a difference to what the other causes. Billy's throw makes it the case that Suzy's throw causes the spot to be green rather than yellow. Suzy's throw makes it the case that Billy's throw causes the spot to be green rather than blue.

A pair of events therefore can only be additive causes if they satisfy two conditions. First, there must be a relevant difference for each event to make. It must be possible under the actual laws that each event failed to cause the effect to occur precisely as it did, either because it caused the effect to occur in a different region or manner or because it did not cause the effect to occur at all. Second, each event must prevent at least one such possibility from being made actual. The first event must prevent some possible outcome of the second event and the second must prevent some possible outcome of the first.

Notice that we can think of these possibilities in terms of world states. In so far as one event makes a difference to what another event causes, the former event must also make a corresponding difference to the state of the world at the time that the effect of the latter event occurs. For example, if Billy's throw makes it the case that Suzy's throw causes a green spot to appear rather than a yellow spot, it must be nomically possible that Suzy's throw could have lead to the obtaining of a

state in which a yellow spot appears. Likewise, when Billy's throw prevents the appearance of the yellow spot, it necessarily prevents the obtaining of such a state (at least at the relevant time). Framed this way, the two previous conditions on additive causation ultimately concern the states that may obtain given an event's occurrence and the states that an event prevents from obtaining. It must be nomically possible that each additive cause could have been followed by a world state in which their effect does not occur when, where, and how it does. Moreover, each additive cause must prevent the obtaining of at least one such state.

It will be helpful to think of states as a species of event. Conceived as such, a state is the largest event to occur at any given time. Moreover, I'll suppose for convenience that states are both maximally fine-grained and fragile. As such, the state of the world right now is the event that exhaustively fixes the intrinsic character of the world at this particular moment. It could not have obtained a moment earlier or later, nor could it have obtained if the intrinsic properties of any part of the world just now had been different.

Whether a pair of events satisfy the above conditions depends on their contributions. An event's contribution helps determine both what states are nomically possible given the event and what states the event prevents from obtaining. Since an event's contribution is the set of worlds that are nomically possible given its occurrence, a world state is nomically possible given that event only if the event's contribution includes a world in which that state obtains. In contrast, an event prevents a world state from obtaining only if its contribution *does not* include a world in which that state obtains.

To better see this last point, consider an event e and a possible world state S . Suppose that e 's contribution includes a world in which S follows e . Now under determinism that is the only nomically possible world in which S obtains. For if

determinism is true, the state of the world at one time entails the state of the world at every other time. Every nomically possible world in which S obtains shares all of its states with every other world in which S obtains. These worlds differ in name only. That being the case, e cannot prevent S from obtaining. Far from impeding S 's obtaining, e is its necessary precursor. So while S may very well fail to obtain in the wake of e , it does not fail to obtain *because* of e . If e is to have any claim at all to preventing a state from obtaining then, its contribution cannot include a world in which that state obtains.

Putting this together allows us to spell out a necessary condition on causal additivity in terms of contributions. Should a pair of events interact additively, they must satisfy the following condition.

Constraint Events c and c^* are additive causes of an event e that occurs at time t only if there are a pair of worlds w and w^* such that (1) states obtain at t in w and w^* in which e does not occur in its actual region and manner, (2a) the contribution of c includes w but not w^* , and (2b) the contribution of c^* includes w^* but not w .

In short, if a pair of events are to interact additively, their contributions must cross-cut one another. The contribution of each event must include some world that the other does not. That world must be one in which the additively caused effect does not occur when, where, and how it in fact does.

The contributions of the vandals' throws bear this out. Imagine that Billy and Suzy both throw their paint at time t_1 and the green spot appears on the wall shortly thereafter at time t_2 . Though the vandals throw together, they each could have thrown alone. The contribution of Suzy's throw includes a world in which she throws alone at t_1 and the wall bears a yellow mark at t_2 . The contribution of Billy's throw includes a world in which he throws alone at t_1 and the wall bears a

blue mark at t_2 . The state of each of these worlds at t_2 is incompatible with the wall's bearing a green mark at t_2 ; they are not states in which a green mark appears. The contribution of Billy's throw excludes the world in which Suzy throws alone and the contribution of Suzy's throw excludes the world in which Billy throws alone. As such, the vandals' throws easily satisfy *Constraint*. Their contributions cross-cut one another in the manner that additivity requires.

To be sure, there is more to additivity than satisfying *Constraint*. *Constraint* only provides a necessary condition on causal additivity (and a weak one at that). It is by no means a sufficient condition. *Constraint* does not demand, and its satisfaction does not entail, that there is any causal relationship at all between the events that satisfy it. However, this is just to say that *Constraint* sets the bar for additive causation quite low. There are surely more exacting demands that we could place upon additive causes if that was our aim. But a low bar is a bar all the same. If a pair of events are to be additive causes, they must measure up.

3.5 Solving the exclusion problem

We're now in a position to address the causal exclusion problem. Recall the challenge we're confronted with. We need to account for the failure of additivity found in the interactions of mental events and their realizers. To do that, we need to explain how a pair of events could bring about the same effect without doing so additively.

One of the lessons to be had from *Constraint* is that causal additivity is in part a matter of how the contributions of additive causes are related. Some interactions may not be additive simply because the contributions of the relevant events stand in the wrong relation to one another. Any relation that precludes the contributions of a pair of events from crosscutting one another also precludes those events from

interacting additively. Take identity for example. Events that make the same contribution clearly cannot be additive causes. Since no world is both included and not included in the same contribution, events that make the same contribution will always run afoul of *Constraint*. Neither event makes a contribution that excludes a world found in the contribution of the other. This is one reason why no event can additively cause an effect with itself. It cannot, in principle, satisfy *Constraint*.

More importantly, causal additivity similarly falls prey to the containment relation — the relation that a set stands in to its subset. If the contribution of one event contains the contribution of another event, every world in the latter contribution is also found in the former contribution. The former, containing contribution does not exclude any world that belongs to the latter, contained contribution. And that's enough to preclude those events from being additive causes.

For example, suppose we have a pair of events *Super* and *Sub* where *Super*'s contribution is the set of worlds $\{w_1, w_2, w_3\}$ and *Sub*'s contribution is the set of worlds $\{w_1, w_2\}$. *Super* and *Sub* partially satisfy *Constraint*. *Sub*'s contribution excludes world w_3 from *Super*'s contribution. Even so, they cannot be additive causes. Because *Super*'s contribution contains *Sub*'s contribution, *Super*'s contribution does not exclude any of the worlds that belong to *Sub*'s contribution. So *Super* and *Sub* do not fully satisfy *Constraint*. *Constraint* demands that the contributions of both additive causes exclude a world that belongs to the other. Containment assures that this condition is not met.

This result is noteworthy because of the close relationship between contribution and supervenience. Should one event supervene on another with at least nomic necessity, the contribution of the former event will contain the contribution of the latter event. Suppose that an event a supervenes on an event b with nomic

necessity. a 's supervening on b consists in part in the upward necessitation of a by b ; in every world that shares our laws, a occurs if b occurs. As such, a 's contribution contains b 's contribution. Every nomically possible world in which b occurs is a world in which a occurs.

This suggests a simple solution to the exclusion problem. Mental events supervene on physical events. That is, necessarily, if a mental event m occurs, some physical event p also occurs such that, necessarily, if p occurs then m occurs.⁴ The strength of those necessities is controversial of course. Physicalists hold that mental events supervene on physical events with metaphysical necessity. Dualists hold that it is only nomic necessity. But most everyone agrees that mental events supervene on physical events with *at least* nomic necessity. If that's the case, the contribution of a mental event contains the contribution of the physical event that necessitates it. Those events cannot be additive causes. Should a mental event and a physical event that necessitates it overdetermine an effect, the failed additivity that this engenders is not problematic. On the contrary, it is to be expected. The contributions of the overdetermining causes would not allow their interactions to be additive.

Unfortunately, the simple solution is a bit *too* simple. The inconvenient fact of the matter is that the exclusion problem does not only arise for mental events and the physical events that necessitate them. We can find examples of this even in paradigmatic cases of mental causation. We often wish to say that mental events are realized by rather localized events in the brain whose constitutive properties are intrinsic to the regions in which they occur. For example, suppose we indulge the philosopher's fiction that pains are realized by c-fiber firings. Imagine that you burn your hand on the stove and reach for the freezer to get ice for the pain. Your

⁴That is, mental events strongly supervene on individual physical events (Kim, 1993a).

reaching is purely overdetermined by your feeling of pain and the c-fiber firing that realizes it. But your pain is not necessitated by the firing c-fibers. As Karen Bennett (2003) has vividly pointed out, a c-fiber firing that realized your pain could have occurred in a petri dish if the conditions were right. But c-fiber firings in petri dishes are not accompanied by pains. So the c-fiber firing that realizes your pain does not itself necessitate your pain. It realizes your pain only under the right circumstances. As such, the contribution of the pain event does not contain the contribution of the c-fiber firing.

Examples like this show that the simple solution to the exclusion problem is inadequate. Causal additivity fails in some cases of mental causation because the contribution of the mental event contains the contribution of its realizer. But this cannot be the sole reason that mental causation leads to a failure of additivity. Even so, I believe that the simple solution is *almost* correct. The contributions of mental events and their realizers are indeed responsible for the failed additivity that arises from mental causation. However, to account for that failed additivity we need to understand those contributions in the circumstances in which they are made.

3.6 Failed additivity in the circumstances

The discussion of additive causation in the previous sections has been ambiguous between two different types of additivity. As we've said, additive causes play a distinctive difference-making role. If the interaction of a pair of events is additive, each event makes a difference to their common effect with respect to the other. However, there are different types of differences that an event might make and still play that role. In particular, we can distinguish the differences that an event makes independently of the circumstances in which it occurs from the differences

that it makes given those circumstances. In the latter case, an event makes a difference to its effect, holding fixed the fact that some further event or events occur. That is, it makes a difference between the world being one way rather than another, both of which are possible given the occurrence of these further events. The former case involves no such restriction. For example, suppose that you accidentally back your car over a glass bottle and two shards of glass puncture your tire. Without taking either puncture as given, it is reasonable to suppose that each puncture makes a difference to the tire's inflation. If neither puncture is taken as given, there is an open possibility that the tire remains inflated. It is reasonable to suppose then that each puncture prevents such a possibility from being actualized and thereby makes a difference to the tire's inflation. In contrast, neither puncture makes a difference to the tire's inflation given that the other puncture occurs. For given that one puncture occurs, it is no longer an open possibility that the tire remains inflated.⁵ Without such an open possibility to prevent, the other puncture cannot make the difference between the tire's remaining inflated and not remaining inflated.

Talk about difference-making in the context of causation typically concerns difference-making of this latter sort. The causal exclusion problem is no different. Mental events and their realizers are not sufficient for their effects entirely on their own. Rather, they are sufficient for their effects given the particular circumstances in which they occur. Recall the way in which we arrived at the problem in Section 1. We imagined that, *given its background conditions*, the physical event that

⁵This is an over-simplification, of course. By presenting the case this way, I am taking quite a lot of the circumstances surrounding the puncture as given. For example, I am taking it as given that there is no one around who will quickly patch one of the punctures before the tire deflates. Were that good-samaritan around, it may well be that the tire's remaining inflated remains an open possibility despite one of the punctures.

realized your decision to applaud was a sufficient cause of your applauding when, where, and how you did. This qualification is important. Had it not been for those background conditions (e.g., your possessing a certain sort of physiology) the event that realized your decision may not have caused you to applaud at all, much less caused to you to applaud precisely as you did. When we regard the realizer as a sufficient cause of your applause, we consider it in the context of the circumstances in which it occurs. Likewise, when your decision fails to make a difference to that effect with respect to its realizer, it fails to make a difference to what its realizer causes given those background conditions.

The failure of additivity that we need to explain in these cases is the failed additivity of mental events and their realizers in light of the conditions in which they occur. We need to show that mental events and their realizers do not interact additively given at least some of the events in their environment. Contributions can provide that explanation. If they are to do so, it will be helpful to extend the notion of contribution to include two different types of influence: individual contributions and collective contributions.

Say that the type of contribution that we have been concerned with thus far is the *individual* contribution of an event. The individual contribution of an event is an influence that the event has on future states of the world independently of any other events in its environment. Additionally then, say that the *collective* contribution of multiple events is the influence that those events have together strictly in virtue of their individual contributions. It is an influence that those events have on future states of the world independently of any further events in their environment. For example, the collective contribution of events *a* and *b* is the influence that *a* and *b* have on future states of the world independently of surrounding events *c*, *d*, and so forth. The account of contribution I offered in

Section 2 easily extends to capture this notion of collective contribution. As before, an event's individual contribution is the set of nomically possible worlds in which that event occurs in its actual region and manner. The collective contribution of multiple events is the intersection of the individual contributions of those events. It is the set of nomically possible worlds in which all of the contributing events occur.

Just as the contribution of a supervenient event contains the individual contribution of any event that necessitates it, its contribution will also contain the collective contribution of any events that jointly necessitate it. Suppose that a mental event m supervenes on a pair of physical events p_1 and p_2 . Since p_1 and p_2 together necessitate the occurrence of m , any world in which p_1 and p_2 occur is a world in which m occurs. As such, the contribution made by m will contain the collective contribution made by p_1 and p_2 . And it will do so even if it does not contain the individual contributions of p_1 or p_2 . It is entirely possible that p_1 and p_2 might each have individually occurred in m 's absence and that m may similarly have occurred in the absence of p_1 or p_2 . If so, the contribution of m does not contain the individual contributions of either p_1 or p_2 . But it does contain their collective contribution.

Here again, containment relations stand in the way of additivity. No supervenient event can additively cause an effect together with *all* of the events on which it supervenes. As we saw in Section 3, if a pair of events are additive causes, each event must make some difference to what the other event causes. By extension, if m , p_1 , and p_2 are all to be additive causes, m must make some difference to what p_1 and p_2 cause together. To make that difference, m 's contribution would need to exclude some world that belongs to the collective contribution of p_1 and p_2 . But m 's contribution contains the collective contribution of p_1 and p_2 . It does not, in principle, exclude worlds in their

contribution. So m , p_1 , and p_2 cannot be additive causes together and, more generally, supervenient events cannot be additive causes with all of the events that necessitate them. Their contributions simply do not allow it.

By the same lights, supervenient events cannot be additive causes with some of the events that necessitate them given the occurrence of the others. A pair of events that are additive causes with one another given some third event must each make a difference to what the other causes, holding fixed the fact that the third event occurs. For example, if m and p_1 are additive causes given that p_2 occurs, m must make some difference to what p_1 causes holding fixed the fact that p_2 occurs. It must prevent the obtaining of a state that is not only incompatible with some effect of p_1 occurring in the region and manner that it does, but one that that may obtain in the wake of *both* p_1 and p_2 . After all, m must make a difference to what p_1 causes when it occurs in an environment that includes p_2 . But m cannot prevent such a state from obtaining. Preventing such a state from obtaining would require the contribution of m to exclude a world from the collective contribution of p_1 and p_2 . Since m does not do that, it cannot make a difference to what p_1 causes given p_2 .

It is containment relations of this sort that lead to the failed additivity that we find in cases of mental causation. While mental events are not always necessitated by their realizers, they are in the very least necessitated by collections of events that include their realizers. This is the case if, as we saw in the example of pains and c-fiber firings, a physical event realizes a mental event under the conditions in which it actually occurs but might not have realized that mental event under different conditions. The conditions in which the realizer would have failed to realize the mental event are conditions in which one or more of these other necessitating events are absent. For example, we needn't suppose that a c-fiber

firing that occurs in a petri dish is intrinsically any different than one that occurs in your brain; it does not fail to realize a pain event on account of any fault of its own. The c-fibers that fire in the petri dish simply aren't in the right environment to realize a pain. Some event or events with which the c-fiber firing could have jointly necessitated a pain event are absent.

The contribution of a mental event therefore either contains the individual contribution of its realizer or the collective contribution of its realizer and some further collection of physical events in their environment. Each of these containment relations is enough to assure the failed additivity that we find in cases of mental causation. If the contribution of the mental event contains the contribution of its realizer, the mental event and its realizer cannot interact additively regardless of their surrounding events. If, instead, the contribution of the mental event contains the collective contribution of its realizer and some further events, the mental event and its realizer cannot interact additively given the occurrence of those further events. Either way, mental events and their realizers cannot interact additively in the circumstances in which they occur. The very conditions that enable the mental event to occur undermine its ability to be an additive cause alongside its realizer.⁶

It bears emphasis that there is nothing remarkable about mental causation lurking behind this result. We needn't suppose that mental events are second-rate causal players, that mental causation is merely "as-if" causation, or that it is somehow derivative of physical causation. On the contrary, the present solution rests on the assumption that mental and physical events exert fundamentally the same type of influence on the world. Each event makes a contribution and could, in principle, serve as an additive cause. The reason that mental events and their

⁶Karen Bennett (2003) raises a similar point.

realizers do not in fact act as additive causes lies in the fact that, at least given their physical background conditions, their contributions simply do not line up in a way that permits it. The contribution of a physical event poses a more severe constraint on the world than that of the realized mental event. It closes off the possibility that the state of the world will evolve in ways that the contribution of the mental event leaves open. But its contribution is not logically or metaphysically prior to (or more important than) the contribution of the mental event as a result.

In this respect, this picture of contribution provides a response to Kim's concern that mental events may be epiphenomenal simply for lack of available causal work. Kim's worry, remember, was that mental events may not do any causal work because the efficacy of their realizers left them with no causal work to do. If the physical realizer of a mental event is a sufficient cause of an event and its properties, it precludes that mental event from doing any causal work for that event. This worry assumes that mental events can only do causal work if they interact additively with their realizers. But this now seems false. Though what it means for an event to do causal work is not entirely clear, it would seem to keep with the spirit of that metaphor that at least one way for an event to do causal work is for it to make a contribution. If so, a mental event and its realizer may both plausibly do causal work that leads to a common effect without interacting additively. The mental event and its realizer each put in causal work in that each makes a contribution. The nature of that work — their contributions — is simply such that it does not lead to an additive interaction. Seen this way, the fact that the realizer is a sufficient cause of an effect and its properties may mean that there is no further work that (in some sense) the mental event *needs* to do for that effect. But it does not follow from this that the mental event does not in fact do

any work, nor that the mental event does not cause that effect.

3.7 The present view in context

It may be helpful to consider my position in the context of more familiar proposals aimed at solving the exclusion problem for mental causation. Quite a few philosophers have thought that the solution to the exclusion problem lies in the close relationship between the mental and the physical (e.g, Yablo (1992), Loewer (2001), Bennett (2003; 2008), Pereboom (2002), Shoemaker (2001), Paul (2007)). Proponents of these views hold that while mental events are not identical to their realizers, they are tied to those realizers through some form of dependence such as supervenience (Loewer, Bennett, Yablo), constitution (Pereboom, Shoemaker), or mereological overlap (Paul). Contrast this with paradigmatic cases of additive causal interactions. Billy's and Suzy's throws, the impacts of the golf balls on the window, and the causes we single out in other textbook interactions are metaphysically independent of one another. And so the thought goes that these dependence relations make a causal difference. Mental events and their realizers are too closely related to interact additively.

The solution I've offered here belongs in part to this tradition. Like these views, it holds that the solution to the exclusion problem begins with a close relation between numerically different events. The supervenience relations that tie mental events to their realizers assure that the contributions of mental events contain the contributions of their realizers and (perhaps) background conditions. Since containment provides a solution to the exclusion problem and supervenience assures containment, it is reasonable to say that supervenience provides a solution to the exclusion problem for mental causation.

Despite this, the present view breaks with this tradition in some obvious and

not so obvious ways. The most significant difference lies in the nature and role of the relation that obtains between the causal relata. Philosophers who claim that the solution to the exclusion problem lies in a close relation between events have generally regarded those relations as a form of ontological dependence. The dependence such events exhibit is at least in part a matter of what they *are* rather than a brute consequence of the laws. This ontic bent applies to their claims regarding supervenience just as readily as it does to the more obviously ontic constitution and mereological relations. While supervenience in its own right needn't carry any ontic import, those philosophers who have claimed that it provides a solution to the exclusion problem have generally restricted their claim to supervenience that holds with at least metaphysical necessity. Weaker, nomic supervenience is generally not considered up to the task.

This is by no means a coincidence. Ontic dependence offers an intuitively compelling solution to the exclusion problem. To appreciate its appeal, recall that the exclusion problem arises in cases where, for one reason or another, we're committed to the view that multiple events cause an effect, but one of those causes is in fact entirely sufficient for that effect and its properties independently of at least some of the others. As such, we can frame the exclusion problem as the problem of explaining how multiple events bring about an effect in a manner that is indistinguishable from its being brought about by a single event. Ontic dependence can seem to provide a simple account of this. The close, ontic connection between mental and physical events is often seen as being in some respect like identity. As it is often remarked, the mental is "nothing over and above" the physical. It's not implausible to think then that if multiple events are *ontologically* like a single event, they are *causally* like a single event. They bring about certain effects as through they were a single cause because in some respect

they are like a single cause.

The view I've offered here doesn't ask ontic dependence to play this role. Should mental events stand in such a relation to their realizers, it matters only in so far as it assures that the appropriate containment relations obtain between their contributions. Ontic dependencies between events may very well be sufficient for this, but they are not necessary. So long as one event supervenes on another with nomic necessity, the containment relation that obtains between their contributions provides a solution to any exclusionary concerns that they engender.

Like these traditional views, however, the key to the approach I've presented is a close, non-nomic relation. Containment is a logical relation. The contribution that an event makes is contingent on the laws that subsume it, but the containment relations that its contribution stands in to other contributions hold with logical necessity. In this respect, the present view does not break outright with these traditional approaches to the exclusion problem so much as it shifts the critical relation from more familiar relata (events) to those that have not garnered much attention until now (contributions). Making this shift allows us to capture the intuitions that have guided these other views while avoiding the central pitfalls and challenges that they face.

The basic difficulty for solutions to the exclusion problem based in ontic dependencies is that they need to make the case that their ontic answer speaks to a fundamentally causal question. Ontic dependencies don't themselves do ready justice to the idea that the mental is nothing over and above the physical.⁷ For

⁷Paul (2007) is a notable exception to this. Paul holds that property instances have parts and that instances of mental properties and their realizers share a part — a common causal core. So understood, the mental property instance itself is something over and above the instance of its realizer, but the efficacious part of the mental property instance is not. Much like the present view, Paul captures the claim that the mental event's causing a given effect is nothing over and

however close the relation linking mental events to their realizers, the two are not identical. Mental events are strictly and literally something over and above the physical events that realize them. So there is at least one respect in which they are not at all like a single cause. As such, any view that claims that mental events and their realizers are like a single event faces a challenge. Granting that the close relations they cite do in fact obtain, they need to both (a) make clear the respect in which multiple events resemble a single event and (b) show how that particular resemblance addresses the exclusion problem. Barring that, the analogy with identity is just talk, however picturesque or intuitively compelling it may be.

In contrast, the present view takes the intuition that the mental is nothing over and above the physical at face value. It allows us to square that claim with the irreducibility of mental events to physical events. Notice that the claim that the mental is nothing over and above the physical is ambiguous between two related claims: the claim that the mental is nothing *ontologically* over and above the physical and the claim that the mental is nothing *causally* over and above the physical. More often than not, philosophers who engage in “nothing over and above” talk have the former, ontic claim in mind. But in the context of the exclusion problem the interesting claim is causal. We want to know how multiple causes could *act* like a single cause. Now one way to argue for the causal claim is to start from the ontic claim. That is the approach that presents a challenge to views rooted in ontic dependence. Shifting our focus from events to causal contributions makes this step unnecessary. It allows us to address the causal claim without the ontic intermediate.

The contributions of mental events are strictly and literally nothing over and above the contributions of their realizers and background conditions.

above its realizer causing that effect.

Contributions are constraints that an event imposes on the way in which the state of the world develops and evolves. For one constraint to be nothing over and above another constraint is just for the former constraint to impose no further limitation given the latter constraint. For example, suppose we have a restaurant that, in typical fashion, requires its patrons to wear shoes and a shirt. This requirement is a constraint. It limits the range of people that may patronize the restaurant. Now suppose that the restaurant introduces an additional requirement: patrons must wear shoes. This new constraint is not identical to the previous constraint but neither is it anything over and above the previous constraint. It does not exclude anyone from patronizing the restaurant that was not already excluded by the previous constraint.

Suppose then that a mental event m supervenes on a physical event p . Supervenience assures that the contribution made by m contains the contribution made by p . Since m supervenes on p , every world in which p occurs is also a world in which m occurs and every world in p 's contribution also belongs to m 's contribution. Unless m could not have been realized by events other than p , their contributions differ. Certain worlds belong only to m 's contribution. These worlds correspond to ways in which the actual world may evolve that are permitted by the contribution of m , but not by p .

To be sure, there is a clear sense here in which m 's contribution *is* something over and above p 's contribution. Seen merely as sets, the contribution of m is larger than the contribution of p . It includes worlds that are not found in p 's contribution. None the less, to focus on contributions only as sets is to neglect the more central notion that contributions are constraints. The critical feature of a contribution is not the worlds that it includes, but those that it rules out. In this way, the size of a contribution qua set is inversely proportional to the severity of

that contribution qua constraint. So while the contribution of m contains the contribution of p — and is therefore a larger set — it is in fact a less severe constraint. The contribution of m does not constrain the way in which the world may evolve in any way that it is not also constrained by the contribution of p . Every world that is excluded by m 's contribution is excluded by p 's contribution. So as a constraint, the contribution of m is quite literally nothing over and above the contribution of p .

With respect to their contributions then, the mental is nothing causally over and above the physical. More importantly, however, mental events do not make a difference to which states of the world obtain in their wake over and above the difference made by their subvenient physical realizers in the circumstances in which they occur. And, as we've seen, accounting for that is sufficient to resolve the causal exclusion problem.

CHAPTER 4

Regularity as a form of constraint

4.1 Introduction

In a frequently cited passage from *Causation*, David Lewis cast doubt on the prospects of a successful regularity theory of causation:

“It remains to be seen whether any regularity analysis can succeed in distinguishing genuine causes from effects, epiphenomena, and preempted potential causes — and whether it can succeed without falling victim to worse problems, without piling on the epicycles, and without departing from the fundamental idea that causation is instantiation of regularities. I have no proof that regularity analyses are beyond repair, nor any space to review the repairs that have been tried. Suffice it to say that the prospects look dark.” (Lewis, 1973)

The landscape for theories of causation has changed significantly in the last 40 years. Active work on regularity theories of causation has largely been abandoned in favor of counterfactual theories inspired by Lewis’s work. This move away from the regularity theory has no doubt been due to any number of factors, the strength of Lewis’s own counterfactual theory surely being among them. But the sea change that has occurred is just as surely a reflection of the sentiment that Lewis captures in the paragraph above: the regularity theory is flawed and likely beyond repair.

There is perhaps a little irony in the fact that the motivation that Lewis offered for that move has in some respects come full circle. In the paragraph

quoted above, Lewis raises three problems for the regularity theory: it wrongly identifies cases of backwards causation, it confuses the causal relation with the relation of having a common cause, and it fails to distinguish the actual causes of an effect from preempted possible causes of that effect. As objections to the regularity theory, these problems retain their bite. As reasons to prefer a counterfactual theory of causation to a regularity theory, however, their merits are less cut and dry. Counterfactual theories of causation famously struggle with cases of preemption and (less famously) face their own problems with backwards causation.¹ Of Lewis' three complaints, only the regularity theory's failure to distinguish the causal relation from the relation of having a common cause — what Lewis calls the problem of epiphenomena — stands as a clear mark in favor of counterfactual theories of causation.

My aim in this paper is to reassess the relative merits of the regularity theory. I will argue that a suitably revised version of the regularity theory can systematically address the problem of epiphenomena and a cluster of related problems that have undermined prior approaches to the regularity theory. My argument is in this respect quite modest. I do not claim that the regularity theory — either the version I propose below or any other — is in fact true. Unlike some recent proponents of the regularity theory (e.g., Strevens, 2007; Baumgartner, 2008), I do not believe that the regularity theory can solve the problem of preemption without drawing on outside resources.² But truth is not the only virtue of a theory. Rather, my claim is that the regularity theory is *comparatively* more robust than it is widely given credit for. Though Lewis's (1973; 1979)

¹See Elga (2001). The problem Elga raises is arguably less severe than the problems that the regularity theory has with backwards causation. My point is merely that problems with backwards causation are not limited to the regularity theory.

²For an argument to this effect, see Maslen (2012).

counterfactual theory of causation is demonstrably false, it continues to do fruitful work by serving as the basis for more nuanced theories. The regularity theory warrants a similar place at the table.

4.2 The problem of epiphenomena and other piggybacking problems

Regularity theories of causation are rooted in the idea that causes are sufficient conditions for their effects given the laws. They need not be sufficient on their own; it is enough that they are parts of larger conditions that are themselves sufficient. To that end, I'll assume here that determinism is true; the state of the world at one time fixes the state of the world at all other times. We can bring out the problems facing the regularity theory with just this basic picture. Call it the Naive Regularity theory (NR):

NR: c causes e iff (a) c is a part of a condition C , (b) C is sufficient for e given the actual laws, and (c) neither C nor the laws are themselves sufficient for e .

For example, suppose that Casey is at bat. The pitch is thrown and Casey swings, hitting the ball into the left-field stands. NR correctly recognizes that Casey's swing (c) causes the ball to land in the stands (e). c is part of a lawfully sufficient condition for e : Casey swings & the ball travels to the plate with a certain momentum (b) & nothing is present that would interfere with the ball's trajectory to the plate, with Casey's swing, or the ball's flight into the stands

$(\neg i)$.³ ⁴ Given the laws, $c \& b \& \neg i$ is sufficient for e , but neither the laws nor $c \& b \& \neg i$ are sufficient for e themselves.

Trouble comes quickly for NR. Suppose that the crowd erupts into a cheer (r) as they watch the ball sail out towards the stands. Though NR correctly recognizes that c causes e , it wrongly entails that r causes e . There are only so many possible ways that r could have been caused. Suppose for convenience that the only other event that could have caused the crowd to cheer at that moment would have been the announcement that everyone present had just won a new car (a). Since the crowd cheered, but were not told that they had won a new car, Casey must have swung. Given the laws then, $r \& \neg a$ is a sufficient condition for c . We can use this fact to derive new sufficient conditions for e . If $c \& b \& \neg i$ is sufficient for e given the laws, then so is $r \& \neg a \& b \& \neg i$. But despite satisfying NR, the crowd's cheer does not cause the ball to land in the stands.

This is the problem of epiphenomena. For NR, it entails that any event causes any other event so long as they have a common cause. And provided we're willing to follow the causal history of a pair of events back far enough, that will often be the case. We can mitigate the problem somewhat by requiring causes to precede their effects. This would block the result that, for instance, the ball's landing in the outfield was a cause of the crowd's earlier cheering, but that is small comfort at best.

³We can suppose that this condition is sufficient in any case. A truly sufficient condition would need to specify additional factors such as the left field stands being in place at the time Casey swings, their failure to be moved or destroyed, the pull of gravity, etc.

⁴By lumping together the absences of any event that would interfere with Casey's swing and the ball's trajectory to and from the plate, I do not mean to suggest that there is a single, monolithic absence that is essentially the absence of any interference as such. Rather, $\neg i$ is simply a catch-all for the myriad absences of events that, as it happens, would have interfered with the course of events (e.g., the absence of metal dome around the playing field at the relevant time).

The problem of epiphenomena is an example of what I'll call a piggybacking problem. A theory of causation suffers from a piggybacking problem in so far as it mistakenly recognizes some entity as a cause of an effect due to the relation that the entity stands in to an actual cause of that effect. The misidentified cause rides piggyback on the legitimate causal status of another entity. For instance, in the previous case the crowd's cheer was counted as a cause of the ball landing in the stands because it piggybacked on the causal status of Casey's swing.

The regularity theory has been hard hit by such problems. Though the problem of epiphenomena is a particularly high-profile piggybacking problem, it is hardly the most embarrassing. The piggybacking involved in the problem of epiphenomena depends on the natural laws. The crowd's cheer and the absence of an announcement that everyone had won a car are jointly sufficient for Casey's swing only because there is a lawful connection between the swing and the cheer. Absent that connection, the problem could not arise. Piggybacking does not always rely on the laws in this way. In fact, it's much worse when it does not. Consider two such piggybacking problems:

- *The piling-on problem.* Minor restrictions aside, NR holds that an event c causes an event e just in case c is part of sufficient condition for e . Under NR, an event can piggyback on the causal status of another event simply by being part of the same condition. If A is sufficient for C , then $A \& B$ is also sufficient for C , for any arbitrary A , B , and C . It follows from NR then that unless an event is entirely uncaused, it is caused by most any event at all. By finishing the previous sentence, you have caused the fall of the Roman empire, the election of the next president of the United States, and every winning goal of every World Cup match ever to be played.
- *The gerrymandering problem.* The mechanics of the piling-on problem are

fairly crude. We can derive the same result in a more subtle way. In the problem of epiphenomena, piggybacking involved replacing part of a sufficient condition with new elements that were sufficient for the replaced part; we derived a new sufficient condition $(r \& \neg a \& b \& \neg i)$ by substituting r and $\neg a$ for c in $c \& b \& \neg i$ rather than simply appending them to $c \& b \& \neg i$. We can do the same for arbitrary events. Take an arbitrary event from the Casey example, such as the crowd's cheer (r). Since Casey swings (c), it's true that either the crowd does not cheer or Casey swings ($\neg r \vee c$). Since $c \& b \& \neg i$ is sufficient for e given the laws, so is $r \& (\neg r \vee c) \& b \& \neg i$. So once again it follows on NR that the crowd's cheer causes the ball to land in the stands. The fact that c causes r , however, is irrelevant.

The fact that these problems arise for NR is not itself cause for alarm. To my knowledge, no one has ever seriously defended NR. As a theory of causation, it is clearly too simple. But this just makes it all the more troubling that more sophisticated approaches to the regularity theory do not fair significantly better.⁵

Perhaps the most successful approach to the regularity theory is J. L. Mackie's (1965; 1974) INUS theory of causation. Mackie's theory holds that an event c causes an event e just in case c is (at least) an insufficient but necessary part of an unnecessary but sufficient condition for e — an INUS condition. In so far as piggybacking problems are concerned, the important aspects of Mackie's analysis are the condition's sufficiency and the part's necessity. A necessary part of a sufficient condition is a part without which the condition would not be sufficient.

⁵Strevens (2007) and Baumgartner (2008) both offer versions of the regularity theory aimed at addressing these problems. I have reservations concerning both of their proposals. For present purposes, however, I only note that *if* they are successful, that success is the exception, not the rule.

For instance, if $A \& B$ is a sufficient condition for C , A is a necessary part of that condition just in case B is not itself sufficient for C . I'll call this simplified view 'the Mackie theory':

c is a cause of e iff (a) c is part of a condition C , (b) C is sufficient for e given the actual laws, and (c) no proper part of C is sufficient for e given the actual laws.

The Mackie theory neatly avoids the piling-on problem. If a condition is sufficient for an effect, any new condition that we derive from it by introducing additional members will also be sufficient for that effect, but cannot be minimally sufficient for that effect. However, this appeal to minimal sufficient conditions does nothing to address either of the other piggybacking problems. Both the gerrymandering problem and the problem of epiphenomena arise because we can derive new sufficient conditions for an effect by replacing parts of a sufficient condition for that effect with proxies that are sufficient for the parts they replace. For example, we can replace c in $c \& b \& \neg i$ with $r \& (\neg r \vee c)$ because $r \& (\neg r \vee c)$ is sufficient for c . Because these proxies are introduced through replacement rather than being crudely tacked on to the original sufficient condition, they needn't be redundant. And so long as they are not redundant, they may be parts of minimal sufficient conditions.

The Mackie theory's inability to cope with these latter two piggybacking problems suggest that its success with the piling-on problem is at least somewhat ad hoc. All three problems arise for the regularity theory because events that are irrelevant to a condition's sufficiency may be recognized as causes by attaching themselves to it. In the case of the piling-on problem, this attachment is explicit; the piggybacking events are directly appended to an already sufficient condition. In the latter two problems, this attachment is left implicit. Though

$r \& (\neg r \vee c) \& b \& \neg i$ does not explicitly state that c occurs, its sufficiency for e is predicated on the fact that c occurs. $r \& (\neg r \vee c) \& b \& \neg i$ is sufficient for e because it is sufficient for $c \& b \& \neg i$. The Mackie theory's appeal to minimal sufficient conditions is enough to avoid the piling-on problem. But it does not address this more general issue of piggybacking that lies behind it.

Even so, the Mackie theory is surely right to distinguish conditions that are, in some sense, minimally sufficient for an outcome from those sufficient conditions that include unnecessary hangers on. But drawing that distinction in terms of the redundancy of a condition's parts — in effect, in terms of its syntax — is not an adequate way to go about this. In what follows, I'll present a different strategy for distinguishing minimal and merely sufficient conditions. Rather than building a theory on the back of conditions that are minimal with respect to their syntax, we ought to build a regularity theory on conditions that are minimal with respect to the demands they place on the world.

4.3 Contribution

One virtue of the regularity theory as it has developed since Mill is that it gives pride of place to the fact that the causal interactions of ordinary events are a collective affair. Very few events cause anything at all by themselves. Though we may say, for instance, that your flipping the light switch was *the cause* of the lights coming on, it is by no means the case that your flipping the switch was the *only* event that caused the lights to come on when they did or that it managed to do so irrespectively of the other events in its environment. The presence of proper wiring in the switch, functioning bulbs in their sockets, and other equally forgettable events also had a role to play. The regularity theory reflects this interdependence. It frames causes as parts of sufficient conditions for their effects. Those parts earn

their causal stripes together or not at all.

In doing so, however, the regularity theory ignores a related intuition. Though what an event causes depends on the other events in its environment, it's very natural to think that events have some influence on the world in their own right, independently of other events. For example, suppose that we each add a drop of dye to a small cup of water. You add a drop of blue dye. I add a drop of yellow dye. In short order, the water turns green. Now as a matter of fact, your adding the drop of blue dye to the water caused it to turn green. But it might not have. If I had instead added red dye or no dye at all, your action would have had a very different effect. In this respect, the outcome of your action — what it causes — is at the mercy of my action. It depends on what I do. None the less, it seems to me that there is an intuitive sense in which the influence of your action is altogether independent of mine. Had I not added any dye to the water, your action would have had a different effect, but its influence, in some sense, would have been the same. The point here is perfectly general. Just as the influence of your action, in this sense, is independent of my action, it is also independent of any other event in its environment. I'll refer to such influence as the contribution of an event. So understood, an event's contribution is the influence that it has on future states of the world given only its nature and the laws.

What influence is that? I've argued elsewhere that we should understand contribution as a kind of constraint.⁶ The contribution of an event is the constraint it imposes on how the state of the world may evolve in its wake. Suppose that an event e occurs at time t_1 . Whatever state of the world later obtains at time t_2 must conform not just to the laws, but to the occurrence of e . In this respect, the range of states that are eligible to obtain at t_2 given the occurrence of e is smaller

⁶See my *Causal contribution and causal exclusion*.

than the range of states that are eligible to obtain at t_2 simpliciter. Simply given its nature and the laws, e constrains what the world may be like at t_2 .

More formally, we can understand contributions as sets of possible worlds. An event's contribution is the set of nomically possible worlds in which it occurs. For example, the contribution of my adding a drop of yellow dye is just the set of worlds that share our laws in which I add that drop of yellow dye. So understood, an event's contribution serves as a constraint on the way in which the state of the world may evolve. Each nomically possible world marks one way in which the state of the world can evolve under the actual laws. Taken together, they mark the ways in which the state of the world can evolve given that the event occurs from those ways in which it cannot. If it is nomically possible for a given state to obtain in the wake of an event, there must be some world in the event's contribution in which that state obtains in the wake of the event. If the event's contribution does not include such a world, the state of the world may not evolve forward from the time that the event occurs such that the state in question obtains. The event exerts an influence on the world in that it permits the state of the world to evolve in some ways and prohibits it from evolving in others.

Going forward, it will be helpful to extend this basic idea in two ways. First, grant that contributions can be made both by events and omissions. The contribution of an omission is the set of nomically possible worlds in which the omission "occurs". It is admittedly a bit odd to think of an absence as exerting an influence on the world or constraining what the world may be like in the future. But this is the same peculiarity that comes from thinking of absences as causes and effects. If omissions can be causes — and it would seem that they can — then they can make contributions.⁷ Along these lines, I'll assume that omissions are

⁷Alternatively, if you prefer to think that omissions are only second-class causes (e.g., Dowe

largely event-like in that they can be causes and effects and are localized to particular spacetime regions.

Second, if events exert an influence on the world individually, it is natural to think that they exert a similar influence collectively. Their collective influence is just the influence that they have together in virtue of the influence that they each have on their own. To reflect this, it will be helpful to distinguish two types of contributions: individual contributions and collective contributions. An individual contribution of an event or omission is the influence that it has on the world simply given the laws. These are the contributions described above. In contrast, the collective contribution of multiple events or omissions is the influence that they have on the world *together* simply given the laws. Say that the collective contribution of multiple events or omissions is the intersection of the individual contributions of those events; it is the set of nomically possible worlds in which all of the contributing events occur. So for a pair of events a and b , the individual contribution of a is the set of all nomically possible a -worlds, the individual contribution of b is the set of all nomically possible b -worlds, and the collective contribution of a and b together is the set of all nomically possible $a\&b$ -worlds.

4.4 Minimal sufficiency two ways

There is a direct connection between contribution and sufficiency under the laws. An event's contribution is a constraint on the way in which the state of the world may evolve under the laws. By constraining that evolution in the right way, a contribution may assure that certain events occur. For example, suppose that the (2000) or that their causal status is derived from events (e.g., Lewis (2004)), then presumably they may make contributions in an analogous fashion. My point is only that omissions are no more poorly suited to making contributions than they are to being causes.

collective contribution of some events only includes worlds in which an event e occurs. Since the contribution assures that one of these e -worlds will be actualized, it also assures that e will occur. We can use this connection between contribution and lawful sufficiency to reframe our notion of sufficient conditions. When the conditions in question are just collections of events and omissions, a condition C is sufficient for an event e given the laws iff every world in the contribution of C 's parts is an e -world.

This is a richer notion of sufficient conditions than regularity theories like Mackie's and NR have employed. Like Mackie, we can distinguish minimal sufficient conditions from mere sufficient conditions by the redundancy of their parts. Call sufficient conditions that are minimal in this respect 'Mackie-minimal sufficient conditions'. We can also recognize a different sort of minimal sufficient condition, one distinguished by the *severity* of the constraint it imposes on the world. Constraints come in different degrees. A carnival that requires ride-goers to be at least five feet tall imposes a more severe constraint on would-be riders than one that only requires that they be at least four feet tall. It allows for less variation in riders' heights. It permits fewer people to ride the rides. We can gauge the severity of a contribution in a similar fashion. A severe contribution permits the state of the world to evolve in a more limited range of ways than one that is comparatively less severe. The more severe the contribution, the fewer worlds it contains.

A condition is minimally sufficient for an event in this latter sense in so far as it imposes a constraint that is severe enough to assure that the event occurs without being any more severe than necessary to assure that the event occurs. Call such conditions *constraint-minimal sufficient conditions*. I'll understand constraint-minimal sufficient conditions as follows. For convenience, say that the

contribution of a condition is the contribution of the events and omissions that constitute that condition.

C is a constraint-minimal sufficient condition for an event e iff (a) every world in the contribution of C is an e -world and (b) there is no condition D such that (i) every world in the contribution of D is an e -world and (ii) the contribution C is a proper subset of the contribution of D .

So construed, constraint-minimal sufficient conditions are sufficient conditions in that their contributions assure that certain events occur. They are minimally sufficient in that those contributions are not proper subsets of larger, less-severe contributions that would also assure that those events occur.⁸ We cannot relax the contribution of a constraint-minimal sufficient condition, as it were, and arrive at a new condition that is sufficient for the same event.

Constraint-minimalism is not just Mackie-minimalism in new clothes. Constraint-minimalism and Mackie-minimalism are logically independent properties. That a condition is constraint-minimally sufficient for some event is neither necessary nor sufficient for it to be Mackie-minimally sufficient for that event. This comes out easily enough when we consider cases that involve pairs of determinate and determinable properties and, by extension, their corresponding events. For example, imagine a point in a bullfight at which the bull is so angry that the matador's waving a red cape assures that the bull will charge. As it

⁸To be clear, constraint-minimal sufficient conditions are *not* minimal in that their contributions are necessarily less severe than those of any other sufficient condition for a given event. For example, if worlds w_1 , w_2 , and w_3 are all e -worlds, a set of events that makes the contribution $\{w_1\}$ and a set of events that makes the contribution $\{w_2, w_3\}$ may both be constraint-minimal sufficient conditions for e despite the fact that the latter contribution is less severe than the former.

happens, the matador's cape is not merely red (nothing is after all) but is in fact *scarlet*. When the matador waves the scarlet cape, the bull charges. Now for the sake of our example, let's assume that waving the scarlet cape and waving the red cape are different events. Call them the scarlet waving and the red waving.

With that in mind, consider the following three conditions:

- A. the red waving occurs & the bull is angry
- B. the scarlet waving occurs & the bull is angry
- C. the red waving occurs & the scarlet waving occurs & the bull is angry

Let's pretend that each of these conditions is lawfully sufficient for the bull's charge. Both A and B then, but not C, are Mackie-minimal sufficient conditions for the charge. However, only A is a constraint-minimal sufficient condition for the charge. The contribution of B is a proper subset of the contribution of A. The matador could have waved a red cape without waving a scarlet cape, but not a scarlet cape without a red one. So given that A is sufficient for the charge, B's contribution is needlessly severe. Suppose, however, that the bull is not so easily provoked. Rather than charge at the waving of any old red cape, he charges only when scarlet capes are waved. In this case, B is the only Mackie-minimal sufficient condition for the charge, but both B and C are constraint-minimal sufficient conditions for it. For though B is a proper part of C, they make the very same contribution.

That constraint-minimal sufficiency and Mackie-minimal sufficiency come apart in these cases points to a substantive difference between these two forms of minimalism. They ultimately concern two quite different kinds of properties of conditions. As I noted previously, Mackie-minimal sufficiency is ultimately a minimalism about syntax.⁹ It is concerned with what we must explicitly specify

⁹My use of syntax-talk is meant to be ontologically innocent. I do not assume that conditions

about the world in order to assure a certain outcome. (This is why C cannot be a Mackie-minimal sufficient condition for the bull charging. It says too much. There is no need for the condition to specify that both the red waving and the scarlet waving occur when one will do just as well.) In contrast, constraint-minimal sufficiency is unconcerned with being tight-lipped. It is not a measure of the condition's size, but rather its influence. It is because of this that a condition that includes just a single event may not be minimal enough and yet a condition that includes redundant members may be just right.

4.5 Adapting the regularity theory

Constraint-minimal sufficient conditions provide the basis for a novel version of the familiar regularity theory. Call this view the constraint regularity theory (CR):

An event c is a cause of an event e iff (i) c and e are distinct events, (ii) c occurs prior to e , and (iii) c belongs to a constraint-minimal sufficient condition at some time t for e .

CR adapts the regularity theory by trading Mackie-minimal sufficient conditions for constraint-minimal sufficient conditions. Specifically, it trades Mackie-minimalism for constraint-minimalism *at a time*. Constraint-minimal sufficiency is a comparative measure. Whether a condition is constraint-minimally sufficient for some effect depends in part on what conditions fall into its comparison class. For a condition to be constraint-minimally sufficient at a time is for it to be constraint-minimally sufficient relative to a comparison class that is restricted with respect to time. Say that a condition C is a constraint-minimal sufficient condition at time t for some effect just in case the events in C occur at t

are linguistic or propositional entities.

and C is a constraint-minimal sufficient condition for the effect with respect to those conditions whose constitutive events also occur at or prior to t .

To illustrate, let's revisit the Casey example from 1. In our example, Casey's swing causes the ball to land where it did in the left-field stands. We noted before that Casey's swing is part of a lawfully sufficient condition for the ball's landing in the stands: Casey swings (c) & the ball is traveling towards the plate with a certain momentum (b) & nothing is present that would interfere with the ball's trajectory to the plate, with Casey's swing, or the ball's flight into the stands ($\neg i$). Suppose that c , b , and $\neg i$ all take place at time t . So $c \& b \& \neg i$ is a prime candidate for being a constraint-minimal sufficient condition at t for e . Moreover, e is distinct from any of the events in $c \& b \& \neg i$ and c precedes e .

According to CR then, Casey's swing causes the ball to land in the stands unless there is a condition with a better claim than $c \& b \& \neg i$ to being constraint-minimally sufficient for e at t and, moreover, c is not a part of that condition. Consider what such a condition would be like. To preclude $c \& b \& \neg i$ from being a constraint-minimally sufficient for e at t , a condition would need to make a contribution that is relevantly less severe than the contribution of $c \& b \& \neg i$ — the contribution of $c \& b \& \neg i$ must be a proper subset of this contribution. This will only be the case if every nomically possible world in which we find c , b , and $\neg i$ together is a world in which we find the constitutive parts of this other condition but not vice versa. In short, this other condition must be necessary but not sufficient for $c \& b \& \neg i$ given the laws, its parts must take place at t , and it must not include c . And of course it must still be sufficient for e .

There is no in-principle prohibition against such a condition. But the most promising candidates do not threaten CR. The obvious candidate conditions are those that replace c in $c \& b \& \neg i$ with some event that supervenes on c , b , or $\neg i$

either individually or collectively. For example, we might distinguish Casey's swing (*c*) from a coincident event, call it *the* swing, on the grounds that the former but not the latter is essentially performed by Casey. Or perhaps we might distinguish both these events from *moving the bat towards the ball* in that the swing was just one way of moving that bat towards the ball, any of which would have resulted in the ball landing in the stands. That is, we might find a more general, determinable event that can take *c*'s place in $c \& \& \neg i$ without rendering the new condition insufficient for *e*. If there are such events, CR entails that they (or their still more determinable contemporaries) cause *e* and *c* does not.

So be it. If Casey's swing, the swing, the moving of the bat, and the like are in fact different events — and not just different descriptions of the same event — the latter are very plausibly better candidate causes of *e* than Casey's swing. They lack the specificity of Casey's swing, specificity that would be irrelevant to bringing about *e*. CR correctly identifies at least one of these other events as a cause of *e* in place of *c*. As such, CR delivers the correct result in our example. Either Casey's swing causes *e* and our original intuition is confirmed or Casey's swing does not cause *e*, but only because a closely related and, on reflection, better-suited event does. Either result does justice to our original intuition that Casey's swing causes the ball to land in the left-field stands.

4.6 Piggybacking problems

We can now return to the piggybacking problems. Like the Mackie theory, CR solves the piling-on problem faced by NR. But unlike the Mackie theory, CR also solves the problems of gerrymandering and epiphenomena. Under CR, ordinary instances of these problems are fundamentally the same. The problems of gerrymandering and epiphenomena are simply special cases of the piling-on

problem. Let's take the problems one at a time.

4.6.1 The piling-on problem

The piling-on problem arose for NR because we can generate arbitrary sufficient conditions for an event simply by appending additional elements to an already sufficient condition. The Mackie theory solved this problem by denying that conditions that could be generated in this way were relevant to causation. By definition, these conditions would have redundant parts and the Mackie theory denies that conditions with redundant parts are causally relevant. The piling-on problem is not quite so straightforward for CR. As we've seen, constraint-minimal sufficient conditions may very well have redundant parts. So by trading a theory based on Mackie-minimal sufficient conditions for one based on constraint-minimal sufficient conditions, we give up Mackie's tidy solution to the piling-on problem. On CR, the mere fact that a condition has redundant parts does not in and of itself tell against its relevance to causation.

What does tell against such a condition on CR, however, is its *arbitrary* redundancy. While constraint-minimal sufficient conditions may include redundant parts, they may not include arbitrary redundant parts. For example, let's suppose that Casey's swing does in fact cause the ball to land in the stands in part because $c \& b \& \neg i$ is a constraint-minimal sufficient condition for e at time t . For the piling-on problem to undermine CR, it must be possible to generate new constraint-minimal sufficient conditions for e at t by appending arbitrary events (or omissions) that take place at t to $c \& b \& \neg i$. But that won't work. Remember, the contribution of a condition is the intersection of the contributions of its parts. When we append an arbitrary event z to $c \& b \& \neg i$, the contribution of this new, expanded condition $z \& c \& b \& \neg i$ is the intersection of the contributions of $c \& b \& \neg i$

and z . By definition, this contribution is a subset of the contribution $c \& b \& \neg i$. It will be a proper subset unless z is at least nomically necessary for $c \& b \& \neg i$ — unless every nomically possible world in which $c \& b \& \neg i$ obtains is a world in which z occurs. And that simply won't be the case for just any arbitrary event that we might tack on to $c \& b \& \neg i$. So while we can generate new sufficient conditions for an effect by arbitrarily padding a constraint-minimal sufficient conditions for that effect with extra events, we cannot generate new constraint-minimal sufficient conditions in that way.

Of course, this does mean that we can derive new constraint-minimal sufficient conditions by adding the right events to a constraint-minimal sufficient condition. I'll return to this issue in 6 below. For the moment, however, note that this is in fact a virtue of the theory. The events that we can safely add to a constraint-minimal sufficient condition are those that supervene on that condition with metaphysical or nomic necessity. For example, we can add events to $c \& b \& \neg i$ such as the conjunctive event $c \& b$ or the more general event *the moving of Casey's arms*. Since these events are both necessary for $c \& b \& \neg i$, adding them to $c \& b \& \neg i$ makes no difference to that condition's contribution. The collective contribution of either (or both) of these events and $c \& b \& \neg i$ just is the contribution of $c \& b \& \neg i$. So, because $c \& b \& \neg i$ is a constraint-minimal sufficient condition for e , these additional events also cause e according to CR. This is the correct result. It would be at best very odd to suppose that c causes e and b causes e , but the complex event $c \& b$ does not.

In any case, if it is a mistake to regard these events as causes, it is a mistake that CR shares with most any counterfactual account of causation. Counterfactual accounts of causation start from the claim that counterfactual dependence between (suitably distinct) events is sufficient for causation. An event c causes e if, had c

not occurred, e would not have occurred.¹⁰ As those counterfactuals are normally understood, if e depends on c , e will also depend on most any event that is nomically necessary for c , at least in so far as that event is a contemporary of c . For example, if Casey had not moved his arms, the ball would not have landed in the stands. And this is precisely because Casey's moving his arms is necessary for him to swing. So according to both CR and counterfactual theories of causation, Casey's moving his arms caused the ball to land in the left-field stands. This is enough to defuse the piling-on problem for CR. To the extent that it poses any problem at all, it is not a special problem.

4.6.2 Replacing and appending

Recall that the problems of gerrymandering and epiphenomena arise because we can derive new sufficient conditions for an effect by replacing one or more parts of a sufficient condition for that effect with new parts that logically or lawfully entail the parts that they replace. Because these new parts replace the old parts rather than augment them, Mackie's prohibition against redundant conditions doesn't speak to these problems. If we start with a Mackie-minimal sufficient condition for the effect, we can derive a Mackie-minimal sufficient condition for that effect.

The difference between replacing parts of a condition and appending additional elements to it doesn't carry the same significance under CR. Consider a toy case. Suppose that $a \& b$ is a constraint-minimal sufficient condition for an event e and that some further event c is at least lawfully sufficient for b . Because c is lawfully sufficient for b , we can derive new sufficient conditions for e either by replacing b with c ($a \& c$) or appending c to b ($a \& b \& c$). The Mackie theory treats these derived

¹⁰For an objection to the claim that counterfactual dependence is sufficient for causation, see Sartorio (2010). For a defense of that claim, see Weslake (2013).

conditions differently; one has redundant parts and the other does not. CR holds them equivalent. Since every nomically possible c -world is a b -world, the contribution of $a \& c$ just is the contribution of $a \& b \& c$. One is a constraint-minimal sufficient condition for e just in case the other is as well.

Under CR then, the problems of gerrymandering and epiphenomena are in an important respect just variations on the piling-on problem. The fact that they rely on replacing rather than appending to derive new sufficient conditions is irrelevant. The sufficient conditions that they derive via replacing will be constraint-minimal sufficient conditions only if the parallel conditions derived via appending are as well. As a result, our solution to the piling-on problem generalizes to these more formidable problems. The details may differ, but their solutions are fundamentally the same.

4.6.3 The gerrymandering problem

The gerrymandering problem arose for Casey because in so far as $c \& b \& \neg i$ was a sufficient condition for e , so was $z \& (\neg z \vee c) \& b \& \neg i$ for any arbitrary event z . Since $z \& (\neg z \vee c) \& b \& \neg i$ is a Mackie-minimal sufficient condition for e , the Mackie theory entails that z causes the ball to land in the stands. CR does not.

$z \& (\neg z \vee c) \& b \& \neg i$ is not a constraint-minimal sufficient condition for e — at least not for any arbitrary event z — for two reasons.

First, $z \& (\neg z \vee c) \& b \& \neg i$ is simply not the right kind of condition to be constraint-minimally sufficient. Constraint-minimal sufficient conditions are collections of events and omissions. Disjunctions of events or omissions are neither events nor omissions. So z cannot earn its causal stripes under CR in virtue of being a part of $z \& (\neg z \vee c) \& b \& \neg i$.

Second, and more importantly, even if constraint-minimal sufficient conditions

could include disjunctions of events or omissions, $z \& (\neg z \vee c) \& b \& \neg i$ would not qualify. Keeping with our treatment of the contributions of events and omissions, suppose that the contribution of a disjunction of events or omissions is the set of nomically possible worlds in which at least one of the disjuncts occurs. So understood, the contribution of $z \& (\neg z \vee c) \& b \& \neg i$ and the contribution of $z \& c \& b \& \neg i$ are identical. As with the piling-on problem, $z \& c \& b \& \neg i$ is not a constraint-minimal sufficient condition for e for just any arbitrary event z . Neither then is $z \& (\neg z \vee c) \& b \& \neg i$. Unless z is at least nomically necessary given $c \& b \& \neg i$, the contribution of $z \& (\neg z \vee c) \& c \& b \& \neg i$ is a proper subset of the contribution of $c \& b \& \neg i$. So CR may address the problem of gerrymandering in precisely the same way that it addresses the problem of piling-on.

4.6.4 The problem of epiphenomena

Finally, let's return to the problem of epiphenomena. Because Casey's swing (c) caused the crowd to cheer (r), r is sufficient for c given certain features of the environment, namely the absence of any other potential causes of r . In our case, this was the absence of an announcement that everyone had won a new car ($\neg a$). Since $r \& \neg a$ is lawfully sufficient for c and $c \& b \& \neg i$ is a Mackie-minimal sufficient condition for e , $r \& \neg a \& b \& \neg i$ is Mackie-minimally sufficient for e as well. So the Mackie theory entails that the crowd's cheer caused the ball to land in the stands.

CR addresses the problem of epiphenomena in much the same manner as it does the other piggybacking problems. As we've noted, the problem arises because effects may be lawfully sufficient for their causes when taken in conjunction with the appropriate context. For Casey, it is $r \& \neg a$ that is lawfully sufficient for c . This entailment allows $r \& \neg a$ to stand in for c : $r \& \neg a \& b \& \neg i$ is sufficient for e because $r \& \neg a \& b \& \neg i$ is sufficient for $c \& b \& \neg i$ and $c \& b \& \neg i$ is sufficient for e . But for that

very reason $r \& \neg a \& b \& \neg i$ is not constraint-minimally sufficient for e . The contribution of $r \& \neg a \& b \& \neg i$ is a proper subset of the contribution of $c \& b \& \neg i$. After all, it could have been the case that Casey swings (c) as the ball heads to the plate (b) in the absence of interference ($\neg i$) and it was announced that everyone won a new car (a). Such worlds belong to the contribution of $c \& b \& \neg i$ but not $r \& \neg a \& b \& \neg i$. Given that $r \& \neg a \& b \& \neg i$ is not a constraint-minimal sufficient condition at the relevant time for e then, CR does not regard r as a cause of e in virtue of being part of $r \& \neg a \& b \& \neg i$.

Now this is in some respects a moot point. By definition, constraint-minimal sufficient conditions at a given time must be constituted by events and omissions that occur at that time. $r \& \neg a \& b \& \neg i$ was therefore ineligible from the outset to be a constraint-minimal sufficient condition at any time for e ; its constituents do not all occur at any one time. Even so, it's difficult to see how any other condition that might give rise to the problem could fair better. To generate the problem of epiphenomena in the Casey example, we would need a collection of events and omissions that occur at the same time as r and, together with r , are lawfully sufficient for $c \& b \& \neg i$. Suppose that we have such a collection then.¹¹ Whatever our new candidate condition may be, if it is sufficient for $c \& b \& \neg i$, it seemingly must be because it is sufficient for $r \& \neg a \& b \& \neg i$. (After all, for r to assure that c occurs, our condition must establish that no event but c could have caused r .) But if that's the case, its contribution is a subset of the contribution of $r \& \neg a \& b \& \neg i$. And since contribution of $r \& \neg a \& b \& \neg i$ is a proper subset of the contribution of

¹¹Identifying such a collection is, to be sure, a tall order. Here is a start. Rather than pairing r with $\neg a$ to entail c , we could instead pair r with the absence, at the time r occurs, of any tell-tale reverberations in the air that would have been left behind from the announcement that everyone had won a car and the similar absence of any traces of those signs being prevented from being present.

$c \& b \& \neg i$, the contribution of our candidate condition will be as well. It cannot be a constraint-minimal sufficient condition for e at the relevant time.

At root then, the problem of epiphenomena is hamstrung under CR by the apparent fact that events can be brought about in multiple ways; there are multiple possible causal histories for any given cause.¹² Establishing a backwards entailment from an effect to its cause requires pairing that effect with a further range of events or omissions that prune away those possible causal antecedents until they assure that the actual cause of the effect occurs. As the Casey example illustrates, however, ruling out these possibilities doesn't come for free. It requires fixing various facts about the world that are not fixed by the cause itself. Relative to the contribution of the cause, the contribution of the would-be problem condition is needlessly severe.

4.7 Limitations of CR

CR provides a simple and uniform treatment for the three piggybacking problems as they ordinarily arise. It does not entail that arbitrary events are causally related or confuse causation with the mere having of a common cause. This is not, of course, enough to establish CR as a viable theory of causation. A viable theory of causation must minimally address the problems raised by redundant causation, the transitivity (or intransitivity) of the causal relation, and causation involving omissions; this paper does none of that. Even so, it suggests that CR might provide a novel and potentially elegant basis for a more fully-developed position.

¹²There may very well be exceptions. Notably, if the obtaining of a complete state of the world at a time is an event, it is an event with only a single nomically possible causal history. Be that as it may, we can ignore such all-encompassing events here. They cannot give rise to the problem of epiphenomena.

To that end, I want to close by flagging an important limitation of CR. As I noted at the end of the last section, CR's treatment of the problem of epiphenomena rests on the assumption that events have multiple possible causal histories; whatever the actual causes of an event may be, it might have been caused by something else. Whether that assumption is ultimately borne out is a brute fact about our laws. For that reason, even if the assumption proves true, I imagine that some may dismiss CR on the grounds that its success is wed to a purely contingent fact about the world. In the interest of transparency then, allow me to briefly lay out what is riding on the assumption and suggest a way to preserve CR without it.

Consider a world where the assumption is false. Suppose that we discover a new type of mineral that, when heated to 120F, emits two unique forms of radiation: r-rays and c-rays. r-rays and c-rays are so unique in fact that they are only produced by this mineral and only then when the mineral is heated to 120F. Imagine then that a group of scientists heat a sample of the mineral to 120F (h). The sample emits r-rays (r) and c-rays (c). A moment later the scientist's c-ray detector registers the presence of c-rays in the room and sounds a ding (d).

Cases like this present a fundamental problem for any account of causation that, like CR, is built around the notion of contribution. We appear to have two events, r and c , that have different effects but make the same contribution. As such, these cases trivially generate instances of the piggybacking problems that CR (as it's presently framed) can't hope to address. Suppose for convenience then that the sample emitting c-rays (c) is itself a constraint-minimal sufficient condition for the ding (d). It follows that $r \& c$, r , and (disjunctions notwithstanding) $r \& (\neg r \vee c)$ are also constraint-minimal sufficient conditions for d . CR entails that r causes d . But presumably r does no such thing.

4.8 Contributions reconsidered

We might conclude from this that no account of causation grounded in contribution can properly address problem cases like the above, at least not without appealing to factors other than contribution itself. But there is another possibility worth considering. The problem may lie in our account of contribution, not our account of causation.

The present account of contribution effectively defines the contribution of an event on the class of nomically possible worlds: the contribution of an event a is the set of a -worlds in that class. The contribution of a is the set of worlds in which a occurs in as wide a range of environments as the laws allow. One natural lesson to draw from the above problem case is that this range is not wide enough. To tease apart the influence of an event like the sample's emitting r-rays from the influence its emitting c-rays, we can extend that range to include both nomically possible and impossible worlds. Doing so, we can define the contribution of an event on a sufficiently diverse class of worlds such that the event and its background conditions may vary independently of one another, even if that variability runs counter to the actual laws.

Nomically impossible worlds have not, for obvious reasons, played a role in traditional regularity theories of causation. However, they play a critical role in counterfactual theories of causation. Consider the counterfactual conditional $\neg r > \neg e$: if the crowd had not cheered, the ball would not have landed in the stands. Counterfactual theories of causation typically hold that the truth of that counterfactual would entail that r causes e . Now there are perfectly reasonable ways to read $\neg r > \neg e$ such that it comes out true. For instance, we might adopt what Lewis (1979) calls a backtracking reading of the counterfactual. Roughly, a backtracking reading of a counterfactual evaluates that counterfactual on the

assumption that the past counterfactually depends on the present. Read this way, we might conclude that if the crowd did not cheer, it must have been because Casey did not hit the ball. And if Casey did not hit the ball, then the ball would not have landed in the stands. So, the thought goes, the ball would not have landed in the stands if the crowd had not cheered. As a result, backtracking readings of counterfactuals like $\neg r > \neg e$ threaten to saddle counterfactual theories of causation with the problem of epiphenomena. Evaluating those counterfactuals on a model that includes only nomically possible worlds would force just such a backtracking reading on us. Holding fixed the laws, if the present had been different, the past must have been different as well.

Lewis's (1973; 1979) counterfactual theory of causation escapes the problem of epiphenomena because it prohibits backtracking readings of counterfactuals. Since r and e both occur, $\neg r > \neg e$ is true just in case the nearest possible $\neg r$ -worlds are $\neg e$ -worlds. Inevitably for Lewis, such worlds are nomically impossible. Per Lewis, the nearest possible worlds are those whose laws differ from our own in the smallest possible way while still being exactly like the actual world in particular matters of fact over the largest spatio-temporal region. To evaluate $\neg r > \neg e$, we need to imagine r gone, as though it was plucked from the world, and then consider the smallest, most localized violation of the laws that would allow for that. Read this way, $\neg r > \neg e$ is false. The nearest $\neg r$ -worlds are exactly like the actual world up to a time just prior to the occurrence of r at which point their laws diverge from ours in a small, local way such that r does not occur. Since these are all worlds in which Casey hits the ball precisely as he actually did, under essentially the same conditions that he actually hit it, governed by laws that — save for a localized difference around r — are the same as our own, they are worlds where the ball lands in the stands (e).

Lewis's metric for the proximity of worlds is both ingenious and altogether separable from the counterfactual theory of causation that appeals to it. We can use it to extend the scope of the domain on which we define an event's contribution. Take the contribution of an event to be defined on a class of worlds in which the event and its background conditions vary independently of one another to the full extent possible, unconstrained by the laws. Some of these variations occur in nomically possible worlds. Others do not. Where a variation occurs in a nomically possible world, the class includes that world. Where a variation is impossible under the laws, the class includes the worlds in which we find that variation that are closest to a nomically possible world under Lewis' metric (though not, necessarily, closest to the actual world). These additional worlds are not nomically possible, but most are very nearly so. They break with the actual laws in as small and localized a way as possible to accommodate the necessary variation.

In this vein, we can understand contribution as follows:

For any time t , there is a class of worlds W such that for every metaphysically possible state of the world s at t and metaphysically possible world w , $w \in W$ iff (a) s obtains in w at t and (b) no world in which s obtains at t is closer to a nomically possible world than w is to its nearest nomically possible world. For any event e that occurs at t , the contribution of e is the set of e -worlds in W .

The contribution of an event on this account is an extension of its contribution on the account I proposed in 3. Since every nomically possible world is, trivially, the world in which its state at a given time obtains that is closest to a nomically possible world, the contribution of an event e includes every nomically possible world in which e occurs. Additionally, however, the contribution of e also includes

worlds whose state at the time that e occurs is nomically impossible. For example, even if the laws are such that, necessarily, an event f occurs when e occurs, e 's contribution will include $\neg f$ -worlds. So long as $e \& \neg f$ -worlds are metaphysically possible, at least some such worlds will belong to e 's contribution, regardless of the laws.

This account of contribution casts CR in a new light. Because the contribution of an event is a larger, more diverse set of worlds on this new account, conditions that were constraint-minimally sufficient for certain effects on the prior account may not be constraint-minimally sufficient for those effects on the current account. In ordinary cases of causation this difference will be negligible. Not so in stubborn problem cases like the radiation example and in instances of the problem of epiphenomena generally. Consider an example of each.

Take Casey's hit. We imagined earlier that, together, Casey's swing (c), the ball's momentum toward the plate (b), and the absence of any events that would interfere with Casey's swing or the ball's flight ($\neg i$) were a constraint-minimal sufficient condition at the time they occur for the ball's landing in the stands (e). They plausibly still are. If $c \& b \& \neg i$ is not constraint-minimally sufficient for e on the new account of contribution, it is because it is no longer sufficient for e ; though every nomically possible world in the contribution of $c \& b \& \neg i$ is an e -world, perhaps some nomically impossible world in the contribution is not. That is, perhaps c and b occur in the absence of interference and yet, somehow, e fails to occur. There are not obviously any such nearby worlds.¹³ But if there are, the

¹³I suggested in 1 that we think of the absence of interfering events ($\neg i$) as a catch-all for the many events that would have interfered with Casey's swing or the ball's flight. The absence of those events assures that there is no nomically possible interference with b and c causing e . Whether they also assure that there is no nomically impossible interference depends on what absences are included in $\neg i$ in the first place.

lesson is only that we need to add to $c \& b \& \neg i$ the absence of any such nomically impossible threat to e . Casey's swing (c) and the ball's momentum towards the plate (b) — the events that intuitively cause e — would still belong to a constraint-minimal sufficient condition for e at the time they occur.

Finally, return to the radiation example. Because the emission of r-rays (r) is nomically necessary and sufficient for the emission of c-rays (c), the original account of contribution held that r and c make the same contribution and opened the door for piggybacking problems. The new account does not. Since r and c are metaphysically independent, the contribution of r includes $\neg c$ -worlds and the contribution of c includes $\neg r$ -worlds. Paired with this new account of contribution then, CR avoids the piggybacking problems that arose under the previous account. Suppose once more for convenience that c is a constraint-minimal sufficient condition for the ding of the c-ray detector (d). The problem of piling on and the problem of gerrymandering now fail here under CR in the same manner that they originally did in more ordinary cases. Both the contribution of $r \& c$ and $r \& (\neg r \vee c)$ are proper subsets of the contribution of c itself. So while $r \& c$ and $r \& (\neg r \vee c)$ are each sufficient for d , neither is a constraint-minimal sufficient condition for d at the relevant time. In contrast, the new account of contribution cuts off the problem of epiphenomena before it can even begin. Though it runs counter to the laws, it is possible that r occurs in the absence of c . At least one such world belongs to the contribution of r . And that is not a world where d occurs. The sample may emit r-rays, conditions may be such that the c-ray detector would have made a ding *if* c-rays had been emitted, but unless the sample emits c-rays — unless c occurs — there is no guarantee that the c-ray detector will ding (d). As such, r it is not a sufficient condition for d under CR, much less a constraint-minimal sufficient condition.

REFERENCES

- Baumgartner, M. (2008). Regularity theories reassessed. *Philosophia*, **36**(3), pp. 327–354.
- Bennett, K. (2003). Why the exclusion problem seems intractable, and how, just maybe, to tract it. *Nôûs*, **37**(3), pp. 471–497.
- Bennett, K. (2008). Exclusion Again. In Hohwy, J. and J. Kallestrup (eds.) *Being Reduced*. Oxford University Press, Oxford.
- Block, N. (2003). Do Causal Powers Drain Away? *Philosophy and Phenomenological Research*, **67**(1), pp. 133–150.
- Dowe, P. (2000). *Physical Causation*. Cambridge University Press.
- Elga, A. (2001). Statistical mechanics and the asymmetry of counterfactual dependence. *Philosophy of Science*, pp. S313–S324.
- Fair, D. (1979). Causation and the Flow of Energy. *Erkenntnis*, **14**(3), pp. 219–250.
- Horgan, T. (2001). Multiple reference, multiple realization, and the reduction of mind. *Reality and Humean supervenience: Essays on the philosophy of David Lewis*, pp. 205–221.
- Kim, J. (1993a). Concepts of supervenience. In Kim, J. (ed.) *Supervenience and Mind*, pp. 53–87. Cambridge University Press.
- Kim, J. (1993b). The Nonreductivist's Troubles with Mental Causation. pp. 336–357. Cambridge University Press, Cambridge.
- Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. Cambridge University Press.
- Kim, J. (2005). *Physicalism, or something near enough*. Princeton University Press, Princeton, N.J. ISBN 0691113750 (hardcover : alk. paper).
- Lewis, D. (1973). Causation. *Journal of Philosophy*, **70**.

- Lewis, D. (1979). Counterfactual Dependence and Time's Arrow. *Nous*, **13**, pp. 455–476.
- Lewis, D. (1986). Postscripts to "Causation". In *Philosophical Papers*, volume 2, pp. 172–213. Oxford University Press.
- Lewis, D. (2004). Void and object. In Hall, N. and L. A. Paul (eds.) *Causation and Counterfactuals*, pp. 277–290. MIT Press.
- Loewer, B. (2001). Review of *Mind in a Physical World*. *Journal of Philosophy*, **98**(6), pp. 315–324.
- Mackie, J. L. (1965). Causes and conditions. *American philosophical quarterly*, **2**(4), pp. 245–264.
- Mackie, J. L. (1974). *The cement of the universe*. Oxford: Clarendon Press.
- Maslen, C. (2012). Regularity Accounts of Causation and the Problem of Pre-emption: Dark Prospects Indeed. *Erkenntnis*, **77**(3), pp. 419–434.
- Ney, A. (2007). Can an appeal to constitution solve the exclusion problem? *Pacific Philosophical Quarterly*, **88**, pp. 486 – 506.
- Papineau, D. (2001). The rise of physicalism. In Gillett, C. and B. Loewer (eds.) *Physicalism and Its Discontents*, pp. 3–36. Cambridge University Press.
- Papineau, D. (2002). *Thinking about consciousness*. Oxford University Press.
- Paul, L. A. (2007). Constitutive overdetermination. In Campbell, J. K., M. O'Rourke, and H. S. Silverstein (eds.) *Topics in Contemporary Philosophy: Causation and Explanation*, volume 4. MIT Press.
- Paul, L. A. and N. Hall (2013). *Causation: A User's Guide*. OUP Oxford.
- Pereboom, D. (2002). Robust Nonreductive Materialism. *Journal of Philosophy*, **99**, pp. 499–531.
- Pereboom, D. (2011). *Consciousness and the Prospects of Physicalism*. Oxford University Press.
- Salmon, W. C. (1998). *Causality and explanation*. Oxford University Press Oxford.
- Sartorio, C. (2010). The Prince of Wales problem for counterfactual theories of causation. In Hazlett, A. (ed.) *New Waves in Metaphysics*, pp. 259–76. Palgrave Macmillan New York.

- Shoemaker, S. (2001). Realization and Mental Causation. In Gillett, C. and B. Loewer (eds.) *Physicalism and Its Discontents*, chapter 4, pp. 74–98. Cambridge University Press.
- Sider, T. (2003). What’s so bad about overdetermination? *Philosophy and Phenomenological Research*, **67**(3), pp. 719–726.
- Strevens, M. (2007). Mackie Remixed. In Campbell, J. K., M. O’Rourke, and H. Silverstein (eds.) *Causation and explanation*, p. 93. The MIT Press.
- Weslake, B. (2013). Proportionality, Contrast and Explanation. *Australasian Journal of Philosophy*, (ahead-of-print), pp. 1–13.
- Yablo, S. (1992). Mental Causation. *Philosophical Review*, **101**, pp. 245–280.