

Mackie's Theory of Causation Revivified: New Oil in Old Jars

Michael Strevens

Draft of April 2003

Abstract

Cases of overdetermination or preemption continue to play an important role in the debate about the proper interpretation of causal claims of the form *C is a cause of E*. I argue that the best treatment of preemption cases is given by Mackie's venerable INUS account of causal claims. The Mackie account suffers, however, from problems of its own. Inspired by its ability to handle preemption, I propose a dramatic revision to the Mackie account—one that Mackie himself would certainly have rejected—to solve its problems. The result is, I contend, the best available account of causal claims.

To INPC Participants: I will focus on the following sections in my presentation: 2, 4.1, 4.2, 5.2, 5.3.

1. Introduction

A *causal claim* is a claim of the form *C is a cause of E*. Much of the literature on the interpretation of causal claims is motivated by the intuition that a causal claim of the form *C is a cause of E* asserts that *C made a difference to E*, in the sense that taking *C* "out of the picture" would result in a situation

where E no longer occurred. The most natural way to interpret this intuition is in terms of natural language counterfactuals: C made a difference to E just in case, had C not occurred, E would not have occurred. This leads immediately to the view that C is a cause of E just in case, had C not occurred, E would not have occurred. I call this the *simple counterfactual account of causal claims*; it is most often associated with Lewis (1973a), although Lewis himself went far beyond the simple account in the course of his career.

The simple counterfactual account's most notorious difficulty is its handling of cases of preemption, that is, cases in which, had the actual cause C not caused E , some other "backup" cause would have done so. A well-known example is the case of the backup assassin; in the interest of reducing the level of violence in the causation literature, I present in its place the case of the backup Circassian:

The grand vizier, seeking to please the sultan, introduces a beautiful Circassian maiden into the harem. She does her duty, and the sultan is well pleased. Her ministrations, then, are the cause of the sultan's good mood. But the grand vizier, fearful for his upper extremities, has the sultan's old favorite, also, coincidentally, from Circassia, in reserve. The favorite is absolutely reliable: had the new Circassian failed to please the sultan, the backup Circassian would have been dispatched immediately, and would certainly have elevated the sultan's mood.¹

1. I learned of this case from an appendix to the augmented edition of Montesquieu's *Persian Letters*.

Had the new Circassian's endeavors failed to please the sultan, the sultan would still have been pleased. On a simple counterfactual understanding of causal claims, then, the new Circassian did not cause the sultan's good mood. But this conclusion is, of course, incorrect. Lewis calls this a case of *preemption*.

On the face of it, preemption seems not only to be a problem for a counterfactual analysis of causal claims, but for any analysis that takes the difference-making intuition seriously, since whenever there is a backup cause, the action of the actual cause in a sense genuinely does make no difference to the occurrence of the effect.

There is, nevertheless, another sense in which it does make a difference, and this sense is captured by the corresponding causal claim, or so defenders of the difference-making intuition maintain. Attempts to make good on the difference-making intuition have tended to use natural language counterfactuals in more subtle and sophisticated ways so as to manufacture an account of causal claims that delivers the same judgment as the simple counterfactual account in the straightforward cases but that gives the right answer in cases of preemption. Examples include Lewis's later accounts of causation (Lewis 1986b, 2000).

An alternative to this strategy is to abandon natural language counterfactuals as the proper technical tool for assessing difference-making, and to look for some other way to remove *C* from the picture and to assess whether *E* still occurs. The most sophisticated alternative to the counterfactual approach is, in my view, John Mackie's INUS account of causal claims (Mackie 1974). The purpose of this paper is to argue that Mackie's account supplies, without

any of the complex amendments now standard in counterfactual theories, a completely satisfactory treatment of cases of preemption.

This is not, I think, a well-publicized fact. Certainly, Mackie himself seems not to have noticed all the virtues of his theory; his own account of cases of preemption is, as I will explain below, far less satisfactory than the account I offer here.²

The success of the INUS approach in vindicating the difference-making intuition while taking care of preemption shows, I think, that it is well worth revamping for the new century. To this end, I examine two serious problems with Mackie's account, unrelated to preemption, and propose a radical reinterpretation of the INUS machinery which solves both problems and transforms Mackie's account into something rather new—too new, I would guess, for Mackie. The result, then, is a novel theory of causal claims.

2. Mackie's Theory of Causation

According to the theory presented in Mackie (1974), the causal claim *C is a cause of E* is true just in case *C* is an insufficient but non-redundant part of an unnecessary but sufficient condition for the occurrence of *E*. When this requirement is satisfied, *C* is said to be an INUS condition for *E*. The truly important aspects of the analysis are the part's non-redundancy and the

2. McDermott (1995) offers a treatment of preemption building on the Mackie account that is different from both Mackie's and my own. The main difference between McDermott's and my treatments is that McDermott makes no use of the violation of negative conditions. I believe that McDermott's account is unable to handle the problem of the colliding cannonballs presented at the end of section 4.1, but my reasons for thinking so will have to wait until another time.

whole's sufficiency; it is adequate to say that in order for C to count as a cause of E , C must be a non-redundant part of a sufficient condition for E . In addition, obviously, the set of sufficient conditions must be *veridical*, that is, each of the conditions must be true.

Consider an example. The mischievous imperial prince's throwing a cannonball at a fine Iznik jar is a cause, let us say, of the jar's breaking. On Mackie's analysis, the reason for this is as follows. There is a set of conditions that were jointly sufficient for the jar's breaking. These include the prince's throwing the cannonball, but also various other elements of the situation: the fact that the prince was close enough to the jar for his throw to connect, the fact that the grand vizier, protecting the jar with a steel-plated fez, failed to parry the cannonball, and the fact that the gravity on Earth is just the right value that the prince's shot was neither too high nor too low. These conditions, together with the relevant laws of nature, are what is sufficient for the jar's breaking. For Mackie, this means that they entail the breaking. Or at least, sufficiency means entailment in those cases where the putative effect is the result of a deterministic process. If the process is probabilistic, the story is more complex. This paper will focus on the deterministic case.

A non-redundant part of a sufficient condition for an event E is a part that cannot be removed from the sufficient condition without invalidating the entailment of E . Removal here is not negation: if I remove the gravity from the sufficient condition I do not leave behind a condition that says there is no gravity; rather, what is left is a condition that says nothing about gravity at all, leaving open the possibility that the gravitational acceleration acting on the cannonball has any value that you like. Removing the gravity in this

sense invalidates the entailment of the jar's breaking: the condition no longer entails breaking, because it is consistent with the possibility that the gravity is so great that the cannonball crashes to the ground long before it reaches the jar. The gravity, then, is a non-redundant part of the sufficient condition; it is therefore a cause of the breaking. The same is true for each of the conditions listed above, so that each of them counts, on Mackie's view, as a cause of the breaking.

It is important for the Mackie approach, as it is for the Lewis approach and indeed for any difference-making approach, that the relata of causal claims are not what Hempel (1965, 421-423) called *concrete events* but are rather what are often called *states of affairs*. A concrete event is individuated by every detail of its happening; the concrete event of the breaking of a jar, for example, depends on the precise trajectory of every shard of ceramic, so that if one such trajectory had been slightly different, a different concrete event would have occurred. A state of affairs has coarser individuation conditions. The state of affairs of the jar's breaking obtains no matter how, exactly, the shards fly. When Mackie talks of a condition sufficient for the jar's breaking, he means a condition sufficient for the state of affairs to obtain, not a condition sufficient for the underlying concrete event that actually realized the state of affairs to obtain. On the latter interpretation absolutely any physical influence on the breaking would, most likely, count as a non-redundant part of a sufficient condition for the breaking (see section 3). In what follows, then, by an *event* I mean a high level event or state of affairs; when I need to

talk about concrete events, I will always refer to them as such.³

Mackie's account is easy to understand as a difference-making account, using the following notion of difference-making: C makes a difference to E if it plays an essential role in entailing E . Mackie himself saw the INUS machinery in this light. The machinery gives us a way of "removing C from the picture", and checking whether E still occurs, that is quite different from the technique we use for evaluating natural language counterfactuals (though this was not, I think, generally appreciated until Stalnaker's and Lewis's work on counterfactuals (Stalnaker 1968; Lewis 1973b) had been fully digested).

I will return to this topic later (section 4), but let me point out two salient differences between Mackie's difference-making and difference-making as defined using natural language counterfactuals. On a counterfactual account, to see whether C makes a difference to E , we move to a "nearby" possible world (or set of possible worlds) in which C does not occur, and we see whether E occurs in that world. Observe that, first, the "removal" of C on the natural language account corresponds to a *negation* of C , rather than, as on Mackie's account, a lack of an assertion as to whether C occurs or not, and second, that on the natural language account, we try to remove C only from a single "sufficient condition" for E , namely, the state of the entire world at the appropriate time, whereas on Mackie's account, we may try to remove C from any number of different sufficient conditions (and there will always be many such conditions). The putative cause C need only be essential to one

3. It is a matter of controversy whether or not the primary meaning of our non-technical term *event* is closer to *concrete event* or to *state of affairs*. Davidson (1969) maintains the former position; Kim (1973) the latter.

of these sufficient conditions in order to qualify as a cause. It is this second difference that accounts, as we will see, for the Mackie account's superior handling of cases of preemption.

3. Mackie's Account of Preemption

Although Mackie's theory of causation contains all of the apparatus necessary for a completely successful treatment of preemption cases, Mackie's own comments on preemption are far from satisfying.⁴

Mackie considers a case much like that of the backup Circassian:⁵

Smith and Jones commit a crime, but if they had not done so the head of the criminal organization would have sent other members to perform it in their stead, and so it would have been committed anyway (p. 44).

Let us say that, acting on orders from the grand vizier, Smith and Jones poisoned the sultan's wine, killing the sultan. Mackie's view is that Smith and Jones's act of putting poison in the wine is not an INUS condition for the sultan's death, and so is not a cause of the death. That is, he bites the bullet: backup causes really do render the events that preempt them causally impotent, in the sense that the claim that *Smith and Jones' poisoning of the wine*

4. In Mackie (1974) these comments appear, by the way, before the presentation of the INUS account itself.

5. As a matter of fact, he spends most of his time on Hart and Honoré's famous case of the desert traveller with the leaky canteen filled with poisoned water. But his treatment of this tricky case is supposed to apply equally to other, uncontroversial cases such as that of the backup Circassian and the case I am about to discuss.

caused the sultan's death is false. We are simply wrong to think otherwise (Mackie 1974, 44–47).

As a consolation prize, Mackie allows that the poisoning is what he calls a *producing cause* of the sultan's mood. To explain the notion of producing cause, I use the notion of a *concrete realizer* of a high level event E , which I define to be the concrete event that realizes E (for the definition of a concrete event, see section 2). The concrete realizer of a particular jar's breaking, for example, is the concrete event of the breaking, that is, the low level event that is individuated by every physical detail of the breaking.

An event C is a producing cause for another event E if C 's concrete realizer is an INUS condition for E 's concrete realizer. This condition will normally be satisfied if C 's realizer had any physical influence at all on E 's realizer. Consider, for example, the gravitational influence of the bulky chief white eunuch. If the eunuch's influence had been slightly different, the paths traced by the molecules in the chemical reaction that killed the sultan would have been slightly different. But then the concrete realizer of E , the sultan's death, would have been slightly different. That particular realizer would not have occurred. Some other concrete event—also a realizer of the sultan's dying—would have occurred in its place. It follows that the removal of the chief white eunuch's gravitational influence from the totality of physical influences on the realizer will invalidate the entailment of that precise realizer. Thus the chief eunuch's influence is an INUS condition for the realizer.

To be a producing cause, then, is not very difficult, and to say that something is a producing cause is not very informative. In particular, to be told that the poisoning is, like the gravitational influence of the chief white eunuch, a

producing cause for the sultan's death, does not provide much consolation.

In any case, our practice in evaluating causal claims such as these is to hold that the poisoning of the sultan's wine is just as much a cause of the death as the prince's throwing the cannonball is a cause of the jar's breaking; the fact of the crime boss's backup plan does not diminish the causal status of the poisoning at all. Mackie's position fails to capture this practice.

Yet Mackie could have done much better. The poisoning is clearly an INUS condition for the sultan's death: it belongs to a set of conditions sufficient to entail the sultan's death, and it cannot be removed from that set without invalidating the entailment. The relevant set of conditions does *not* mention the backup plan, but it does not need to: Mackie's sets of sufficient conditions need only be sufficient; unlike Lewis's possible worlds, they need not be maximally detailed. The next section goes on to develop an INUS approach to preemption in greater detail.

4. Preemption with the Mackie Account

4.1 Actual Causes Are Not Discounted

Solving the preemption problem using the Mackie account will involve a careful scrutiny of the form of the conditions sufficient for the occurrence of a given event. I will, therefore, consider a causal process that is simpler than a poisoning or a caress.

The imperial prince heaves a cannonball at an Iznik jar, breaking it. The sultan's mother was standing by in case the prince fumbled his throw; had the prince failed to break the jar, the sultan's mother would have thrown

her cannonball and smashed it for sure. The prince's throw fails the simple counterfactual test for difference-making: had he not thrown, the jar would have been broken anyway. What does Mackie's INUS account say about this case?

The prince's throw is part of a set of conditions sufficient for the jar's breaking, namely, the same set of conditions that would have obtained had the sultan's mother, the backup thrower, not been present. The list perhaps looks something like this:

1. The prince threw his ball at time t from such and such a point with such and such a velocity,
2. Nothing interfered with the trajectory of the ball,
3. The jar was in such and such a position at time $t + 1$, and
4. The laws of physics imply that a ball thrown in this fashion at time t will strike a jar in this position at time $t + 1$ hard enough to break the jar, provided that nothing interferes with its flight.

I have suppressed reference to the necessary assumptions about the structure of the ceramic in the jar.

Clearly, condition (1) is a non-redundant part of the sufficient conditions: if it is removed, the conditions no longer entail the jar's breaking. Thus, the prince's throwing the ball is a non-redundant part of a set of conditions sufficient for the breaking, and so is, on Mackie's account, a cause of the breaking. The fact of the backup at no stage enters into the calculation,

which is, I think, as it should be: the fact of the backup is irrelevant to the causal status of the prince's throw.

Is it really this simple? Suppose that we add to the list of sufficient conditions the following condition:

5. The sultan's mother was standing by ready to throw her cannonball, if the prince failed to break the jar with his ball, from such and such a position with such and such a velocity and so on.

Then, if condition (1) is removed from the list, the breaking of the jar is not invalidated: the new condition (5) contributes just enough to make up for the absence of (1). Thus (1) is not a non-redundant part of this set of sufficient conditions.

It is more or less this fact—(1)'s redundancy in the presence of (5)—that is responsible for the failure of the simple counterfactual account to handle cases of preemption. But it is irrelevant to Mackie's account. For *C* to count as a cause of *E*, Mackie's account requires that there exist a veridical set of conditions sufficient for *E* of which *C* is a non-redundant part. This allows that there are other sets of veridical conditions sufficient for *E* in which *C* is redundant. Provided that there is at least one set that fulfils Mackie's requirements, *C* is a cause of *E*. In normal circumstances, no matter how many backups are in place there will always be one such set, a set that mentions none of the backups. Thus backups will normally make no difference to an event's causal status.

Why the *normally*? Keen-eyed readers will have noted a potential difficulty for the Mackie account in a case where a single state of affairs both acts

as a backup and plays an essential role in the actual causal production of the effect, so that any set of sufficient conditions mentioning the actual cause must also mention the backup cause.

Let me give an example. Suppose that the imperial prince and the sultan's mother both throw cannonballs at a jar. The prince's is off target, but the sultan's mother's is deadly accurate. The balls collide in mid-air, however, and the mother's is directed away from the jar, whereas the prince's is deflected towards the jar, which it was otherwise going to miss. The trajectories of the balls are shown in figure 1.




Figure 1: The sultan's mother's throw is both a backup cause and an INUS condition for the jar's breaking.

Observe that any set of sufficient conditions that contains the prince's throw non-redundantly will also have to contain the sultan's mother's throw, since without her throw, the prince's throw would not have been redirected towards the jar. But, the mother's throw is a backup cause for the jar's breaking: had the prince not thrown, her throw would certainly have broken the jar. Thus, it seems that the backup cause must be mentioned in any set of sufficient conditions containing the prince's throw, making his throw a redundant part of all such sets of conditions, and so not a cause of the breaking.

Almost, but not quite. Let us look more closely at the claim that the prince's throw is redundant. Consider conditions (1)–(5) above, which are, I will suppose, sufficient for the breaking. The prince's throw is redundant if conditions (2)–(5) are sufficient in themselves to entail the breaking. But despite appearances, conditions (2)–(5) do not entail breaking, and there is no way to tweak the conditions or the example to make do so. The reason is this: in order for the conditions to entail the breaking, they must entail that the sultan's mother's throw hits the jar. In effect, they are describing what would have happened if the prince had not thrown. But for this entailment to go through, the conditions would have to include a condition of the form:

6. Nothing interfered with the trajectory of the sultan's mother's ball.

They do not include such a condition. More importantly, they could *never* contain this condition while satisfying the requirements of Mackie's account, for the account requires that the sufficient conditions be veridical—that is, true of the actual world—but (6) is false.

It follows that (1) is not, after all redundant: it cannot be removed from the conditions without invalidating the entailment of the breaking. Thus it is a cause of the breaking, as desired.

This is not due to any peculiarity of the scenario. A backup cause is *merely* a backup cause because the conditions necessary for it to exert its characteristic effect (I mean the effect it would have if it were called on as backup) are not all present. The backup Circassian never cavorts with the sultan, Smith and Jones' boss, the grand vizier, never dispatches the B team, and the sultan's mother's cannonball never hits the jar. A set of conditions

sufficient for the backup cause to have its characteristic effect would have to assert otherwise, but then it would assert falsely. The requirement of veridicality would not be met.

4.2 Backup and Other Ineffective Causes Are Discounted

The success of the Mackie account in coping with preemption where the simple counterfactual account fails might be explained as follows: the Mackie account is far more liberal than the counterfactual account. The Mackie account declares an event *C* a cause of *E* as long as *C* is a nonredundant part of just one out of the great number of sets of conditions sufficient for *E*, whereas the counterfactual account demands that *C*'s role be essential in a very particular set of sufficient conditions for *E*, namely, the set describing every aspect of the actual world. (The counterfactual test means something different by *essential* than the Mackie account means by *non-redundant*, of course—I am speaking rather loosely here.)

If the Mackie account works its magic in virtue of its liberality, it makes sense to worry that it is perhaps too liberal. Does it classify as causes events that are not causes at all? I will consider two cases that present *prima facie* problems for the Mackie account.

Consider jar-breaking again. This time, the sultan's mother throws her cannonball immediately after the prince throws his, with her usual deadly accuracy. The prince is on target this time, and his ball breaks the jar, but the prince's grandmother's ball was close on its tail. Had he missed, her throw would have connected.

The question I want to ask is, not whether the prince's throw counts as

a cause, which it does for the reasons given above, but whether the Mackie account erroneously counts the sultan's mother's throw, too, as a cause. Is there a set of conditions, sufficient for the jar's breaking, of which her throw is an essential part? At first, it may seem so. Her ball is thrown in the right direction with the right amount of heft, and—in contrast to the case of the colliding balls considered above—nothing interferes with its flight. Put together these conditions, saying nothing about the prince's throw, and do we not have what is needed to elevate the mother's throw to the status of a cause?

Again, not quite. One of the conditions that must be present for a throw to entail a jar's breaking is, recall from above,

The jar was in such and such a position at time $t + 1$,

where the relevant ball was thrown at time t and, given the relevant laws of nature and other physical parameters, the ball takes one unit of time to reach the jar. It is this condition that does not hold in the present scenario: if time t is the time of the sultan's mother's throw, then at time $t + 1$ the jar is no longer at the required position; indeed, the jar is no longer anywhere at all, since it has been destroyed by the prince's throw. There is no set of veridical conditions, then, that entails that the mother's ball strikes the jar, thus it is not a cause.

If there *were* a set of veridical conditions entailing that the mother's ball hits the jar, it would indeed have hit the jar. But then it would, intuitively, be a cause of the breaking, in which case Mackie's account would be correct in deeming it so.⁶

6. Unless, perhaps, this were a case of true overdetermination (see section 4.3).

That takes care of putative causes of an event E that come too late to be real causes of E . What of putative causes that come early, but that fail to bring about E for some other reason? Suppose, for example, that the grand vizier poisons the sultan, but that he survives. The next day, the sultan slips on a jar shard and breaks his neck. Will the Mackie account correctly discount the grand vizier's poisoning the sultan as a cause of the sultan's death? It will: since the poison did not kill the sultan, and we are assuming that all underlying processes are deterministic, some condition must have failed to hold that was required in the circumstances for the poison to have its effect. Perhaps the sultan had just taken the antidote; perhaps he has an iron constitution; perhaps he drank the poison on a full stomach. Whatever this condition, its negation, or something that entails its negation, would have to be a part of any set of conditions in which the poisoning played an essential part in entailing his death. But this negation would be non-veridical: it would assert that some condition—the actual condition responsible for the sultan's survival—did not obtain, when in fact it did. The poisoning, then, cannot be a cause of death.

4.3 True Overdetermination

In order to bolster the INUS account's claim to handle difficult cases better than any other account of causal claims, let me consider another kind of scenario that has interested philosophers of causation: true overdetermination.

The chief white eunuch and the chief black eunuch simultaneously hurl their scimitars at an intruder in the palace. One scimitar impales the left lobe of the intruder's heart, the other the right lobe. Are both throws, or

either, or neither, causes of the intruder's death? This is the question raised by overdetermination

On the Mackie approach, it appears that both throws are causes. Each belongs to a set of conditions sufficient for the sultan's death, and plays an essential part in the entailment. The relevant set of conditions is, of course, the usual set of conditions that we would use to establish the causal status of such a throw, with the fact of the other throw excluded.

Is this the right answer to the question about overdetermining causes? Some writers believe so. Yet there is something a little odd about these cases: we feel, for each throw, that the existence of the other throw somehow does detract a little, though far from completely and in a very hard to define way, from its causal status, in a way that a mere backup cause never would.

The Mackie account can explain the oddness, if it is extended in a certain way. I will briefly sketch the extension I have in mind; it will not be developed or defended, however, nor will the extension play any further role in this paper.⁷

The Mackie account uncovers causes of an event E by removing from a set of sufficient conditions for E all those conditions not essential to the entailment of E . I propose the following extension: we can remove not only conditions, but *parts* of conditions. More exactly, we can remove particular details from a condition, leaving behind something more vague, provided that the removal does not invalidate the entailment of E .

For example, suppose that a set of conditions sufficient for the breaking

7. For the necessary development and defense, see Strevens (2003).

of an Iznik jar specifies that the weight of the cannonball hefted at the jar by the sultan's mother was exactly 2 kg. On the original Mackie account, we would have to leave this condition in the description; taking it out would leave the weight of the cannonball unspecified, creating the possibility that the ball might weigh only 2 g, in which case the jar would not have been broken. Thus, on the Mackie account, the exact weight of the ball is a cause of the breaking.

If the account is extended, however, we are allowed to make the description of the weight more vague. We can replace the condition stating an exact weight of 2 kg with a vaguer specification, say that the weight was between 1 kg and 10 kg. (Why an upper limit? Even the sultan's mother's strength has its limits.) This replacement can be carried out without invalidating the entailment of the breaking. Result: we can claim, on the extended account, that though the exact weight of the cannonball did not make a difference to the jar's breaking, the approximate weight did; the ball's being quite heavy, but not too heavy to lift, was the difference-maker.

Now we can understand the peculiarity of true overdetermination by considering the way that the extended Mackie account deals with two sets of conditions sufficient for the intruder's death.

First, consider a set of sufficient conditions for the death that includes the chief black eunuch's scimitar throw non-redundantly, but that does not mention the chief white eunuch's throw. The conditions entail the scimitar's flying towards the intruder's chest, entering the heart, causing traumatic damage to the heart tissue, and so on. Now remove all the redundant detail, as envisaged by the extended Mackie account. The pared-down sufficient

conditions will not specify the precise trajectory of the scimitar; rather, they will say just enough about the trajectory to entail that the scimitar strikes the heart. Consequently, they will not entail the precise details of the damage caused by the scimitar; rather, they will entail only that massive damage is done to the heart.

But observe: everything that the pared-down sufficient conditions say of the chief black eunuch's throw, is also true of the chief white eunuch's throw, because of the vagueness in the description of the trajectory, the damage, and so on. Looking at the conditions alone, you cannot tell which scimitar throw they refer to. The conditions say that a scimitar throw caused the death, and they give some details about that throw, but the details are not sufficient to specify one of the throws rather than the other.

Second, consider a set of conditions sufficient for the death that describes both scimitar throws. It is possible to remove entirely the conditions describing one of the throws provided that the conditions describing the other are left behind. The pared-down conditions will mention only one throw, then.⁸ But, for the reasons given in the last paragraph, there will be nothing left in the conditions to determine which of the two throws is being mentioned.

So what is the cause of the intruder's death? The two throws are not a joint cause of the death, because a set of conditions that mentions both throws can always be pared-down so that it mentions only one. But the throws are not separate causes of the death, either. For each to be a cause, there would

8. Though they will leave open the possibility that there was more than one throw, on pain of non-veridicality.

have to be two sets of pared-down sufficient conditions, one describing how the chief white eunuch's throw lead to death and one describing how the chief black eunuch's throw lead to death. Instead we have a single set of pared-down conditions mentioning a single throw—but which one, is undetermined. This result, I propose, precisely captures the ambivalence we feel in cases of true overdetermination.

5. The Mackie Account Reconstituted

For all the advantages that it enjoys in handling cases of preemption, the Mackie account faces deep problems of its own. If we are to take advantage of the Mackie approach to preemption, these problems must be addressed; the result is an outline of a theory of causal claims that has a quite different flavor from Mackie's.

I will focus on two shortcomings of Mackie's account. The first is independent of, whereas the second exposes serious difficulties with, the handling of preemption problems discussed above. The solution to the first will, however, point the way to the solution of the second.

5.1 Distinguishing Correlation and Causation

The first problem is that of the notorious Manchester hooters. When the hooters go off in Manchester, the factory workers in Manchester go home. Similarly, when the hooters go off in London, the workers in London go home. It seems that the Mackie account classes the sounding of the Manchester hooters as a cause of the London workers' going home, for the following reason.

The relevant generalizations about the workings of the hooter system in industrial society, and some facts about London and Manchester, entail the truth of the following generalization: when the hooters go off in Manchester, the factory workers in London go home. The generalization, together with the fact that the Manchester hooters sound at a particular time t , entails that the London workers go home at time $t + 1$. Removing the fact of the Manchester hooters' sounding from this set of sufficient conditions invalidates the entailment. Thus the sounding of the Manchester hooters is an INUS condition for the London workers' return home. But we do not want to say that the Manchester hooters cause the London workers to go home.

Mackie's solution to this problem emerges from a rather informal discussion in which he suggests that the Manchester hooters do not qualify as a cause of the London workers' return because of some combination of (a) considerations concerning the times at which events become "fixed", which are redolent of the screening off criterion proposed by Reichenbach (1956), and (b) a negative answer to the question whether the sounding of the Manchester hooters is linked to the London workers' return by a "continuous causal process" (Mackie 1974, 190–192).

It is difficult to extract from this discussion a canonical solution to the hooters problem, and I will not try to do so here. (If anything, one gets the impression that for Mackie, empiricist that he is, our distinction between the effect of the Manchester and the London hooters is more a human foible than a desirable feature of an ideal science.) Let me point instead to Mackie's own admission (p. 191) that his solution to the hooters problem will not work in a completely deterministic world. This represents, I think, a fundamental weak-

ness in Mackie's account, if it is construed as an account of the truth conditions for causal claims, rather than as a piece of revisionary metaphysics. Clearly, we do distinguish the effects of the London and Manchester hooters, and clearly, the question of determinism has no bearing on the distinctions we make.

Why do we deny that the Manchester hooters cause the London workers' return? Mackie is correct, I think, when he points to our beliefs about continuous causal processes as the source of our denial. There is no causal process linking the sounding of the Manchester hooters to the London workers' actions, whereas there is such a process linking the sounding of the Manchester hooters to the Manchester workers' actions. It is for this reason, I suggest, that although the sounding of the Manchester hooters is an INUS condition for both the London and the Manchester workers' return, it is right to say that the Manchester hooters caused the Manchester workers' return, but wrong to say that the Manchester hooters caused the London workers' return.

If this approach is correct, then there must be facts about continuous causal connections that are prior to, therefore independent of, the facts asserted by causal claims. Mackie, the reader will recall, has an appropriate definition of a continuous causal connection, namely, his relation of causal production described in section 3, characterized in terms of INUS conditions for concrete realizers. It seems that Mackie builds causal connection into his account of causal claims simply by adding to the INUS account a requirement that cause and effect be casually connected in his proprietary sense.

My proposed revision to Mackie's account differs from Mackie's own suggestion in two ways. First, the facts about causal connection are not defined

by INUS conditions for concrete realizers, but are rather read off directly from the relevant causal laws. Second, the requirement of causal connection is not added to the INUS account; rather, the definition of an INUS condition is itself modified to reflect facts about causal connection. I discuss each requirement in turn.

First, the source of the facts about causal connection. We hold that the sound of the Manchester hooters is causally connected to the Manchester workers' returning home, but not to the London workers' return. Why? Because physics tells us that the sound of the Manchester hooters reaches the ears of the Manchester workers, but not the ears of the London workers. Thus the Manchester hooters are causally connected to the Manchester workers' actions, but not to the London workers' actions. In general, I claim, we can read the facts about causal connection off from the nomological dependencies spelled out in the laws of physics. There is more to say, of course, about how it is done, but as it is not directly relevant to my defense of the Mackie approach, I will not say it here. Some possible approaches are found in "process" accounts of causation, for example, Dowe (1992).

Second, the question of how to build a requirement of causal connect-edness into the Mackie account. A part of the answer, sufficient to solve the hooters problem, is given here; the full answer will be given in section 5.3.

In Mackie's original account, we begin with a set of conditions sufficient to entail that the putative effect E occurred. By contrast, I propose that we begin with a set of conditions *causally sufficient* for E . The full definition of causal sufficiency will be stated in section 5.3; for now I give just a necessary condition for causal sufficiency: a set of conditions sufficient for E is causally

sufficient only if each condition describes a causal influence on E , by which I mean that each condition describes either (a) an event that had some causal influence on E , (b) a causal law (or set of laws) in virtue of which an event had such an influence, or (c) a background condition necessary for the operation of such a causal law.

Note that, because the background conditions required for the operation of a law are sometimes negative states of affairs—such as nothing’s having interfered with the flight of the prince’s cannonball—an absence can count as a “causal influence” in my technical sense. This opens the door to causation by omission; however, the treatment of omissions requires some adjustments to my account that I will not pursue here, and so I will have to leave that very interesting topic to another time.

The hooters problem is to be solved as follows. The problem arises because it is possible to assemble a set of true facts and generalizations in virtue of which the sounding of the Manchester hooters entails the London workers’ return. Among these generalizations is something more or less equivalent to the following generalization: *Whenever the Manchester hooters sound, the London hooters sound*. This generalization, though true, does not state a fact about causal influence. It cannot, because it is not a statement of a causal law, since there is no underlying physical process by which its antecedent leads to its consequent. Thus, the conditions sufficient for the London workers’ return are not *causally sufficient* for their return; consequently, non-redundant parts of the sufficient conditions, including the sounding of the Manchester hooters, are not thereby causes of the return.

5.2 The Role of Causal Claims in Understanding

Let me pause to sketch the role, as I see it, that causal claims play in our scientific and ordinary discourse. This discussion will provide the basis for solving the second problem with Mackie's account in section 5.3.

The world, according to physics, is vast and complex causal web. For the most part, the elements of the web are completely determined by the initial conditions of the universe and the fundamental laws of nature. Find any property of any particular region of space-time, and there is some combination of physical facts and laws that is causally sufficient for the region's instantiating the property. This much physics tells us, or so I suppose.

If we were fully satisfied with knowledge of the facts about causal sufficiency, we would have no need for causal claims. It would be enough, for any event E , to know, concerning any other event C , that C causally influenced E , meaning that we can trace, by way of a series of instantiations of causal laws, a chain of events causally connecting C to E . We would know, to use Mackie's term, that there was a "continuous causal process" connecting C and E .

Just how meager this knowledge is can be seen from my earlier discussion of Mackie's relation of causal production (not quite the same relation that I have described here). The vast mass of the chief white eunuch is causally connected to the sultan's death by poisoning (Smith and Jones being the perpetrators, you will recall), due to the gravitational influence it exerts on the event, as on every other event in the vicinity. Using the laws of physics, that is, we can trace a line of causal influence from the eunuch to the dying

sultan. But what does this tell us? Almost nothing of interest.

What we really want to know is what, of all the physical influences on the sultan's dying, *made a difference* to the fact that he died. It is here, I claim, that the INUS apparatus comes into its own. What we want to find are the parts of the causal network that play an essential role in the causation of the death. These may be ascertained by looking to a description of the web of causal influences in which the death is embedded and finding the parts that play an essential role in entailing the death's occurrence, that is, roughly, the parts that cannot be removed from the description of the web without invalidating the causal entailment of the death.

What facts about the world, then, are our causal claims supposed to capture? Not the facts about fundamental causal relations, about what causally influenced what, as Mackie, Lewis, and others assume. Rather, causal claims capture higher level facts about which causal influences played a critical role, which were decisive, in bringing about some high level (almost never concrete) event. These are the causal influences that—unlike the gravitational influence of sundry large bodies—made a difference between the event's occurring and its failing to occur.

I concur with Lewis and Mackie, then, that causal claims are claims about difference-making. I disagree that this difference-making is itself the fundamental causal relation. The web of causal influence is presented to us, by fundamental physics, as a given. The role of causal claims is to single out the elements of the web that are relevant to whatever high level events interest us: the breakings, the pleasurings, the dyings, and all the rest.

5.3 Spurious Non-Redundancy

Late twentieth century philosophy has discovered that over-reliance on the entailment relation can endanger one's analyses. For Mackie's account of causal claims, the danger is exemplified by the following recipe.

To show that any event R whatsoever is a cause of a given event E : Take a set of non-redundant conditions jointly sufficient for E . Replace one of these conditions C with the following two conditions: R and $R \supset C$. The new set of conditions is also sufficient for C . The intuitively irrelevant event R cannot be removed from this set without invalidating the entailment of E . Therefore R is an INUS condition for E , and so, according to Mackie's account, R is a cause of E .

Let me give a specific example. The sultan has been murdered. A set of sufficient conditions for his death involved, non-redundantly, the fact of the grand vizier's poisoning his wine. Earlier, the chief black eunuch sneezed. To show that the sneeze is a cause of the death, take a list of sufficient conditions for death non-redundantly including the grand vizier's poisoning the wine, and replace the poisoning with the fact of the eunuch's sneeze and the following disjunction:

Either the chief black eunuch did not sneeze, or the grand vizier poisoned the sultan's wine.

Then the sneeze is, in virtue of the set of sufficient conditions so constructed, an INUS condition for the sultan's death.

The problem exists because it is so easy to play an essential role in an entailment. If the Mackie account is to be saved, some sort of constraint must

be imposed on the kinds of entailments that count for the purpose of determining causes. An entailment involving an irrelevant R and the disjunction $R \supset C$ must be declared, for some reason, illegitimate.

For what reason, then? Let me take as my starting point the amendment made to the Mackie account in response to the hooter problem. In the determination of the causes of an event E , I proposed, it is not enough that a set of conditions be sufficient for the occurrence of E ; it must be *causally sufficient* for E . In section 5.1, I gave a necessary condition for causal sufficiency: all conditions must describe either events that causally influence E or causal laws or background conditions in virtue of which they do so.

This seems not enough in itself to solve the problem of spurious non-redundancy, since the eunuch's sneeze is a causal influence, in the same negligible way as any bystander's gravitational influence, on the sultan's death. In what follows, I complete my account of causal sufficiency so as to rule out the sneeze as a cause of death, taking as inspiration the picture of the role of causal claims sketched in section 5.2.

According to that picture, a causal claim picks out a piece of the causal web essential for the production of some event E . The revised Mackie account promises to determine such causes by, first finding a part of the causal web sufficient for the production of E , described by a sufficient condition for E , and then discarding those elements that are not essential to the production of E , the redundant parts of the sufficient condition.

The sufficient condition for E , then, is supposed to represent a part of the causal process that produced E . Many sets of conditions may entail E but do not represent any part of the process that caused E ; the case of the

Manchester hooters provides a salient example. It is these conditions that I am trying to rule out of contention for the INUS treatment by my requirement that the conditions not only be sufficient for E , but be causally sufficient.

Causal sufficiency ought to be defined, then, so that a set of conditions is causally sufficient for an event E only if the conditions represent a causal process that produces E . A set of conditions entailing E represents a causal process producing E , I propose, just in case each step in the entailment corresponds to a strand in the relevant causal web.

Take, to choose the simplest possible example, an event C and a law *All events of the same type as C cause events of the same type as E* . These two conditions not only entail E ; they correspond to a step in the causal chain that produced E , namely, the step from C to E . Call such an entailment a *causal entailment*.

Now consider by contrast the case of the eunuch's sneeze. The step in the entailment of the sultan's death that involves the sneeze is the step from the sneeze and the disjunction *Either the chief black eunuch did not sneeze, or the grand vizier poisoned the sultan's wine* to the conclusion that the grand vizier poisoned the sultan's wine. This is not a causal entailment, as it does not correspond to a causal process recognized by the laws of physics. Indeed, it is hard to imagine a physics in which something in the world captured by the description $\neg C \vee A$ could be a part of any story about causal influence.

I define causal sufficiency, then, as follows: a set of conditions is causally sufficient for an event E just in case each step in the entailment of E is a causal entailment. It follows that, because the chief black eunuch's sneeze is an INUS condition for the sultan's death only by way of an entailment that is

not causal, the sneeze does not count as a cause of the death. This solution goes to the heart of the problem: when what we regard as an intuitively irrelevant factor R is made essential to the entailment of an event E , it is always by way of a disjunction or other logical construction that links R and E truth-functionally but not causally.

The approach to causal claims taken by my revision of Mackie's account puts a considerable burden on the physical laws: they must determine what primitive causal connections there are in the world. I think that they are quite capable of bearing the load, and that we do indeed look to the laws as the final arbiters on any question of causal connection. But I will not try to make the case here; it is enough for my present purposes to show that an updated INUS account is not defenseless against the old objections to the Mackie account.

In the course of the defense, Mackie's account has been transformed into something that he would likely not endorse. It is no longer explicitly empiricist—though it is compatible with empiricism, since we may give an empiricist account of causal influence. More important, although it makes use of deductive logic, and in particular the entailment relation, logical constructs and relations do not, as they do in the logical empiricist tradition, replace metaphysics. Rather, they are used to represent metaphysics. No longer does logical necessity take over from some forsaken relation of nomic dependence. Its role in the new account is far more humble: it is used to represent the species of nomic dependence that I am calling causal influence. The Mackie account, by picking out certain propositions as essential to the entailment of the proposition that E occurred, also seeks out the real object

of our inquiry, the causal influences essential to the causal production of E itself.

References

- Davidson, D. (1969). The individuation of events. In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, pp. 216–34. D. Reidel, Dordrecht.
- Dowe, P. (1992). Wesley Salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science* 59:195–216.
- Hempel, C. (1965). Aspects of scientific explanation. In *Aspects of Scientific Explanation*, chap. 12. Free Press, New York.
- Kim, J. (1973). Causation, nomic subsumption, and the concept of event. *Journal of Philosophy* 70:217–236.
- Lewis, D. (1973a). Causation. *Journal of Philosophy* 70:556–67. Reprinted in Lewis (1986a).
- . (1973b). *Counterfactuals*. Harvard University Press, Cambridge, MA.
- . (1986a). *Philosophical Papers*, volume 2. Oxford University Press, Oxford.
- . (1986b). Postscript to "Causation". In Lewis (1986a).
- . (2000). Causation as influence. *Journal of Philosophy* 97:182–97.
- Mackie, J. (1974). *The Cement of the Universe*. Oxford University Press, Oxford.
- McDermott, M. (1995). Redundant causation. *British Journal for the Philosophy of Science* 46:523–544.
- Reichenbach, H. (1956). *The Direction of Time*. University of California Press, Berkeley, CA.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (ed.), *Studies in Logical Theory*. Blackwell, Oxford.
- Strevens, M. (2003). The causal and unification accounts of explanation unified—causally. Forthcoming, *Noûs*.