

Mackie Remastered

Michael Strevens

Topics in Contemporary Philosophy, vol. 4, MIT Press

ABSTRACT

Cases of overdetermination or preemption continue to play an important role in the debate about the proper interpretation of singular causal claims of the form *c was 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 these difficulties. The result is, I contend, a very attractive account of singular causal claims.

1. INTRODUCTION

A *singular causal claim* is a claim of the form *c was a cause of e*, where *c* and *e* are token events. Much of the literature on the interpretation of causal claims is motivated by the intuition that *c was a cause of e* asserts that *c made a difference* to whether or not *e* occurred, 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 difference-making 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 a simple counterfactual account of causal claims, according to which *c* was a cause of *e* just in case, had *c* not occurred, *e* would not have occurred. The simple counterfactual account is most often associated with Lewis (1973a), although Lewis himself went far beyond the simple account in the course of his writings on causation.

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 acquits herself superbly, and the sultan is well satisfied. Her ministrations, then, are the cause of the sultan’s good mood. But the grand vizier, ever mindful of the contingency of his position, 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 transported his epicurean majesty to a higher plane of happi-

ness.¹

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—the wrong answer.

It is now standard to follow Lewis in calling such scenarios cases of *preemption*, since what is intuitively the actual cause preempts the backup cause. (For an excellent and wide-ranging discussion of the significance of preemption scenarios for the causal claims literature, see Hall and Paul (2003).) On the face of it, preemption seems not only to be a problem for a simple 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. Among the best known examples are, of course, Lewis's various accounts of causation (Lewis 1973a, 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. One promising removal technique is an approach based on

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

work in computer science and other disciplines on “causal Bayesian networks” Pearl (2000); see Hitchcock (2001) and in a different vein Yablo (2002) for the application to preemption. Another sophisticated alternative to the counterfactual approach is 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 the standard cases of preemption.

This is not, I think, a well-publicized fact. Certainly, Mackie himself seems not to have noticed all of the virtues of his theory; his own account of preempting causes is, as I will explain below, far less satisfactory than the account I offer here on his behalf.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, drawing inspiration from related work on causal explanation (Strevens 2004), I propose a radical reinterpretation of the INUS machinery that 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 the meaning of causal claims.

For the sake of truth in advertising, let me remark on two ways in which this new theory is limited. The first limitation, common to much of the preemption literature, is a restriction to causal claims about processes governed by deterministic laws.

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 occasion.

The second is that the theory of causal claims on offer is not also a reductive metaphysics of causation. It does not translate causal claims into claims about something not intrinsically causal; rather, it takes for granted a world permeated by a certain kind of primitive physical causal connection, and offers an account of what causal claims are saying about such a world. This interpretation of the role of causal claims in our causal understanding of the world is developed in section 5; first, however, my appreciation of the virtues of Mackie.

2. MACKIE'S THEORY OF CAUSATION

According to the theory presented in Mackie (1974), the causal claim *c was 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 critical aspects of the analysis are the part's non-redundancy and the whole's sufficiency; for my present purposes, it is enough to say that *c* is a cause of *e* just in case *c* is a non-redundant part of a sufficient condition for *e*. Though it is not always made explicit, 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 was a cause, 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 exerted just the right pull on the ball 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. As noted above, this paper will assume determinism.

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 way 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 (and Mackie) called *concrete events* but are rather what are often called *states of affairs* (Hempel 1965, 421–423). A concrete event is individuated by every intrinsic detail of its happening; the concrete event of the breaking of a jar, for example, is individuated by 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 con-

dition 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 (as explained in 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 can and should be understood as a difference-making account, using the following notion of difference-making: *c* makes a difference to whether or not *e* occurs just in case it plays an essential role in entailing *e*. Mackie himself saw the INUS machinery in this light, but he thought of it as providing an analysis of natural language counterfactuals rather than as an alternative to the counterfactual characterization of difference-making. In fact, Mackie's machinery constitutes 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. (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* made a difference to *e*, you move to a "nearby" possible world (or set of possible worlds) in which *c* did not occur, and you 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, you try to remove *c* only from

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.

a single “sufficient condition” for e , namely, the state of the entire world at the appropriate time, whereas on Mackie’s account, you 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 of these sufficient conditions in order to qualify as a cause. It is this second difference that accounts, as you 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).

Suppose that Smith and Jones, acting on orders from the grand vizier, 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 before the presentation of the INUS account itself.

5. Most of Mackie’s discussion concerns Hart and Honoré’s famous case of the desert traveler with the leaky canteen filled with poisoned water. But his treatment of this complex 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 consolation, Mackie allows that the poisoning is what he calls a *producing cause* of the sultan's death. To characterize a producing cause, I will use the notion of the *concrete realizer* of a high level event e , which I define to be just what you would think: it is the concrete event that realizes e (concrete events having been defined in 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 actual concrete realizer of e , the sultan's death, 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 mass 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 compensation for the poisoning's being stripped of its causehood.

In any case, our practice in evaluating causal claims such as these is to hold that the poisoning of the sultan's wine was just as much a cause of

the death as the prince's throwing the cannonball was a cause of the jar's breaking; the fact of the grand vizier'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. This point will be explained at greater length in the next section, where I develop an INUS approach to the standard cases of preemption.

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 the backup Circassian's favorite 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. In this standard example of *early preemption*, 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 have to say about the throw?

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 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 presence of the Sultan's mother is irrelevant to the causal status of the prince's throw.

Is handling preemption really this easy? Suppose that you 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 had failed to break the jar with his ball, she would have launched hers 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 fulfills Mackie’s requirements, c was a cause of e . In normal circumstances, no matter how many backups are in place there will always be one set featuring the actual cause non-redundantly, 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 (taken from Strevens 2003). 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.

Observe that any set of sufficient conditions for the jar’s breaking 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




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

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. 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 them 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 was 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.⁶

For this reason, the Mackie account is sure to succeed where the simple counterfactual account does not, counting what are intuitively real causes as real in any case of preemption, including all the variants of early preemption discussed above as well as what is called *late preemption*, an example of which will be given in the next section.

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 non-redundant 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

6. Ramachandran (1997) proposes a variant of the counterfactual account of causation based on a similar observation (p. 273). This account cannot, I think, handle the case of the colliding balls.

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, you might 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 Mackie.

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 his grandmother's ball was close on its tail. Had he missed, her throw would have connected. This is what Lewis calls *late preemption*.

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 you not have what is needed to elevate the mother's throw to the status of a cause?

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. The mother's ball 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.⁷

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 the sultan survives. The next day, he 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 all underlying processes are by working assumption 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 an antidote; perhaps he had 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 the sultan's 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.

One further case of preemption that can cause trouble for the Mackie approach will be discussed in section 5.4.

7. Unless, perhaps, this was a case of symmetric overdetermination (see section 4.3).

4.3 *Symmetric 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: symmetric 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 you 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 (Schaffer 2003). Yet there is something a little odd about these cases: you 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. Compare and contrast with cases of preemption, where a backup cause in no way compromises the status of the actual cause.

The Mackie account can explain what is going on, 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. (For the necessary development and defense, see Strevens (2004).)

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: you can remove not

only conditions, but also *parts* of conditions. More exactly, you can remove particular details from a condition, leaving behind something strictly more abstract, provided that the removal does not invalidate the entailment of *e*.

For example, suppose that a set of conditions sufficient for the breaking 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, you 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 was a cause of the breaking.

If the account is extended, however, you are allowed to make the description of the weight more abstract. You can replace the condition stating an exact weight of 2 kg with a looser specification, say that the weight was between 1 kg and 10 kg. (Why an upper bound? Even the sultan's mother's strength has its limits.) This replacement can be carried out without invalidating the entailment of the breaking. Result: you 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. Mackie himself, it should be noted, comes very close to spelling out the extended account in chapter 10 of *The Cement of the Universe* (see especially pp. 260–265).

The extended Mackie account explains the peculiarity of symmetric overdetermination. Consider a set of sufficient conditions for the intruder's 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. The pared-down conditions say that a scimitar throw caused the death, and they give some details about that throw, yet looking at the conditions in isolation, you cannot tell which throw is being described.

Now 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 have to be two sets of pared-down sufficient conditions, one describing how the chief white eunuch's throw led to death and one describing how the chief black eunuch's throw led to death. Instead you 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 symmetric overdetermination.

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

5. MACKIE REMASTERED

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

According to this reconception of Mackie's theory, the role of causal claims in science and everyday life is not to express basic metaphysical facts about causal connections, but rather to extract from the basic causal facts an understanding of how causal connections work together to bring about certain states of affairs. The question of the meaning of causal claims turns out to be less a question about metaphysics, then, and more a question about understanding or explanation. (Appropriately, the theory of causal claims sketched in what follows was first presented as an account of causal explanation in Strevens (2004).)

If I am right, it will turn out to have been an error—a major and pervasive error—for causal metaphysicians to have focused so great a part of their energies on causal claims. Let me put these strategic remarks to one side, however, to concentrate on the task of providing an account of the truth conditions of causal claims that is adequate to the data, that is, to our intuitive judgments of causal claims' truth and falsehood.

In order to motivate my theory, 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 (Mackie 1974, 81). Let me tell the story in a post-industrial guise. When the *boru*, or horn, sounds on the European side of the Bosphorus, the Rumelian janissaries assemble outside the sultan's New Palace. Similarly, when the *boru* sounds on the Asian side of the Bosphorus, the Anatolian janissaries assemble outside the Old Palace. Seeking the sense of order that evades him in his own harem, the sultan insists that the *boru* sound at exactly the same time on both sides of the Bosphorus, although the sound does not carry over the water. As a result of this very regular timing, it seems that the Mackie account classes the sounding of the Rumelian *boru* as a cause of the Anatolian janissaries' assembly, even though the Anatolian janissaries cannot hear the Rumelian *boru*.

At the core of the problem is the truth, indeed robustness, of the following generalization: when the *boru* is sounded on the Rumelian side of the Bosphorus, the Anatolian janissaries assemble. The generalization, together with the fact that the Rumelian *boru* sounds at a particular time t , entails that the Anatolian janissaries assemble at time $t + 1$. Removing the fact of the Rumelian *boru*'s sounding from this set of sufficient conditions invalidates the entailment. Thus the sounding of the Rumelian *boru* is an INUS condition for the Anatolian janissaries' assembly. But we do not say that the Rumelian *boru* causes the Anatolian janissaries to assemble.

Mackie's solution to this problem emerges from a rather informal discussion in which he suggests that the Rumelian *boru* does not qualify as a cause of the Anatolian janissaries' assembly 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 Rumelian *boru* is linked to the Anatolian janissaries' assembly by a "continuous causal process" (Mackie 1974, 190–192).

It is difficult to extract from this discussion a canonical solution to the *boru* 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 Rumelian and the Anatolian *boru*s 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 *boru* problem will not work in a completely deterministic world. This represents, I think, a fundamental weakness 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 Rumelian and the Anatolian *boru*, and clearly, the question of determinism has no bearing on the distinctions we make.

Why do we deny that the Rumelian *boru* causes the Anatolian janissaries' assembly? Mackie is correct, I think, when he points to our beliefs about continuous causal processes as lying at the root of the denial. There is no causal process linking the sounding of the Rumelian *boru* to the Anatolian janissaries' actions, whereas there is such a process linking the sounding of the Anatolian *boru* to the Anatolian janissaries' actions. It is for this reason, I suggest, that although the sounding of the Rumelian *boru* is an INUS condition for both the Rumelian and the Anatolian janissaries' assembly, it is right to say that the Rumelian *boru* caused the Rumelian janissaries' assembly, but wrong to say that the Rumelian *boru* caused the Anatolian janissaries' assembly.

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 causally 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 amendment in turn.

First, the source of the facts about causal connection. We hold that the sound of the Rumelian *boru* is causally connected to the Rumelian janissaries' assembly, but not to the Anatolian janissaries' assembly. Why? Because physics says that the sound of the Rumelian *boru* reaches the ears of the Rumelian janissaries, but not the ears of the Anatolian janissaries. Thus the Rumelian *boru* is causally connected to the Rumelian janissaries' actions, but not to the Anatolian janissaries' actions. In general, I claim, you can read off the facts about causal connection from the nomological dependencies spelled out in the laws of physics.

This is a controversial view. But there are various ways to extract causal relations from fundamental physics: one strategy is to found the facts about causal connection using the "process" approach to causation (Dowe 2000); another might employ Lewis's recent account of causal influence (Lewis 2000). (Both authors, of course, have higher ambitions for their theories than this.) Rather than endorse any particular approach, let me simply assume as given the primitive facts about causal connections between events and about the laws in virtue of which the connections exist.

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 *boru* problem, is given here; the full answer will be given in section 5.3.

In Mackie's original account, you begin with a set of conditions sufficient to entail that the putative effect *e* occurred. By contrast, I propose that you

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 characterizes 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 occasion.

The *boru* problem is then solved as you would expect. The Rumelian *boru* is not a causal influence on the Anatolian janissaries’ actions, so a set of conditions sufficient for the Anatolian janissaries’ assembly that mentions the Rumelian *boru* is not *causally sufficient* for the assembly. Such conditions do not, then, confer causehood on their non-redundant parts.

Why do I appeal to facts about causation in the middle of an account of facts about causes? What big picture legitimates such a move? In the next section I pause to sketch the role, as I see it, played by causal claims 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.

5.2 *The Role of Causal Claims in Understanding*

The world, according to physics, is a vast and complex causal web. By this paper’s working assumption, it is a deterministic web: the elements and structure 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 whose combined causal influence is sufficient for the region's instantiating the property. This much fundamental physics says, or so I suppose.

If we were fully satisfied with knowledge of the facts about causal influence, 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 you 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, you can trace a line of causal influence from the eunuch to the dying sultan. But what does this tell you? 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 many philosophers suppose. 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 do not agree that this difference-making is itself the fundamental causal relation. What is fundamental is the web of causal influence. This is the province of fundamental physics. The role of causal claims is to single out the elements of the web that are relevant to the high level events that matter to us humans: the breakings, the pleasurings, the dyings, and all the rest.⁹

5.3 *Spurious Non-Redundancy*

Over-reliance on the entailment relation can put a philosophical analysis in real peril. For Mackie’s account of causal claims, the threat is exemplified by the following recipe.

To show that any event r whatsoever was a cause of a given event e : Take a set of conditions jointly sufficient for e , none redundant. 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 e . 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 was a cause of e .

Let me give a specific example. The sultan has been murdered. A set of

9. Here, by the way, is a way of understanding the claim of Hall (2004) and others that there are two different notions of causation at work in our cognitive economy, without concluding that our causal thinking is ambiguous or confused.

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 was 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 *boru* 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 a guide 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 ac-

count 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 Rumelian *boru* 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.¹⁰

Take, to choose the simplest possible example, events c and e and a law *All events of type C cause events of type E*, where c and e are of types C and E respectively. The occurrence of c and the law entail (when fleshed out) the occurrence of e ; but also, this entailment corresponds to a link in the causal chain that produced e , namely, the link between c and 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

10. Talk of "steps in the entailment" implies an intended proof. The proof can be understood as a causal model, as described in Strevens (2004). For the purposes of this paper, say that causal sufficiency requires the existence of at least one proof, intended or otherwise, in which each step corresponds to a strand in the web.

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 $\neg r \vee c$ 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 the occurrence of 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, hence determine the structure of the causal web. 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 you 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 influ-

ence. 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.

5.4 One More Preemption Scenario

In section 4.2 I mentioned a preemption scenario that would be handled later. Now is the time.¹¹ The story is as follows. The sultan's mother is determined that a certain hated Iznik jar should be broken. She stations herself near the jar with a cannonball and resolves that, if the jar is still intact in ten minutes time, she will hurl the cannonball at the jar, of course breaking it. Before the time is up, the prince arrives and breaks the jar himself. Now, just from the fact of the sultan's mother's resolution (and her ability to make good on it), it is sure that the jar will be broken. Therefore, it is possible to construct a set of conditions that is sufficient for the breaking and that mentions only the resolution, saying nothing about the prince. The mother's resolution is non-redundant: it cannot be removed from the conditions without invalidating the entailment of the breaking. It seems that her resolution must be counted as a cause of the breaking.

Although this is a case of preemption, I think that the problem it uncovers in the Mackie account belongs with the Rumelian *boru* and the chief black eunuch's sneeze. In any case, the solution is the same: the sufficient conditions of which the mother's resolution is a non-redundant part are not, given the way the story unfolds, *causally sufficient* for the breaking.

This case is somewhat more complex than the case of the *boru*, since the mother's resolution would have been a straightforward causal influence on the breaking if the prince had never turned up. There are two quite different

11. Christopher Hitchcock has urged the importance of this sort of case for Mackie's account.

ways, then, that the resolution-including sufficient conditions for breaking could be satisfied. One way, the conditions are not just sufficient but causally sufficient: the prince stays away and the resolution causes the mother to hurl the ball. The other way, the conditions are not causally sufficient: they identify neither a cause of the breaking nor a condition, negative or positive, necessary for such a cause to have its effect. In short, a single, apparently quite univocal set of sufficient conditions turns out to be satisfiable by two different causal scenarios; in one case but not the other the conditions identify elements of the causal process in question.

It remains only to say that the definitions of causal entailment and so of causal sufficiency should be understood so as not to apply to such cases, by specifying, for example, that a set of conditions sufficient for an event *e* is causally sufficient only if it identifies aspects of the world that play a role in causally producing *e* in *every one* of its instantiations. It then follows that the identified aspects play a causal role in the actual scenario, as desired.

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. (2000). *Physical Causation*. Cambridge University Press, Cambridge.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*. MIT Press, Cambridge, MA.
- Hall, N. and L. A. Paul. (2003). Causation and preemption. In P. Clark and K. Hawley (eds.), *Philosophy of Science Today*. Oxford University Press, Oxford.
- Hempel, C. (1965). Aspects of scientific explanation. In *Aspects of Scientific Explanation*, chap. 12. Free Press, New York.
- Hitchcock, C. R. (2001). A tale of two effects. *Philosophical Review* 110:361–396.
- 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.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge.
- Ramachandran, M. (1997). A counterfactual analysis of causation. *Mind* 106:263–277.
- Reichenbach, H. (1956). *The Direction of Time*. University of California Press, Berkeley, CA.
- Schaffer, J. (2003). Overdetermining causes. *Philosophical Studies* 114:23–45.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (ed.), *Studies in Logical Theory*. Blackwell, Oxford.
- Strevens, M. (2003). Against Lewis's new theory of causation. *Pacific Philosophical Quarterly* 84:398–412.
- . (2004). The causal and unification accounts of explanation unified—causally. *Noûs* 38:154–176.
- Yablo, S. (2002). De facto dependence. *Journal of Philosophy* 99:130–148.