JONATHAN SCHAFFER

OVERDETERMINING CAUSES*

When two rocks shatter the window at once, what causes the window to shatter? Is the throwing of each individual rock a cause of the window shattering, or are the throwings only causes collectively? This question bears on the analysis of causation, and the metaphysics of macro-causation. I argue that the throwing of each individual rock is a cause of the window shattering, and generally that *individual overdeterminers are causes*.

1. OVERDETERMINATION: INDIVIDUALISM OR COLLECTIVISM?

Terminology (from Lewis, 1986): *Redundant causation* occurs whenever there are multiple actual distinct events *c1*, *c2*, . . . , *cn*, *e*, such that each *cj* without the other *c*s would cause *e*. For simplicity I focus on the case with just two redundant factors, *c1* and *c2*. In such a case, *preemption* (asymmetric redundancy) occurs whenever just one of the *c*s actually causes *e*; *overdetermination* (symmetric redundancy) occurs whenever both of the *c*s are causally on par with respect to *e*.[1]

So when two vandals throw rocks that simultaneously shatter the window, there are three actual distinct events: *c1*, the throwing of one rock; *c2*, the throwing of the other rock; and *e*, the shattering of the window. Here *c1* and *c2* are redundant causes of *e*. And since both *c1* and *c2* are causally on par with respect to *e* (neither rock arrives first, or knocks the other off course, etc.), *c1* and *c2* are overdetermining causes of *e*.

Likewise when two wizards cast spells that turn the prince to a frog at midnight, there are three actual distinct events: *c1*, the casting of one spell; *c2*, the casting of the other spell; and *e*, the metamorphosis of the prince. Here *c1* and *c2* are redundant causes of *e*. And since both *c1* and *c2* are causally on par with respect to

Philosophical Studies **114:** 23–45, 2003.

e (neither spell acts first, or trumps the other, etc.), *c1* and *c2* are overdetermining causes of *e.*

And when properties from two levels lawfully suffice for the same effect, such as when my mental and physical properties both suffice for my subsequent behavior, perhaps there are three actual distinct events: *c1*, my being in mental state *m*; *c2*, my being in physical state *p*; and *e*, my subsequent behavior. Perhaps *c1* and *c2* are redundant causes of *e*. And perhaps both are causally on par with respect to *e*. If so then *c1* and *c2* are overdetermining causes of *e*.[2]

Question: When *c1* and *c2* are overdetermining causes of *e*, what causes *e*? To begin with, *c1* and *c2* have *something* to do with causing *e* – it is not as if *e* just erupts independently. The two rock throwings, for instance, have *something* to do with the window shattering – it is not as if the window just ruptures independently.

So the question becomes: When *c1* and *c2* are overdetermining causes of *e*, how do *c1* and *c2* cause *e*? One plausible answer is that *c1* and *c2* each cause *e individually*. Another plausible answer (the only other I know of, and the only other entertained in the literature) is that *c1* and *c2* only cause *e collectively*, which is to say that (i) *c1* and *c2* do not each cause *e* individually, but (ii) the mereological sum $c1 \oplus c2$ causes *e*.[3]

So the question becomes: When *c1* and *c2* are overdetermining causes of *e*, are *c1* and *c2* each causes of *e* individually, or are *c1* and *c2* only a cause of *e* collectively? I label the first answer *individualism*, and the second *collectivism*. There is also logical room for a *mixed view*, on which some distinction is drawn between one type of overdetermination case which deserves individualistic treatment, and another type which deserves collectivist treatment. But I see little prospect for such a mixed view, since I see no principled and plausible way to draw the needed distinction (§6), and since the arguments offered below cut across any would-be distinction anyway. Mixed views may be (provisionally) ignored.

There is also room for disagreement within the individualist camp as to whether $c1 \oplus c2$ also causes *e* (alongside *c1* and *c2*). I think that $c1 \oplus c2$ does also cause *e*, and some of the arguments I offer below (§3, §6) support this. But I see comparatively little interest in this issue: (i) the individualist has a decent answer to "What causes *e*?" either way; and (ii) $c1 \oplus c2$ is not distinct from *c1*

and from *c2*, so there is no proliferation of redundant factors either way. This issue may be left undecided.

So the question becomes, ultimately: Individualism or collectivism?

Stakes: The question of individualism or collectivism bears on *the analysis of causation*, and *the metaphysics of macro-causation*. As to the analysis of causation, a simple counterfactual analysis, on which *c* causes *e* iff $\sim O(c) > \sim O(e)$ (had *c* not occurred then *e* would not have occurred), entails collectivism. Had one of the rocks not been thrown the window would still have shattered. A simple nomological analysis, on which *c* causes *e* iff $O(c) \& L \rightarrow O(e)$ (the conjunction of the occurrence of *c* and the laws of nature entails the occurrence of *e*), entails individualism. Each rock throw is lawfully sufficient for a shattering. Different variants of these views have different implications. For instance, if the simple counterfactual account is extended to allow *quasi-dependence* (Lewis, 1986: *e* quasi-depends on *c* iff *c* and *e* are part of a process, the great majority of intrinsic duplicates of which stand in the ancestral of counterfactual dependence to each other), then individualism follows. Whereas if the simple nomological account is restricted to require *non-redundancy* (Mackie, 1965: *c* is non-redundant for *e* iff no minimally sufficient condition for *e* actually occurs in which *c* does not feature), then collectivism follows. Thus, arguments favoring individualism or collectivism constrain which analyses of causation are acceptable.

As to the metaphysics of macro-causation, given certain assumptions (see note 2), *macro-causes are overdetermining causes*. Or at least, one response to the question of how the causal efficacy of the macro is compatible with the completeness and closure of micro-causality is via overdetermination (Mellor, 1995; Mills, 1996; Sturgeon, 1998, Schaffer, forthcoming b, *inter alia*). This response presupposes individualism. Given collectivism, overdetermination about macro-causation would not lend the macro any causal efficacy: my behavior would have no mental cause, but only a collective cause involving mental and physical events fused together. Indeed, given collectivism, overdetermination about macro-causation would conflict with the causal closure of the micro: my behavior would have no physical cause either, but only a

collective mental⊕physical cause. Thus, arguments favoring individualism or collectivism constrain which views of macro-causation are defensible.

As important as resolving the causal status of overdeterminers may be, there has been little substantive discussion in the literature. Theorists such as Fair (1979), McDermott (1995), Menzies (1996), and Dowe (2000) have intuited in favor of individualism, whereas theorists such as Mackie (1965), Lewis (1986), and Hausmann (1998) have considered the issue intuitively unclear while concluding in favor of collectivism on theoretical grounds.[4] By and large the discussion of the issue has yet to progress beyond the mere brutings of intuitions and applications of preferred theories (with the notable exceptions of Loeb (1974) and Mellor (1995, §3 and §6)). And the issue has been clouded by claims that overdetermination is strictly impossible (Bunzl, 1979, §2) or at least severely improbable (Kim, 1998, §2).

Roadmap: In what follows I argue, first, that overdetermination is a genuine and widespread phenomenon (§2). I then offer four arguments for individualism: arguments from the theoretical role of causation (§3), from the presence of complete processes (§4), from the source of collective powers (§5), and from the pragmatics of causal discourse (§6).

2. OVERDETERMINATION EVERYWHERE

The question of "Individualism or collectivism?" has been clouded by claims that overdetermination is strictly impossible, or at least severely improbable. I aim to clear the air. There is nothing, I argue, secretly contradictory or otherwise suspicious about overdetermination – in fact, overdetermination is *everywhere*.

Impossible? Bunzl (1979) argues that overdetermination is impossible, as follows. Given (1) the actual laws of nature, and (2) maximal modal fragility for events,[5] Bunzl concludes that each alleged overdeterminer is really necessary for the effect. Each rock throwing, for instance, lawfully impacts the manner in which the window shatters, by contributing a certain quantity of rock; and so without one rock throwing, what would have occurred would have been some *different* window shattering e^* (and so e would not

have occurred). Thus Bunzl concludes that (3) all alleged overdetermination cases are really cases of *joint causation underdescribed*: each rock is as necessary for the actual window shattering as the striking of the match and the presence of oxygen are necessary for the ignition.

I agree with Bunzl that *given* the actual laws and *given* maximal modal fragility, the *two-rocks case* is really a case of joint causation underdescribed. But I reply, first, that the limitation to actual laws should be rejected. As any child can attest, we have perfectly clear intuitions about causation involving spell castings, and other fairy tale devices. Rock throwings and spell castings have something in common as instances of causality, and an analysis of causation should explain why. Without the chauvinistic limitation to actual laws, the two-spells case serves as a possible case of overdetermination.

I reply, second, that maximal modal fragility should be rejected, for conjuring spurious causality (Lewis, 1986, pp. 198–199). The preempted backup, the innocent bystander, and the hound baying in the distance will all presumably make some tiny difference to exactly when and how the murder is committed, through the contribution of a bit of light or sound (–on this poor test, every event in the back lightcone of e winds up counting as a cause of e). Without the punctilious insistence on maximal modal fragility, the two-rocks case serves as an actual case of overdetermination.

I reply, third, that even with the actual law limitation and maximal modal fragility, the two-levels case may still serve as an actual example of overdetermination, because the model of composition of causes does not apply here. However it is that the micro and the macro may be thought to overdetermine the subsequent effect, it is *not* on the model of each contributing a component force vector. Without the hasty assumption that all causes must compose, the two-levels case may still serve as an actual case of overdetermination.

Improbable? Kim (1989, 1998) maintains that overdetermination, while indeed possible, is severely improbable. Kim claims that two rocks hitting a window at once could only be due to a perfectly timed conspiracy or cosmic coincidence. Kim is primarily concerned to argue that overdetermination is a *poor model for the*

two-levels case, since macro-causation is far too widespread to be explicable via the kind of improbable conspiracies and coincidences that overdetermination requires.[6]

I agree with Kim that the two-rocks case is improbable. But I reply, first, that this has no bearing on the two-levels case. If the two-levels case meets the requirements for overdetermination (see note 2), then it is overdetermination, *end of story*. Given that the two-levels case is far more widespread than the two-rocks case, all that would follow would be that some overdetermination cases are far more widespread than others. And so it would follow that at least one sort of overdetermination is indeed everywhere.

I reply, second, that there are non-question-begging (intra-level) examples of overdetermination that are widespread. When *one big rock* hits the window flying northwards, the rock's eastern and western hemispheres are overdetermining causes of the window shattering.[7] This case is of a sort that Mackie labels *quantitative overdetermination*. Mackie's example is that of a sledgehammer flattening a chestnut:

> [T]he whole of the blow was not necessary for [the flattening of the chestnut], though it was more than sufficient: a somewhat lighter blow would have sufficed. Even if part of the hammer-head had been absent, this result would still have come about (1974, p. 43).

Quantitative overdetermination occurs whenever the cause is decomposable into distinct and independently sufficient parts.[8]

Quantitative overdetermination involves no conspiracy or coincidence (at least in the relevant sense in which such are severely improbable). Rather the explanation for why quantitative overdetermination occurs, as opposed to each part acting alone, is that the overdetermining parts are *lawfully yoked*: there are forces that hold the hemispheres of the big rock and the halves of the hammerhead together. And so some cases of overdetermination are due not to improbable conspiracy or coincidence, but rather to everyday lawful correlation. Such overdetermination is everywhere.[9]

To return to the two-levels case, when any set of higher-level properties is instantiated, presumably a subvening set of lower-level properties is also instantiated. And perhaps these are distinct existences that are lawfully correlated. On this model there is, as it were, *one big person* with a *physical part* and a *psychological part*

that are lawfully yoked together. If this model is apt then the best precedent for two-levels overdetermination is the lawful correlation model, rather than the conspiracy or coincidence models. The one-big-rock case thus provides an intra-level precedent for widespread inter-level overdetermination.

3. THEORETICAL ROLE

Causation is among the most central of our concepts, and part of the point of analyzing causation is to shed light on such connected concepts as *prediction*, *explanation*, *manipulability*, and *responsibility*. To a crude approximation, causes license predictions of their effects, causes serve to explain their effects, causes serve as means to manipulate their effects, and causes bestow moral responsibility on agents for their foreseeable effects. While these principles no doubt need refinement, I submit that any event that satisfies every one of these principles (relative to a given effect) deserves to be considered clearly cause-worthy – that event plays all the roles causes play.

Here then is my first argument for individualism: *individual overdeterminers play the predictive, explanatory, manipulative, and moral roles of causes*. Starting with prediction, knowledge that *c1* occurs is sufficient to license a prediction to the occurrence of *e*. For instance, knowledge that rock1 is thrown at the window (on an accurate trajectory, with enough force) is sufficient to license a prediction that the window will shatter.

Turning to explanation, citing the occurrence of *c1* serves to explain the occurrence of *e*. For instance, if one wonders why the window shattered, then it seems that "because rock1 was thrown at it" suffices to answer the question (–one should no longer feel surprised that the window shattered), even if other factors like the presence of rock2 go unmentioned.

Moving on to manipulation, were *c1* different in any of many ways then *e* would have been different in any of many ways. For instance, if rock1 were thrown faster in any way, then the window would have shattered at a correspondingly earlier time, and if rock1 were thrown more forcefully in any way, then the window would have shattered in a correspondingly more forceful manner. Thus there is counterfactual variation between individual overdeterminers

and their effects; in a way that Lewis (1986, p. 165) points out is typical of such intuitively causal processes as measurement, perception, and control.[10]

Turning finally to responsibility, an agent who performs *c1* is liable to praise or blame for *e*. For instance, each individual rock-thrower would be held morally responsible for an act of vandalism. If the first vandal were to plead before the court, "You can't blame me – I didn't do it" I doubt the judge would be swayed[11] (intuitions of moral responsibility can in turn be buttressed by intuitions about the appropriateness of *remorse*, and here I suggest that it would be appropriate for each individual vandal to feel remorse).

Loeb: Louis Loeb (1974), however, has invoked theoretical role considerations to argue against wholesale individualism. Loeb's overall position is that (1) theoretical role considerations bearing on *remembering* and *inferring* demand individualism, but (2) considerations bearing on *killing* and *coercing* demand collectivism, so that (3) we need two concepts to split the difference, one of which (*c-condition*) is individualist, and the other of which (*causation*) is collectivist.

I agree with Loeb that *remembering* and *inferring* demand individualism, and pause to rehearse his arguments.[12] As to memory, Loeb (i) formulates a plausible *necessary* condition on remembering involving causality: *s* remembers *p* from time *t* only if *s*'s knowing that *p* at *t* causes *s*'s present knowledge that *p* (1974, p. 526); he then (ii) *intuits* that, in an overdetermination case in which *s* recalls an acquaintance's name just as he hears his friend whisper it, *s* should count as remembering; from which he (iii) concludes that an *individualistic* concept is needed to play the role of causes in remembering. As to inference, Loeb presents a parallel argument. He (i) formulates a plausible *necessary* condition on inferring involving causality: *s* inferred *q* at time *t* from *p* only if *s*'s believing *p* prior to *t* causes *s*'s believing *q* at t (1974, p. 527); he then (ii) *intuits* that, in an overdetermination case in which *s* reasons that the dog must be outside just as he sees the dog out the window, *s* should count as inferring; from which he (iii) concludes that an *individualist* concept is needed to play the role of causes in inferring.

But I deny Loeb's claim that *killing* and *coercing* demand collectivism. As to killing, Loeb (i) formulates a plausible *necessary*

condition on killing involving causality: *s* kills *p* only if *s* causes *p*'s death (1974, p. 528); he then (ii) *intuits* that, in an overdetermination case involving two assassins, "it is not clear that *s did kill p*" (p. 529); from which he (iii) concludes that a *collectivist* concept is needed to play the role of causes in killing. As to coercion, Loeb presents a parallel argument. He (i) formulates a plausible *necessary* condition on coercing involving causality: *s* coerces *p* into not doing *a* only if *s*'s threatening causes *p*'s not doing *a* (p. 529); he then (ii) *intuits* that, in an overdetermination case involving two hoodlums, "it is not clear" that *s* coerced *p* (p. 529); from which he (iii) concludes that a *collectivist* concept is needed to play the role of causes in coercion.

I reply to Loeb's arguments for collectivism that, first, that even if it is intuitively *unclear* that some individual overdeterminers can kill or coerce as Loeb maintains in (ii), this in no way entails that it is *wrong* to count them as such (actually, my own intuitions are that the individual assassins should count as killers, but I do not wish to insist on this[13]). An 'unclear' verdict is as compatible with individualism as it is with collectivism (this should be contrasted with the clear positive verdict Loeb proffers for remembering and inferring).

I reply, second, that even *if* it were intuitively wrong to count these individual overdeterminers as killing or coercing, this in no way conflicts with the plausible *necessary* conditions for killing and for coercing that Loeb formulates in (i). It merely entails that these conditions are not *sufficient* for killing and for coercing, as was already obvious (this should be contrasted with the genuine conflict between it being intuitively right to count individual overdeterminers as remembering and inferring, and Loeb's plausible necessary conditions for remembering and inferring).

(I also deny that Loeb's concluding position splits any differences. Loeb's definition of *causation* is collectivist. I see no mitigation in 'individualism' about *being a c-condition*, since that invention has no conceptual role to play, and does not even appear in any of the plausible necessary conditions which Loeb himself formulates.)

Mellor: Mellor (1995) invokes theoretical role considerations to defend wholesale collectivism. Mellor argues that (1) *prediction*, *explanation*, and the agential *means* relation require *probability-*

raising; (2) individual overdeterminers do not raise the objective token probabilities of their alleged effects (since the presence of the other overdeterminer already drives the objective token probability of *e* all the way to 1); so (3) individual overdeterminers fail to play the predictive, explanatory, and agential roles of causes.[14]

I agree with Mellor that individual overdeterminers do not raise the objective token probabilities of their effects. But I reply, first, that Mellor's conclusion is already refuted by *preemption* (the *a*symmetric species of redundant causation: §1). There can be both probability-lowering causes and probability-raising non-causes, in cases in which an unreliable source preempts a reliable source (Good, 1961; Menzies, 1989, *inter alia*). For instance, if the first vandal is weak of hand but strong of heart, while the second vandal is strong of hand but weak of heart, then the first vandal's throw might prevent the second vandal from throwing and thereby lower the overall chance that the window will shatter, while shattering it all the same. So the fact that individual overdeterminers do not raise the probabilities of their alleged effects shows nothing.[15]

I reply, second, that Mellor's probability-raising requirement on prediction, explanation, and the means relation represents an *over-simplification of all these concepts.* Prediction is a wholly epistemic concept. Explanation has an epistemic aspect concerning reducing surprise and rendering events more comprehensible, an ontological aspect concerning subsumption under generalizations (Hempel and Oppenheim, 1948), and an ontological aspect concerning tracing processes (Salmon, 1984). The means relation has an epistemic aspect concerning subjective rationality, an ontological aspect concerning general rules for effective behavior, an ontological aspect concerning the expectations of specific acts, and an ontological aspect concerning the retrospective impacts of specific acts (Beebee, 1997). All of these aspects of all of these roles (save the ontological aspect of the means relation concerning the expectation of specific acts) can diverge from the raising of objective token probabilities.

Once this multiplicity of aspects is noted, it emerges that individual overdeterminers are capable of playing *virtually every aspect* of the predictive, explanatory, and agential roles of causes, and of doing so *in exactly the same ways as preempting causes*. Epistemic-

ally, the subject unaware of either rock should raise her expectation of a window shattering on the information that one rock (preempting or overdetermining) was thrown at the window. This enables both overdeterminers and preemptors to play the predictive role of causes, to fit that aspect of the explanatory role concerning reducing surprise and rendering the world more comprehensible, and to fit that aspect of the agential role concerning rationalizing behavior from the subject's perspective. Type-wise, rock-throwing events (which both preempting and overdetermining throws fall under) are positively correlated with window-shattering events. This enables both overdeterminers and preemptors to fit that aspect of the explanatory role concerning subsumption under generalization, and to fit that aspect of the agential role concerning general rules for effective behavior. Process-wise, there is a complete path from rock (preempting or overdetermining) to window. This enables both overdeterminers and preemptors to fit that aspect of the explanatory role concerning tracing physical processes, and to fit that aspect of the agential role concerning retrospective impacts of specific acts.

And so Mellor's mishandling of preemption and collectivism about overdetermination are of a piece. It is no argument against individualism that an account of causation and its roles that cannot accommodate preempting causes cannot accommodate individual overdeterminers either.

4. COMPLETE PROCESSES

While the final analysis of *causation* remains elusive, some decently reliable generalizations have emerged. Here is a decent quasi-sufficient condition for causality: *c* causes *e* if *c* and *e* are connected by a *complete process*. While this quasi-sufficient condition no doubt needs refinement, I submit that most event-pairs that meet this condition deserve to be considered causally related – they stand in a causally suggestive relation.

Here then is my second argument for individualism: *individual overdeterminers and their effects are connected by complete processes*. There is a complete process from the throwing of the rock, through the various stages of the flight of the rock, all the way into the window – no needed intermediaries go missing.

What exactly is a complete process? When dealing with rocks, it will do to think of processes as energy-momentum flows. There is a complete energy-momentum line from the throwing of the rock all the way into the window. But when dealing with spells or with macro-properties, where the connection is magical or macro-level, a more general understanding of processes is needed.[16] Perhaps (as I myself think) the best general understanding of a complete process is in terms of lawful subsumption across sequences of events: energy flows, magical spells, and mental processes can all be understood as lawful sequences. Or perhaps the best understanding of a complete process is in terms of intact sequences of events with a certain intrinsically cause-like pattern: energy flows, magical rays, and mental trains can all be understood as intrinsically intact sequences. No need to decide provided the understanding is suitably general to extend to magical and macro-level processes. On any decent understanding of processes, each individual rock, spell, and property is associated with a complete process – no needed intermediaries go missing.

Quasi-sufficiency: Being associated with a complete process is nearly sufficient for being a cause. There are counterexamples, but none applicable to overdeterminers.

I know of four sorts of counterexamples to the claim that *c* causes *e* if *c* and *e* are connected by a complete process. The first counterexample involves mereological rather than causal connectedness. Every event is connected to itself, and any larger wholes of which it is a part. For example, my writing 'rr' is connected to my writing 'Larry', but the connection is mereological rather than causal (Kim, 1973). I think the best solution to this sort of counterexample is to add a requirement that *c* and *e* be actual distinct events (Lewis, 1986). In any case, however this problem is to be solved, the connection between individual overdeterminers and their alleged effects is clearly not mereological.

The second counterexample involves effects and byproducts rather than causes. For example, if the cue ball impacts the eight and nine balls simultaneously, sinking both, then there is a causal process linking the striking of the cue to the sinking of the eight, and a causal process linking the striking of the cue to the sinking of the nine. But there is also an 'effectual' line linking the sinking of the eight (/nine) to the striking of the cue, and a 'byproduct' line

linking the sinking of the eight to the sinking of the nine (and *vica versa*). I think the best solution to this sort of counterexample is to introduce an independent analysis of the causal order, and add a requirement that *c* be causally prior to *e* (Mackie, 1974; Dowe, 2000), which would effectively force the causal relations to 'follow the arrows'. In any case, however this problem is to be solved, individual overdeterminers are clearly neither effects nor byproducts of their alleged effects.

The third counterexample involves mismatch of properties. For example, if I paint the rock purple and you (for reasons unrelated to the color of the rock) throw it at the window, there is a complete process from my painting the rock purple to the window shattering (Dowe (2000) calls this sort of case *misconnection*). I think the best solution to this sort of counterexample is to add a requirement of essentiality-to-the-process ($\sim O(c) > \sim O(e\text{-process})$), and maintain that the purpleness of the rock is inessential to the window-shattering process: without the painting of the rock the window would have shattered by much the same route. In any case, however this problem is to be solved, individual overdeterminers are clearly not mismatched with their alleged effects.

The fourth counterexample involves minute processes. The preempted backup, the innocent bystander, and the hound baying in the distance will presumably be connected to the victim's death by processes involving photons and other physical minutiae (in Schaffer (2001a) I call this the problem of *traces*). I think the best solution to this sort of counterexample is the above idea of essentiality-to-the-process: these minutiae are inessential parts of sufficient processes, just as a strand of hair is an inessential part of one's body. In any case, however this problem is solved, individual overdeterminers are clearly impactful enough for their alleged effects.

Bringing the discussion together, the following is, to the best of my knowledge, a genuine sufficient condition for causality: *c* causes *e* if (i) *c* and *e* are actual distinct events, (ii) *c* is causally prior to *e*, (iii) *c* and *e* are connected by a complete process, and (iv) $\sim O(c) > \sim O(e\text{-process})$. Individual overdeterminers satisfy this condition. So individual overdeterminers are causes. Or at least, individual overdeterminers stand in a highly causally suggestive relation to their effects.

The collectivist might rejoin as follows: Granted that process-connection is quasi-sufficient for causality, it is also true that counterfactual dependence is quasi-necessary for causality. There is also a decently reliable necessary condition that: *c* causes *e* only if $\sim O(c) > \sim O(e)$. So decent generalizations about causation are equivocal here. Why give the quasi-sufficient condition involving process-connection any greater weight then the quasi-necessary condition involving counterfactual dependence?

I reply that both quasi-conditions (process-connection as sufficient, counterfactual-dependence as necessary) are indeed decently reliable, but that the process-connection condition ought to be considered *more reliable for redundant causation*. One sort of counterexample to the counterfactual dependence condition is preemption: the preempting cause is not necessary for the effect, given the presence of a backup. While the counterexamples to the connection condition are inapplicable to overdetermination, the preemption counterexample to the dependence condition is applicable, especially since preemption and overdetermination are the two species of redundant causation. So the reason why the quasi-sufficient condition involving process-connection deserves far greater weight than the quasi-necessary condition involving counterfactual dependence is that we already know that only the former works in cases of redundant causation.

Preemption: There is an emerging consensus that the best way to treat preemption cases is to test for complete processes: preemptors are causes because they are associated with processes that run to completion, while backups are noncauses because they are associated with processes that are cut short.

Almost all of the leading extant approaches to causation have now converged on the complete process treatment of preemption. This treatment first emerged from the process-view of causation associated with Salmon, Dowe, and Douglas Ehring (1997):

> Many cases of casual preemption can be handled [by tracing causal processes] . . . Suppose that a well-hit ball is traveling on a trajectory that would surely strike and break a certain window pane, if the glass were intact when the ball arrived. As it happens, however, a stray bullet shatters the window just prior to the arrival of the ball. In this case, the ball never intersects the window pane, and thus plays no part in the explanation of the breaking of the window (Salmon, 1997, pp. 474–475).

This treatment has also emerged in almost all sophisticated counterfactual accounts of causation, such as that of Lewis (1986) (*quasi-dependence*), Menzies (1989), McDermott (1995), Ramachandran (1997), and Noordhof (1999):[17]

> It seems true in all genuine cases of causal pre-emption, . . . that the pre-empted processes do not run their full course. . . . All genuine causes, on the other hand, *do* seem to run their full course; indeed, they presumably count as genuine precisely because they do so (Ramachandran, 1997, p. 273).

Hybrid counterfactual-process accounts have incorporated this insight (Fair, 1979; Schaffer, 2001a): "[T]he preemptor comes out a cause because it is an essential part of the *e*-process: without the preemptor, even though *e* still occurs it is by a different process entirely, namely, by the backup process" (Schaffer, 2001a, p. 87).

The collectivist might rejoin that this argument is at best programmatic: I have not established a final answer to preemption here, and I have not argued against the possibility of other (collectivist-friendly) solutions. Why be hasty?

I agree that the argument is programmatic, but I reply that the program is promising. Or at least, the program is far more promising than anything else going. Given (i) the amount of effort that has gone into addressing the preemption problem, (ii) the convergence of otherwise diverse approaches, and (iii) the paucity of decent alternative treatments, there is at least a strong *prima facie* case for individualism here.[18]

Putting previous points together, deep parallels between preempting causes and individual overdeterminers emerge. First, both are species of redundant causation (§1). Second, both play exactly the same causal roles in exactly the same ways (§3). Third, both are associated with complete processes in a seemingly theoretically salient way (§4). Together, these points mutually reinforce the following case for individualism: *individual overdeterminers deserve the same status as preempting causes*.

5. COLLECTIVE POWERS

Does the collectivist think that *c1* or *c2* individually cause *anything*? Here I think that the collectivist faces a dilemma: she must either say

that (i) *c1* (/*c2*) individually causes *nothing*, in which case the power of *c1* v *c2* is a mysterious emergent power;[19] or say that (ii) *c1* (/*c2*) individually causes *part of e*, in which case it is hard to understand how *c1* individually can fail to count as a cause of *e*.

Here then is my third argument for individualism: *the collectivist cannot offer a stable account of the causal contribution of individual overdeterminers.* Does, for instance, the throwing of either individual rock cause *anything* with respect to the window shattering? If the answer is *no* then it is hard to understand how the sum of the rock throwings could shatter the window. Do fusions of rocks have novel powers not present in individual rocks? This is vastly *implausible*.[20]

If the answer is *yes*, then I ask: What then does each individual rock throwing cause? I expect the collectivist to answer: that *part* of the window shattering associated with the component force that its rock contributes (this is the answer that *explains* why the sum of the rock throwings has the power it has). But now it is hard to understand how the throwing of each individual rock could fail to count as a cause of the window shattering. After all, the component contribution of each rock is significant – so significant, in fact, that it alone suffices to shatter the window.

The collectivist might rejoin that the component contribution of each rock to the shattering is still *not significant enough* for each individual rock throwing to count as a cause of the shattering. But this rejoinder seems unprincipled. When the one normal rock shatters the window (no redundancy involved) of course there are other forces impinging on the window from the surrounding environment. And when the straw breaks the camel's back, that one straw contributes but a straw to the total force bearing on the poor camel's spine.

The right thing to say, I maintain, is that (i) each individual rock throwing does cause components of the window shattering, such that (ii) each individual rock throwing is thereby a cause of the window shattering. But this is individualism.

6. PRAGMATICS

Intuitions about overdetermination turn on *how the case is described*. In particular, *coincidental descriptions* in which the overdeterminers are described separately evoke individualist intuitions,

while *conspiratorial descriptions* in which the overdeterminers are described together evoke collectivist intuitions. For example, if the case is described as involving nine riflemen, one is likely to think individualistically; while if the case is described as involving one firing squad, one is likely to think collectively.[21]

Here then is my fourth and final argument for individualism: *the individualist can explain away collectivist intuitions triggered by conspiratorial descriptions, while the collectivist cannot explain away individualistic intuitions triggered by coincidental descriptions*. The individualist has a precedented and principled pragmatic story to tell as to why conspiratorial descriptions trigger collectivist intuitions. The precedent is our known penchant for selecting *one* cause from among the many, calling it *the* cause, and dismissing the remainder (Mill, 1941 [1846]). As Lewis says: "We may select the abnormal or extraordinary causes, or those under human control, or those we deem good or bad, or just those we want to talk about" (1986, p. 162). Conspiratorial descriptions ('the one firing squad') evoke collectivist intuitions, on this precedent, because they invite us to select the whole as *the* cause, and dismiss the contributions of the parts as not under discussion.

The pragmatic principle behind this sort of causal selection is Grice's (1967) maxim of Relevance, which enjoins the co-operative speaker to stick to the topic. To shift the topic is to *flout* Relevance, and thereby *implicate* a rejection of the preceding claim. If the topic is the causal contribution of the firing squad, then it may be deemed *infelicitous* to shift the discussion to other causes, for fear of implicating rejection.

So when the firing squad executes the prisoner, among the causes of the prisoner's death are (i) the individual shooting of the first rifleman, (ii) the individual shooting of the second rifleman, (iii) the presence of gunpowder in the guns, (iv) the absence of a bulletproof barrier between the squad and the victim, and (v) the restraints that hold the prisoner in place. The individualist claims that our occasional intuitive rejection of (i)–(ii) as causes is of a piece with our occasional intuitive rejection of (iii)–(v) as causes.

The collectivist appears to have no comparable story to tell as to why coincidental descriptions trigger individualistic intuitions. None of Grice's maxims, for instance, look to provide any help at all here.

The only collectivist who has made any attempt to explain away individualistic intuitions is Mellor. Mellor offers two ways "to make [collectivism] more palatable" (1995, p. 102). First, Mellor suggests that one may be misled by one of the following three truths: (i) *c1* and *c2* are each individually sufficient for *e*, (ii) *c1* and *c2* would each cause *e* without the other, and (iii) *c1* and *c2* each entail the collective disjunction O(*c1*) v O(*c2*), which is a cause of *e*. Second, Mellor suggests that one might be ignorant as to whether *c1* truly determines *e*, or rather merely raises the probability of *e* short of 1, and that one's willingness to call *c1* and *c2* causes may reflect a secret belief in indeterminism.

As to Mellor's first claim about potential misleaders, I reply that the alleged misleadingness of all three truths is (thematically enough!) disproved by *preemption*. If one were prone to be misled by any of the truths he cites, one would thereby be prone to labeling preempted backups as causes, since (i) the preemptor and the backup are each individually sufficient for *e*, (ii) the preemptor and the backup would each cause *e* without the other, and (iii) the fact that the preemptor occurs and that backup occurs each entails the collective disjunction. Yet no one seems tempted to label preempted backups as causes.

As to Mellor's second claim about secret indeterministic sympathies, I reply that if anything we naturally think of causation as deterministic, and if anything our secret sympathies are likely to lie with determinism.[22] And in any case such ignorance should not detain the competent philosopher who appreciates the *stipulation* that the case in hand is *deterministic*.[23]

I conclude that (i) overdetermination is everywhere, and (ii) individual overdeterminers are causes. If so then (i) the way is barred for accounts of causation that entail collectivism, and (ii) the way is open for treatments of macro-causation that involve overdetermination.

NOTES

* Thanks to Michael Fara, Chris Hitchcock, Carolina Sartorio, and to audiences at the Werkmeister Conference on Causation and Free Will and at the Bellingham Summer Philosophy Conference.

[1] When there are three or more redundant factors, there may also be what I would call *overpreemption*, in which more than one of the *c*s preempt one or more of the other *c*s. Overpreempted backups are like preempted backups in not being causes, while overpreempting causes are like overdetermining causes in sharing whatever causal status the latter may have.

[2] The three "perhaps"s of the paragraph in the main text flag three worries with treating the two-levels case as overdetermination. First, *c1* and *c2* must be (not just different but) *distinct*. Here one might worry that the lower-level physical event metaphysically necessitates the higher-level mental event, so as to preclude distinctness. Second, *c1* and *c2* must be *redundant* vis-à-vis *e*. Here one might worry that the mental event would not cause anything without the physical event, perhaps because the causal powers of macro-events are parasitic on the causal powers of micro-events (Kim, 1998). Third, *c1* and *c2* must be *symmetric* with respect to causing *e*. Here one might worry that the mental event is preempted by the physical event, perhaps because the candidate causes at each level are in competition for some sort of exclusive status (Kim, 1989). Putting this together, the overdeterminative treatment of the two-levels case requires a strongly non-reductive metaphysic, on which the properties of each level are metaphysically independent existences with their own causal powers, and on which the relation between levels is one of mere nomological correlation via contingent 'bridge laws'. To discuss this further would plunge us deeper into the metaphysics of reduction and of the levels of nature than space permits here (I discuss these issues in my forthcoming a, b, and c). Rather here I focus on the question of what would follow *if* the two-levels case fits the model of overdetermination.

[3] Mereology is not mandatory – one might take the collective cause to be the logical disjunction O(*c1*) v O(*c2*) instead. The issue of how to understand collective causes turns on the issue of what the relata of causation are. If the relata of causation are concrete existences like events then collectives should be understood as mereological sums, while if the relata of causation are abstract language-like entities like facts then collectives should be understood as logical disjunctions. I prefer the eventive-mereological interpretation, and employ it in the main text for the sake of definiteness, but nothing here turns on this.

[4] Hausmann says that collectivism is the "most common approach in the literature" (1998, p. 263).

[5] Modal fragility measures the stringency of counterpart relations for events, with respect to time and manner (Lewis, 1986, pp. 196–197). Maximal modal fragility is the rule that: Counterpart (*e1*, *e2*) only if $\text{time}_{w1}(e1) = \text{time}_{w2}(e2)$, and $\text{manner}_{w1}(e1) = \text{manner}_{w2}(e2)$. Minimal modal fragility places no time or manner restrictions on event counterpart-hood. Intermediate views place intermediate restrictions.

[6] Kim offers two other objections to the overdetermination model for mental causation: (i) it renders mental causes redundant, (ii) it entails that the counterfactual supposition of a mental event occurring without any physical basis would take us to a world at which the physical is not causally closed. I simply fail to see anything objectionable in either consequence.

[7] And the rock's northern and southern hemispheres are a preempting cause and a preempted backup, respectively, of the window shattering.
[8] As Michael Strevens pointed out to me, engineers seek out quantitative overdetermination for safety (structural redundancy). Thus ten columns may support the roof, when any five would suffice.
[9] In fairness to Kim, he is aware of the possibility of lawfully correlated overdetermination (1989, p.247). But he claims that such correlation would obviate the needed distinctness between the overdeterminers. Not so: the overdeterminers may be *metaphysically* distinct, and merely *nomologically* correlated.
[10] Lewis (2000) develops the idea of counterfactual variance into a new account of causation as influence, which is individualistic for the reasons provided in the main text.
[11] Each individual vandal would not *merely* be blamed for the lesser crime of *attempted* vandalism. That, it seems, is the most collectivism could allow. As A.M. Honore explains: "In reply it is argued (Mackie) that in these cases all the agencies that are singly or jointly sufficient for the outcome together constitute its cause. But in law that does not solve the problem because, unless the agents are acting in concert, the responsibility of each agency has to be independently established" (2001, p. 6).
[12] Loeb's arguments from remembering and inferring help buttress the role argument for individualism. But to my mind they are merely buttresses: prediction, explanation, manipulation, and responsibility are far more central roles.
[13] Honore says: "If two huntsmen independently but simultaneously shoot and kill a third person, . . . it is intuitively clear that each should be held responsible for the death . . ." (2001, p. 5).
[14] Mellor also argues that probability-raising suffices for prediction, explanation, and agency, and so concludes that causation, prediction, explanation, and agency are merely distinct guises of one and the same probability-raising relation.
[15] Mellor does not discuss preemption, but the severity of the preemption problem for his overall position has been noted by several commentators, including Dorothy Edgington (1997), Helen Beebee (1998), and Paul Noordhof (1998).
[16] There is a tradition, running from David Fair (1979) through Salmon (1984, 1994) to Dowe (1992, 2000), of attempting a reductive identification of causal processes within an antecedently assumed reductive-physicalist ontology. I have no objection to the question: what is the actual world, physical-level realizer of causal processes? But this question should not be confused with distinct questions such as: What across all possible worlds are the conditions for something being a causal processes? And: Are there chemical-/biological-/psychological-level realizers of causal processes as well? It is no objection to magical and macro-level processes that they are not included in attempts to reductively identify processes with actual physical phenomena.
[17] Perhaps the most notable exception is Lewis 2000 (though see my 2001b for a critique of Lewis 2000 based exactly on the lack of any notion of process). In any case Lewis's influence account is individualistic (§3), so even if influence were to

yield the best treatment of preemption, the argument to individualism would still go through, albeit by a different route.

[18] In this context it is worth noting that those theorists who have preferred collectivism on theoretical grounds (Mackie, 1965; Lewis, 1986; Mellor, 1995; Hausmann, 1998, §1) have stood on theoretical grounds that have been undermined by preemption.

[19] If one thinks of collective causes as logical disjunctions rather than mereological sums (see footnote 3) then the mystery only deepens. It seems like a plausible principle that the causal powers of a disjunction are inherited from the causal powers of its disjuncts (Armstrong, 1978, 1997). – How could (p v q) have powers that neither p nor q possesses?

[20] This is not to endorse any metaphysical principle about the causal relations of parts and wholes. I think this is purely a nomological matter in ordinary cases. And even then, as Carolina Sartorio pointed out to me, it is not clear what to think about extraordinary cases like omissions. My point is simply that, *at actuality*, it is implausible that fusions of *rocks* should have novel powers not present in their parts. Magic spells and omissions may have somewhat spooky and counterintuitive powers, but rocks don't.

[21] This is part of why a mixed view (§1) is so implausible. It is not as if there is any ontologically interesting difference between the cases which trigger individualist versus collectivist intuitions. Rather our intuitions turn on superficial matters of description.

[22] Here it is worth noting that Reichenbach (1956) is generally credited with being among the first philosophers to allow that *indeterministic causation* is not a contradiction in terms.

[23] Mellor actually offers a third suggestion, based on his *claim* that his account is not purely collectivist, but rather delivers a mixed view, going individualistic when the overdeterminers share a common cause (Mellor then goes on, pp. 103–105, to claim that his account goes individualistic for mental-physical overdetermination). I think Mellor's claim is mistaken – his account really is purely collectivist. Mellor's reason for claiming that his account goes individualistic given a common cause *b* is that he thinks that this makes *c1* a probability-raiser of *c2*, since had *c1* not occurred then *b* would not have so *c2* would not have. But this reasoning would, for Mellor, entail that *c1* causes *c2* when by stipulation they are rather correlates of a common cause. It would also entail that *c2* equally causes *c1* (a tight causal loop!) when Mellor himself argues that causal loops are impossible. Really Mellor has slipped and mistakenly invoked a 'backtracking' counterfactual. The right thing to say, if the counterfactual account of causation is to have any hope, is that had *c1* not occurred then both *b* and *c2* still would have (Lewis, 1986, pp. 170–171).

REFERENCES

Armstrong, D.M. (1978): *A Theory of Universals*, Cambridge: Cambridge University Press.

Armstrong, D.M. (1997): *A World of States of Affairs*, Cambridge: Cambridge University Press.

Beebee, H. (1997): 'Taking Hindrance Seriously', *Philosophical Studies* 88, 59–79.

Beebee, H. (1998): 'Do Causes Raise the Chances of Effects?', *Analysis* 58, 182–190.

Bunzl, M. (1979): 'Causal Overdetermination', *The Journal of Philosophy* 76, 134–150.

Dowe, P. (1992): 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science* 59, 195–216.

Dowe, P. (2000): *Physical Causation*, Cambridge: Cambridge University Press.

Edgington, D. (1997): 'Mellor on Chance and Causation', *British Journal for the Philosophy of Science* 48, 411–433.

Ehring, D. (1997): *Causation and Persistence*, Oxford: Oxford University Press.

Fair, D. (1979): 'Causation and the Flow of Energy', *Erkenntnis* 14, 219–250.

Good, I.J. (1961): 'A Causal Calculus I', *British Journal for the Philosophy of Science* 11, 305–318.

Grice, H.P. (1967): *Logic and Conversation*, William James Lectures: Harvard University.

Hausmann, D. (1998): *Causal Asymmetries*, Cambridge: Cambridge University Press.

Hempel, C. and Oppenheim, P. (1948): 'Studies in the Logic of Explanation', reprinted in J. Pitt (ed.), *Theories of Explanation* (1988) (pp. 9–50), Oxford: Oxford University Press.

Honore, A.M. (2001): 'Causation in the Law', *Stanford Encyclopedia of Philosophy*, at http://plato.stanford.edu/entries/causation-law/.

Kim, J. (1973): 'Causes and Counterfactuals', *The Journal of Philosophy* 70, 570–572.

Kim, J. (1989): 'Mechanism, Purpose, and Explanatory Exclusion', reprinted in *Supervenience and Mind: Selected Philosophical Essays* (1993) (pp. 237–264), Cambridge: Cambridge University Press.

Kim, J. (1998): *Mind in a Physical World*, Massachusetts: MIT Press.

Lewis, D. (1986): 'Causation', *Philosophical Papers II* (pp. 159–213), Oxford: Oxford University Press.

Lewis, D. (2000): 'Causation as Influence', *The Journal of Philosophy* 97, 182–197.

Loeb, L. (1974): 'Causal Theories and Causal Overdetermination', *The Journal of Philosophy* 71, 525–544.

Mackie, J.L. (1965): 'Causes and Conditions', reprinted in E. Sosa (ed.), *Causation and Conditionals* (1975) (pp. 15–38), Oxford: Oxford University Press.

Mackie, J.L. (1974): *The Cement of the Universe*, Oxford: Oxford University Press.
McDermott, M. (1995): 'Redundant Causation', *British Journal for the Philosophy of Science* 40, 523–544.
Mellor, D.H. (1995): *The Facts of Causation*, London: Routledge.
Menzies, P. (1989): 'Probabilistic Causation and Causal Processes: A Critique of Lewis', *Philosophy of Science* 56, 642–663.
Menzies, P. (1996): 'Probabilistic Causation and the Pre-emption Problem', *Mind* 105, 85–116.
Mill, J.S. (1941): *A System of Logic* (orig. 1846), London, Parker.
Mills, E. (1996): 'Interaction and Overdetermination', *American Philosophical Quarterly* 33, 105–117.
Noordhof, P. (1998): 'Causation, Probability, and Chance: Review of Mellor', *Mind* 107, 855–875.
Noordhof, P. (1999): 'Probabilistic Causation, Preemption, and Counterfactuals', *Mind* 108, 95–125.
Ramachandran, M. (1997): 'A Counterfactual Analysis of Causation', *Mind* 106, 263–277.
Reichenbach, H. (1956): *The Direction of Time*, Berkeley: University of California Press.
Salmon, W. (1984): *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.
Salmon, W. (1994): 'Causality without Counterfactuals', *Philosophy of Science* 61, 297–312.
Salmon, W. (1997): 'Causality and Explanation: A Reply to Two Critiques', *Philosophy of Science* 64, 461–477.
Schaffer, J. (2001a): 'Causes as Probability-Raisers of Processes', *The Journal of Philosophy* 98, 75–92.
Schaffer, J. (2001b): 'Causation, Influence, and Effluence', *Analysis* 61, 11–19.
Schaffer, J. (forthcoming a): 'Is There a Fundamental Level?', *Nous*.
Schaffer, J. (forthcoming b): 'Of Ghostly and Mechanical Events', *Philosophy and Phenomenological Research*.
Schaffer, J. (forthcoming c): 'Two Conceptions of Sparse Properties', *Pacific Philosophical Quarterly*.
Sturgeon, S. (1998): 'Physicalism and Overdetermination', *Mind* 107, 411–432.

Department of Philosphy
359 Bartlett Hall
University of Massachusetts
130 Hicks Way
Amherst, MA 01003-9269
USA
E-mail: schaffer@philos.umass.edu