

Kim's Causal Efficacy

Scott Walden
March 2001

For some time now, Jaegwon Kim has been concerning himself with the problem of mental causation.¹ As he acknowledges, his concern is very much of the sort that vexed Descartes. How is it that the mental can causally affect the physical? How can our beliefs and desires causally affect our bodily motions? Suspecting it was Descartes's stringent dualism that got him into trouble, Kim suggests that we adopt a version of reductionism in order to render mental causation intelligible. If mental properties are reducible to physical properties then there's no problem about how the mental affects the physical, for at bottom there really are only physical properties, and few fret over how the physical affects the physical.

However, as Kim is well aware, things are not quite this simple. There are well-known reasons for thinking that mental properties are not reducible to physical properties, reasons which Kim is under an obligation to address. And address he does, to the tune of a substantial portion of his most recent book.² Predictably there has been a response from the anti-reductionist camp, with both Ned Block (1997) and Jerry Fodor (1997) offering reasons why Kim's new reductionism is unacceptable. And thus has begun the latest manifestation of the reductionist-anti-reductionist debate.

Still, as much as Block and Fodor differ from Kim with respect to the issue of reduction, there is an assumption that all three share. In their various publications on the issue, Block, Fodor and Kim each assumes that physical properties can be causally efficacious, and then each in his own way works to show that mental properties can inherit this causal efficacy.³ Kim is merely the most extreme of the three, arguing that nothing short of a mental-physical reduction will suffice in this regard.

Recently, however, others have challenged this common assumption in a striking way.⁴ These philosophers suggest we look at the issue of mental causation, not from a metaphysical point of view that places a premium on establishing the causal efficacy of the mental, but rather from a commonsense point of view that places a premium on simple intuitions regarding day-to-day causal explanation. All we need to do in order to vindicate mental properties, these philosophers suggest, is establish that they often figure in correct causal explanations, arguably a trivial task. For it is very much part of our commonsense conception of the world that 'because' statements such as

[F] Fred's arm went up because he wanted to get the waiter's attention often constitute correct causal explanations.

Predictably, Kim is unmoved by this suggestion. Can the venerable problem of mental causation be solved simply by pointing to such mundane examples? Clearly not, he insists, calling this a "free lunch" strategy, and reminding us that "there are no free lunches in philosophy any more than in real life."⁵

Kim's admonishment notwithstanding, my own sympathies lie with the free-lunchers. It's not the effortlessness of the strategy that I find appealing, but rather the methodological orientation it embodies. I like the idea of taking as our investigative starting point intuitions concerning a commonplace human activity (in this instance, causal explanation) and then postulating only as much by way of metaphysical apparatus (causal efficacy, reduction, etc.) as is necessary to account for those intuitions. But I concede to Kim that more has to be done by way of vindicating mental properties than merely pointing out that they often are involved in correct causal explanations. For an obvious rejoinder to the free-lunch strategy is that causal explanation, once analyzed, will be found to require that the properties involved be causally efficacious. According to this rejoinder, [F] expresses a correct causal explanation only if the want referred to was causally efficacious in bringing about the raising of Fred's arm. And whether this want was indeed causally efficacious (or how it could be) is precisely what's at issue. Free-lunchers are under an obligation to deal with this prima facie powerful rejoinder.

Ideally this would be done by first fully analyzing causal explanation, then fully analyzing causal efficacy, and then finally showing that the former does not require the latter. This won't be accomplished here, however, as there is little consensus regarding the nature of the relevant notions. Instead, here I'll undertake the more modest task of establishing that Kim's understanding of causal explanation does not incorporate his own understanding of causal efficacy. While admittedly success in this regard falls short of dealing conclusively with the rejoinder, finding such a gap at the center of Kim's program would at the very least be highly suggestive.

(I) Kim on Causal Efficacy

Consider two distinct ways of understanding the causal relation, one championed by Kim and the other by Donald Davidson. According to Kim

... the causal relation obtains between a pair of events because they are events of certain kinds, or have certain properties.⁶

He adds that

... we ... need a way of talking about the causal role of properties, the role of properties of events in generating, or grounding, these two-termed causal relations between concrete events.⁷

By way of contrast, Davidson has recently reiterated that

it is events that have causes and effects. ...[I]t makes no literal sense ... to speak of an event causing something ... by virtue of its mental properties, or as described in one way or another.⁸

On Kim's understanding, properties of cause-effect pairs have an important role to play in bringing about instances of causation insofar as they "generate" or "ground" such instances, whereas according to Davidson's understanding properties have some other role to play in our metaphysics. Of these two, it's Kim's understanding that embodies a notion of causal efficacy, as it's natural to label "causally efficacious" those properties of cause-effect pairs that "generated" or "grounded" the instances of causation. So let's put the Davidsonian construal to one side for now, focus on the Kimean construal, and try to get a better grip on what it means to say that some properties "generate" or "ground" instances of causation in this way.⁹

Surprisingly, Kim nowhere provides a more precise characterization of causal efficacy than that which is conveyed by these metaphors. But a more precise characterization is required if we are to determine whether his understanding of causal explanation requires causal efficacy. One place to look for help in sharpening up the notion is in the work of one of Kim's closest allies, Brian McLaughlin, who offers the following:

If c causes e, then c is of some type, P, and e is of some type, Q, P and Q are appropriately related, and c causes e in virtue of this.¹⁰

McLaughlin's contribution is to add another relation to the metaphysical picture, a relation holding, not between events, but between properties instanced by those events. Causally efficacious properties--properties that "ground" or "generate" instances of causation--would be ones that stand in this higher-order relation. To illustrate, suppose that an electrical short causes a house fire, and does so on a Tuesday. The cause event here has at least two properties, that of being an electrical short and that of being a Tuesday-event. Intuitively the properties of being an electrical short and of being a house fire are special in a way that the properties of being a Tuesday-event and being a house fire are not. The former pair, but not the latter, "grounded" or "generated" the instance of causation; the former, but not the latter, were causally efficacious. On McLaughlin's understanding it's the former pair that instances the proposed higher-order relation.

But merely postulating such a higher-order relation falls short of telling us what that relation is, and thus McLaughlin's augmented picture falls short of providing the non-metaphorical understanding of

causal efficacy required. In an effort to remedy this, I propose to run through all available candidate analyses for the office of this new relation, and critically evaluate each.

The first analysis to be considered involves mere counterfactual dependence. Consider the following two options in this regard:

- [A] \underline{c} is \underline{P} and \underline{e} is \underline{Q} , and \underline{c} caused \underline{e} , and in all nearby possible worlds in which \underline{c} is not \underline{P} , \underline{c} does not cause \underline{e} , and \underline{e} is not \underline{Q} .
- [B] \underline{c} is \underline{P} and \underline{e} is \underline{Q} , and \underline{c} caused \underline{e} , and in all nearby possible worlds in which \underline{c} is \underline{P} , \underline{c} causes \underline{e} , and \underline{e} is \underline{Q} .

The idea here is that two properties, \underline{P} and \underline{Q} , are related in the manner sought by McLaughlin if and only if they yield true instances of [A] and [B]. For example, suppose a piece of metal is heated and as a result it expands. Here we have two events, a heating of the metal and its expansion. It's true that if the metal had not been heated it would not have caused the expansion. It's also true that in all nearby worlds in which the metal is heated the expansion takes place. Metal heatings and metal expansions are, on this conjecture, related in the right sort of way and, hence, are causally efficacious properties.

Still, I doubt that Kim or McLaughlin or any other advocate of causal efficacy would be satisfied with such a counterfactual construal. Properties that are merely epiphenomenal relative to instances of causation can yield true instances of [A] and [B], and yet such properties are by definition not causally efficacious. To see that such properties can satisfy [A] and [B], consider again the metaphysical scenario described in the quotation from McLaughlin, but now add to the picture that strongly supervening on \underline{P} is another property \underline{P}' , and strongly supervening on \underline{Q} is another property \underline{Q}' . Suppose further that, while \underline{P} and \underline{Q} are causally efficacious, \underline{P}' and \underline{Q}' are not. \underline{P}' and \underline{Q}' are--to use a manner of speaking current in the literature--mere "epiphenomenal dangles,"¹¹ doing no "causal work" in bringing about the instance of the causal relation holding between \underline{c} and \underline{e} . However, if we plug \underline{P}' and \underline{Q}' into [A] a true instance results, for in those worlds in which \underline{c} lacks \underline{P}' it must (because of the supervenience relation) also lack \underline{P} , and if it lacks \underline{P} then \underline{c} will not cause \underline{e} . And although matters are less clear-cut in the case of [B], there is little reason to think that it would not be satisfied by \underline{P}' and \underline{Q}' as well. Nearby worlds in which \underline{c} is \underline{P}' are likely to be ones in which it is also \underline{P} , and thus in these worlds \underline{c} will cause \underline{e} , \underline{e} will be \underline{Q} , and \underline{e} will be \underline{Q}' therefore as well.

I don't offer this discussion of the counterfactual construal of causal efficacy in the spirit of novelty. Indeed Kim himself has discussed it at length under the heading of "supervenient causation" and, while he at one time flirted with the idea of construing causal efficacy in this way, he has long since distanced himself from it.¹² Given that even Kim rejects this construal of causal efficacy, it seems safe

to move on to another, much stronger candidate, one that analyzes the sought-after relation in terms of causal lawhood.¹³

The idea behind the causal-lawhood analysis is, initially at least, quite straightforward. On this construal the sought-after relation between properties is simply that of standing in a causally lawful relationship. If it's a causal law that $(\forall x)(\underline{P}x \rightarrow (\exists y)\underline{Q}y)$, then \underline{P} and \underline{Q} are causally efficacious properties, and that's that.

Such an approach has a number of appealing features. First of all, it does justice to Kim's metaphors. Suppose that \underline{c} occurs and is \underline{P} . With the causal law in place it then follows with some degree of nomic necessity that another event \underline{e} will take place, an event which will be \underline{Q} , and it seems entirely natural to say that \underline{c} 's being \underline{P} "generated" or "grounded" this instance of causation. Furthermore, this construal can distinguish between genuinely causally efficacious properties and merely epiphenomenal ones. Consider again the metaphysical scenario involving \underline{P} , \underline{P}' , \underline{Q} and \underline{Q}' , but this time imagine that \underline{P} and \underline{Q} are covered by a causal law. If \underline{c} occurs and is \underline{P} then \underline{e} will occur and be \underline{Q} , and, given that \underline{P}' supervenes on \underline{P} and \underline{Q}' supervenes on \underline{Q} , \underline{c} will also be \underline{P}' and \underline{e} will also be \underline{Q}' . But it does not necessarily follow from this that \underline{P}' and \underline{Q}' be covered by a causal law, and, hence, it does not follow from the construal of causal efficacy at hand that they are causally efficacious. And this is precisely the result that advocates of causal efficacy want, given that \underline{P}' and \underline{Q}' are, by hypothesis, merely epiphenomenal.

Tempting as it is, however, Kim does not embrace this nomological construal, and he is wise not to do so, for there are several serious problems facing it. First of all, on this construal if mental properties are to be causally efficacious then they must be covered by causal laws, and it's highly doubtful that in general they are. Consider again Fred, who enters a restaurant, sits at a table, and raises his arm to get the waiter's attention. On the standard mentalistic picture Fred's desire to get the waiter's attention caused his arm to go up. But on the construal of causal efficacy under consideration, if that desire is to be causally efficacious in bringing about the raising of Fred's arm, there must be a causal law covering desires to get the waiter's attention and arm-raising. This is surprising. It might be true that such desires often are accompanied by such arm-raising, but this falls far short of establishing that anything worth calling a causal law is present. And this point in no way turns on the mentalistic character of the example, as a moment's reflection reveals myriad examples of non-mentalistic causation in which we would be hard-pressed to find the requisite covering law. Suppose, for example, that upon leaving the restaurant Fred steps on a banana peel, and this causes him to slip.¹⁴ Do we really want to postulate the existence of a causal law covering banana-peel steppings and slippings? Granted it might

be true that banana-peel steppings and slippings often come together, and it might be true that such co-occurrences are well known, but none of this provides any significant justification for postulating a law covering the two.

And there is a deeper problem than this, a problem I'll only gesture toward here. Consider the distinction between regularity and realist construals of lawhood couched in terms of supervenience of laws, or failure of supervenience of laws, on the distribution of property instances those laws cover. Very roughly, according to the regularity construal, laws supervene on the distribution of property instances they cover and amount to sophisticated summaries of those distributions. In contrast, according to a realist construal, laws fail to supervene on the distribution of property instances they cover and (in some sense) guide those distributions.¹⁵ On the regularity construal laws are metaphysically posterior to the distribution of property instances insofar as there are no such laws unless some properties are instanced. On the realist construal laws are metaphysically prior to the distribution of property instances insofar as there are laws in place which then function to structure the distribution of property instances. Now, causal efficacy is what helps "ground" or "generate" instances of causation, and, while these metaphors are obscure, one does have the sense that, in general, if X grounds or generates Y, X is in some relevant sense prior to Y. 'Precedent grounded the legal decision' seems natural insofar as the precedent was temporally and authoritatively prior to the decision, and 'The rotation of the armature generated the electrical potential' seems natural insofar as the rotation was entropically prior to the potential. Given this, one would expect that an unpacking of Kim's metaphors would require that causal efficacy be metaphysically prior to the instances of causation it helps bring about, and this in turn suggests that any nomological construal of causal efficacy acceptable to Kim would be one that relies essentially on a realist construal of lawhood. In this way there appears to be a link between the nomological construal of causal efficacy and realism about causal laws.

Now there's nothing wrong per se with needing realist laws for the purposes of one's metaphysics. Indeed, the realist construal of lawhood is currently the most popular, at least when it comes to discussions of causal laws in physics. But it's worth noting how poorly the realist construal mixes with the idea that there are causal laws in psychology. It's one thing to believe that there are laws fundamental to the universe that guide the distribution of physical property instances, quite another to believe that there are such things guiding the distribution of properties such as a desire to get the waiter's attention. Intuitively one would think that, if indeed there are laws covering the properties involved in intentional states, their existence is posterior to the evolution of creatures manifesting such states, exactly the opposite of the metaphysical situation demanded by Kim's metaphors.¹⁶

On the whole, then, the nomological construal of the special relation between properties sought by Kim and McLaughlin looks problematic. We have also seen that the counterfactual construal is inadequate, and so it's now time to look elsewhere for a construal of the relation that gives the desired substance to Kim's metaphors. The problem is, however, that there are no other such construals. And this leaves us with a methodological problem, as the plan, recall, is to find out what Kim thinks causal efficacy is, find out what he thinks causal explanation is, and then see if the latter entails the former. How can we proceed if we don't know what the former is? Well, we have learned at least this much. Whatever causal efficacy is exactly, it must be some notion stronger than that which satisfies the counterfactuals of the type at [A] and [B], stronger in the sense that, while it entails the satisfaction of such counterfactuals for causally efficacious properties, it contains some extra ingredient, one that epiphenomenal properties do not satisfy. Knowing at least this much, we can now proceed to Kim's analysis of causal explanation and ask whether it requires any notion stronger than one that satisfies instances of [A] and [B]. If so, then I concede to Kim that the free-lunch strategy is in trouble and he is right that we need to show that mental properties are causally efficacious. If not, then perhaps the free-lunchers are on the right track--mental properties need no more by way of vindication than pointing out that they often figure in correct causal explanations.

(II) Kim's Causal-Explanatory Realism

Kim endorses "causal-explanatory realism," a model of causal explanation according to which the correctness conditions for instances of causal explanation fall into one or the other of two categories, epistemic and ontic. On the epistemic side are familiar requirements such as "... intellectual satisfaction, making things intelligible, dispelling puzzles and apparent inconsistencies, etc."¹⁷ On the ontic side is the requirement that

a certain determinate relationship [hold] between the explanandum and what is adduced as an explanation of it.¹⁸

This "determinate relationship" is causation:

... a causal explanation of E in terms of C is a 'correct explanation' only if C is in reality a cause of E.¹⁹

It is this ontic correctness condition that injects realism into the model of causal explanation:

... the view that explanations must be 'real' or 'objective' in this sense can be called explanatory realism.²⁰

Kim's explanatory realism falls within the parameters set by David Lewis's understanding of causal

explanation, one according to which a successful causal explanation is an act in which an explainer informs an explainee about some aspect of the causal history of the explanandum event.²¹ In general, acts of informing can be resolved into ontic and epistemic components. If Fred is successfully to inform Mary about the weather he must not only state truths about the weather (ontic component), he must as well select from the many such truths those in which Mary is interested, those of which she is ignorant, etc. (epistemic component). According to Lewis and Kim matters are the same with respect to causal explanation, the primary difference between informing in general and informing via causal explanation in particular being that the ontic component associated with the latter activity involves relatively complex causal facts.

Putting Lewis to one side and focusing on Kim, recall from the beginning of Section I that Kim builds causal efficacy into his understanding of causation insofar as causally efficacious properties “ground” or “generate” instances of causation. Given this close connection between causation and causal efficacy, it’s tempting to assume that Kim wants also to build the requirement of causal efficacy into his understanding of causal-explanatory realism. Specifically, the temptation is to assume that Kim factors the ontic correctness condition into two components so that an offered causal explanation is correct only if

[O₁] the explanans event caused the explanandum event

and

[O₂] the explanans and explanandum events were described in terms of the properties that grounded or generated that instance of causation.

Call this construal of Kim’s understanding of causal explanation Strong Explanatory Realism (SER).

And call an alternative understanding, the one that includes O₁ but omits O₂ (and that, therefore, relies on a Davidsonian understanding of causation) Weak Explanatory Realism (WER).

Does Kim indeed endorse SER as opposed to some other version of explanatory realism such as WER? There is no explicit discussion of this topic, but there are good reasons for thinking he does. One such reason is that it would explain why Kim is so impatient with the free-lunchers. If causal efficacy really is a precondition for involvement in correct causal explanation, and if the connection between causal efficacy and causal explanation really is as simple and direct as SER has it, then anyone who accepts SER would regard it as obvious that the causal efficacy of a class of problematic properties has to be established prior to their vindication for causal-explanatory purposes. And anyone thinking along these lines would regard attempts to establish the causal efficacy of properties by simply pointing to their involvement in putatively correct causal explanations as merely question begging--just as does Kim in the

case of intentional properties. Furthermore, it would be surprising if Kim made the effort to articulate his metaphysics to the extent that it incorporates causal efficacy in addition to simple causation, and then not go on to avail himself of this articulation in the construction of his model of realist causal explanation.

This last point occasions methodological reflection, as it's worth dealing explicitly with the question of how to adjudicate between competing models of causal explanation. How are we to decide if something like WER is an adequate model, or if instead we need to embrace something more complex like SER? Presumably the answer is that we take common intuitions regarding the correctness or incorrectness of real or imagined instances of causal explanation and compare them with the predictions regarding correctness or incorrectness entailed by the model under evaluation. Except in rare instances, if there's substantial mismatch in this regard the model gets rejected. Or, if two models are in competition, the model that accounts for the wider range of such intuitions is to be preferred. The situation is analogous to standard methodology in scientific investigation, with the candidate model of explanation taking the place of the scientific theory to be tested, and our intuitions regarding correctness or incorrectness taking the place of experimental observations.

Such methodological considerations make it even clearer why Kim might be implicitly assuming SER, for there are familiar examples that prima facie suggest that SER captures a wider range of causal-explanatory intuitions than does something like WER. For example, suppose again that an electrical short caused a fire, and did so on a Tuesday. The widespread intuition is that

[T] The fire occurred because of a Tuesday-event
suffers in comparison to

[S] The fire occurred because of an electrical short.

SER neatly accounts for this intuition insofar as it assumes that being an electrical short, as opposed to being a Tuesday-event, is a causally efficacious property. WER, by contrast, predicts that the two are on par.

All in all, then, there is good reason for thinking Kim implicitly accepts SER. Indeed these considerations give all of us reasons for joining Kim in embracing SER and the notion of causal efficacy it incorporates. Thus I concede that, unless a way can be found of dealing with such intuitions--a way that does not require expanding our ontology to include causal efficacy--SER will have to be accepted, and the prospects for a free lunch look dim.

(III) Strong, Weak and Moderate Causal-Explanatory Realisms

Causal efficacy, whatever exactly it might be, is an ontic notion. If property P was efficacious in

bringing about an instance of causation, then it was so regardless of whether or not anyone believed that it did, knew that it did, cared that it did, etc. With this in mind an heroic advocate of WER might try to explain away the counterexample facing his model by claiming that the ontic aspect of causal explanation is satisfied in the case of the Tuesday-event description, and then claim further that our intuitive unease is due to violations of requirements on the epistemic side of the model. On this heroic view, it's not that [T] is false, rather it's inappropriate given the epistemic status of a typical explainee. This move would obviate the need to postulate causal efficacy since, whatever epistemic factors engender our intuitions, by definition they won't be ontic ones such as causal efficacy.

But such a move is too extreme for me, and I'll not pursue it here. My own intuition is that [T] is not merely inappropriate but rather false. I concede to Kim that the truth conditions for 'because' statements--the very conditions that comprise the ontic component of the model--are highly sensitive to the properties in terms of which the explanans and explanandum events are described. I concede that we need a basis for discrimination between the various properties of explanans and explanandum events that is at least in part ontic in nature.

Notice, however, that this concession falls short of mandating SER, for it's not true that the only ontic basis we have for such discrimination is the causal efficacy or causal inertness of properties. There is room for a compromise position according to which we do indeed discriminate amongst the various properties of cause events on an ontic basis, but according to which this ontic basis involves something other than causal efficacy. One could, for example, discriminate on the basis of which properties satisfy one or the other (or both) of the counterfactual schemas at [A] and [B] above. Doing so would fall short of embracing SER, for we saw in Section I that satisfaction of these schemas does not require causal efficacy.

Indeed, this is the route I propose to take. I advocate a version of what I call Moderate Explanatory Realism (MER), the view that the ontic component of causal explanation is satisfied only if

[O₁'] the explanans and explanandum events are related as cause and effect
and

[O₂'] the properties in terms of which these events are described yield true instances of the schemas at [A] or [B].

More specifically, I advocate the version of MER that requires only instances of [A] to be satisfied, as I suspect that adding [B] to the ontic component is unnecessary (see below). I thus call the specific position I advocate A-type MER.²²

With matters set up this way, the substance of this discussion turns on whether we need full-

blown SER in order to account for our core intuitions regarding causal explanation, or whether we can make do with the much weaker A-type MER. If the former, then causal explanation requires causal efficacy and Kim is correct, whereas if we can make do with the latter then the free-lunchers would appear to be on the right track.

A-type MER easily handles cases of the sort that eliminate WER from contention. [T] and [S] are pulled apart on this model simply because it's not true that if the event that caused the fire had not been a Tuesday-event it would not have caused the fire, whereas it is true (on any reasonable construal of the truth conditions for counterfactuals) that if the event that caused the fire had not been an electrical short it would not have caused the fire. All this is just as intuition demands, and yet no notion stronger than satisfaction of the counterfactual is invoked.

But there are other imaginable scenarios that might be thought to render A-type MER problematic. Consider cases of causal overdetermination. Suppose that a lightning strike occurs just as the electrical short occurs, and that either occurrence on its own would have been sufficient for the fire. Then the instance of [A] involving the electrical short would be false, as the fire would still occur in nearby worlds in which the short is absent. If someone were to claim that [S] remains a correct causal explanation under such circumstances the result would be a putative counterexample to A-type MER.

Notice, however, that this counterexample is only as threatening as the intuition that [S] remains correct under such circumstances is firm, and this intuition is far from firm. It's certainly not clear to me that [S] remains correct under the scenarios described. If anything, I lean in the direction of thinking that an explanation in terms of the electrical short is inferior under circumstances of overdetermination. But I really have no firm convictions. Consider an analogous case, one involving less fiction. Suppose that in fact Cuban operatives fired from the grassy knoll at the same instant that Oswald fired from the book depository, and that Kennedy's death was in this way overdetermined. Under such a scenario, would it be correct to say that Kennedy is dead because Oswald shot him? It seems to me that there's no clear answer, although I sense that members of the Warren Commission would at the very least be disingenuous if they offered such a 'because' statement in full knowledge of the whole story. At any rate, firm intuitions that such 'because' statements are correct are what would be needed to seriously threaten A-type MER, and these we don't have.

It's worth noting in passing that such softness of intuition suggests that we are dealing in these instances with a familiar problem, that of vagueness of concepts becoming apparent in their application to unusual cases. Typically our concepts are refined only to the degree necessary to deal with commonly encountered, day-to-day scenarios, and as a result they leave us with indistinct intuitions regarding their

application in circumstances not commonly encountered. Our concept of causal explanation should be no exception in this regard, and, given that genuine cases of causal overdetermination constitute highly unusual circumstances, it would be no surprise if intuition went flat when considering them. If this is right then cases of overdetermination will at worst remain neutral between SER and A-type-MER analyses.

So far, then, we are not compelled to adopt SER. But what about the cases of supervenient epiphenomena described in Section I? Recall that we imagined two events, \underline{c} and \underline{e} , where \underline{c} was \underline{P} and \underline{P}' and \underline{e} was \underline{Q} and \underline{Q}' , and \underline{P}' supervenes on \underline{P} and \underline{Q}' supervenes on \underline{Q} , and, finally, \underline{P} and \underline{Q} are causally efficacious whereas \underline{P}' and \underline{Q}' are not. As we saw, \underline{P}' and \underline{Q}' yield true instances of [A] and so, according to A-type MER, they are candidates for involvement in correct causal explanations. But one might have the intuition that causal explanations couched in terms of merely epiphenomenal properties must be incorrect, and if so we have an apparent counterexample to the model.

Here again the intuition is problematic. The problem this time is that the intuition is based on the idea that some properties are epiphenomenal, and the idea that some properties are epiphenomenal in turn entails that some properties (i.e., the ones relative to which the epiphenomenal properties are epiphenomenal) are causally efficacious. The intuition thus implicitly assumes that there is such a thing as causal efficacy, and whether we need to expand our metaphysics to include a notion of causal efficacy (and what exactly causal efficacy is) is precisely what's at issue in this discussion. Thus to take the intuition seriously is to beg the question against the free-lunchers. Indeed, if a free-luncher were to attempt to deal with the intuition the situation would be akin to making the mistake of answering a "loaded question"--in both cases the mere act of answering would implicitly concede to the questioner precisely what the questionee denies. More generally, in the context of this investigation it is not a legitimate move to adjudicate between competing models of causal explanation on the basis of whether or not they satisfy intuitions arising from the notion of causal efficacy itself. The sought-after argument must move from our commonsense understanding of causal explanation in the direction of the theoretical notion of causal efficacy, not the reverse.

Another class of cases that might be thought to militate against A-type MER is one involving overgeneral properties. Consider [S] again, but now compare it with

[P] The fire occurred because of an electrical problem.

It's true that if there hadn't been an electrical problem the fire would not have occurred, and so A-type MER entails that [P] is true. But a common intuition is that [P] is inferior to [S], and so the dictates of A-type MER and those of intuition appear to be at odds. SER, by way of contrast, need not face any such

apparent counterexample, as advocates of the notion of causal efficacy have the option of insisting that the mere property of being an electrical problem (as opposed to the property of being an electrical short) is not causally efficacious, and that it's for this reason that [P] fails where [S] succeeds. More generally, those who avail themselves of causal efficacy can claim that overgeneral properties are never causally efficacious, and then use this to account for the intuition that causal explanations couched in terms of such properties fail generally.

A-type MER can, however, accommodate our intuitive unease with explanations involving overgeneral properties. To see how, recall that A-type MER must be seen as embedded within a larger understanding of causal explanation according to which an act of explaining is an act of informing. The causal-counterfactual aspect of A-type MER constitutes only the ontic component of this larger understanding, a component that is intended to capture the truth conditions of 'because' statements. There remains the complex epistemic aspect of the larger understanding about which nothing has been said. Given this, it's open to the advocate of A-type MER to trace our intuitive unease with explanations couched in terms of overgeneral properties to violations of one or more pragmatic aspects of the overall model. And, indeed, it's plausible that such violations are the source of the trouble.

Intuitively [P] is inferior to [S], but it's not at all clear that [P] is false. Indeed, given that the fire would not have occurred if there hadn't been an electrical problem, [P] appears to be true, and this leaves violations of pragmatic requirements as the only remaining likely source of our intuitive unease. Furthermore, if such violations are the source of our unease, we should detect changes in our intuitions about the example as we ring changes on the epistemic status of imagined explainees, and this appears to be the case. Suppose it was Fred's house that burned down and that he wants a causal explanation of why this happened. And suppose further that Fred is not sophisticated when it comes to electrical problems and, indeed, that he finds such matters tedious--short circuits, overloads, power surges, faulty grounds, corrosion due to aluminum wiring, etc. are all the same to him. But suppose Fred does know that electrical problems, lightning strikes, smoking in bed, playing with matches, etc. are all kinds of events that can cause fires. Relative to Fred, the causal explanation provided by [P] would be at least acceptable and indeed arguably superior to that provided by [S], as Fred wants to know only the general type of cause of the house fire, without being bothered by details regarding the specific type within the culprit category. By way of comparison consider Susan, an electrical engineer hired by the Fire Marshall's Office to investigate the fire, who is au courant regarding the various types of electrical problems, and who has a professional interest in finding out which type of electrical problem was involved in this instance. Relative to Susan [S] will be superior to [P]; indeed, for her, even [S] might

not be good enough, as she'll likely want to know exactly what kind of electrical short was the source of the trouble. Was it a short between two wires? Did a metal object fall onto exposed conductors?

In this example the cause of the fire remains the same and yet our intuitions regarding the offered explanations differ. Since the only relevant difference between Fred and Susan lies in their respective epistemic states, it seems that our intuitions regarding the acceptability of causal explanations are coupled in part to our estimations of the epistemic status of the explainees involved. This in turn suggests that it's the pragmatics of causal explanation that is driving intuitions when it comes to overgeneral properties, thus leaving the advocate of A-type MER a way of dealing with such intuitions without surrendering to SER.

But merely noting that the A-type-MER advocate has a way out falls short of establishing that she is right. What positive reason is there for thinking that the pragmatics-based treatment of intuitions regarding overgeneral properties is superior to the causal-efficacy based treatment? The answer is not far afield. Recall that causal efficacy, whatever its exact nature, is an ontic notion, that is, one not sensitive to the epistemic status of explainees. Given this, the advocate of SER will have a tough time accounting for the variability of our intuitions regarding [P] relative to Fred and Susan. If SER is correct then the property of being an electrical problem is or is not causally efficacious regardless of the epistemic status of parties involved. But in the Fred-and-Susan example the only changes in the imagined scenario are epistemic, and yet our intuitions about the cases differ. So there's a conflict between a central aspect of the notion of causal efficacy on the one hand, and our intuitions arising from the Fred-and-Susan example on the other. A-type MER, by way of contrast, faces no such conflict.

It's worth noting in passing that if this sort of reasoning is sound then postulating causal efficacy in order to deal with puzzles arising from consideration of overgeneral properties amounts to making the mistake of reifying merely pragmatic aspects of the activity of causal explanation. Intuitions arising from consideration of overgeneral properties need to be accounted for. If we are not alert to the possibility that these intuitions are grounded in the pragmatics of causal explanation we might be tempted--in the manner of SER--to build into the truth conditions for 'because' statements the requirement that the properties involved be causally efficacious. This in turn leads us to augment the stock of truth makers by postulating causal efficacy in order to open the possibility that the truth conditions be satisfied. The general pattern of error here will be familiar to anyone who has attempted to reveal the logical form of some problematic class of natural-language statements--belief ascriptions are an excellent example. It should come as little surprise if we fall into the same pattern of error with regard to natural-language statements that provide causal explanations.

Finally, if the inferiority of causal explanations in terms of overgeneral properties can indeed be traced to the pragmatics of causal explanation then a way is opened of settling a minor intramural dispute among MER advocates. The dispute I refer to is that between those who advocate A-type MER on the one hand, and those who advocate the stronger A-and-B-type MER on the other.²³ Those who prefer MER to SER and who are concerned about overgeneral properties but who pay insufficient attention to the pragmatics of causal explanation might conclude that we need to augment the truth conditions for ‘because’ statements in the manner of A-and-B-type MER, for the latter does enable us to account for our intuitive unease with [P] as opposed to [S]. To see how, consider that it’s not the case that in all possible worlds in which Fred’s house suffers from a mere electrical problem that there’s also a fire, for the category of mere electrical problems includes, in addition to electrical shorts, harmless occurrences such as power outages. Given this, on a reasonable construal of the truth conditions for counterfactuals, [B] will be true when ‘electrical short’ is plugged in, but false when ‘electrical problem’ is substituted. This difference can then be used to account for our intuitive unease with [P] as opposed to [S]. If, like the SER advocates, we feel that such unease must be accounted for at the level of truth conditions, but, unlike the SER advocates, we don’t want to postulate causal efficacy, we are compelled to augment the truth conditions by building in instances of [B].

However, the problem with going this route should be obvious given the discussion above. [P] is not false as A-and-B-type MER requires; rather, it’s simply inappropriate given our default assumptions regarding the epistemic status of typical explainees. A-and-B-type MER is thus too strong. A-type MER, by way of contrast, allows [P] to be true, but at the same time has a plausible story to tell regarding the etiology of our intuitive unease. Thus the intramural dispute should be resolved in favor of the advocates of simple A-type MER.

Introducing some new terminology might help forestall such confusions in the future. I suggest we reserve ‘incorrect explanation’ for those instances in which the causal-counterfactual truth conditions are not satisfied, and ‘inappropriate explanation’ for those instances in which the pragmatic requirements have been violated. Thus [S] would be correct and appropriate (given default assumptions regarding typical explainees), [P] would be correct but inappropriate (given those same default assumptions), and [T] would be incorrect (although possibly appropriate).

* * *

The assumption that properties function to “ground” or “generate” instances of causation leads to problems. One such problem was mentioned at the outset, the problem of vindicating mental properties by showing that they can have causal efficacy just like physical properties. Another problem--not

discussed here--is that of "causal/explanatory exclusion."²⁴ If an event can have both causally efficacious mental properties and causally efficacious physical properties, and if that event enters into the causal relation, then the question arises which of the two kinds of properties was doing the "causal work." If both, our ontology is burdened with an unappealing sort of overdetermination. If the former but not the latter, then it's arguable that causal explanations in terms of mental properties are at best inferior to explanations in terms of physical properties, and at worst out-and-out false. A further problem is the one that has been a recurring theme throughout this discussion, the problem of saying in non-metaphorical terms what causal efficacy is. Advocates of the notion have soft-pedaled this problem, but eventually it will have to be faced, and for the various reasons noted above I am not sanguine about the prospects for its analysis in a way that coheres with other domains of metaphysical investigation.

As always with troublemakers at a certain point one has to stop and ask whether they are worth keeping around. If not, if the advantages associated with their presence are outweighed by the troubles they create, then they should be gotten rid of, and I believe that this is precisely the situation with regard to causal efficacy. If my reasoning in this discussion has been sound then even Kim's construal of causal explanation does not require causal efficacy, and the rejoinder to the free-lunch strategy is correspondingly weakened. If there are no other reasons available for thinking that we need causal efficacy then we should simply abandon the notion and enjoy the free lunch.^{25, 26}

Works Cited

- Baker, L. (1993). "Metaphysics and Mental Causation." In Heil and Mele (1993), 75-95.
- (1995). *Explaining Attitudes: A Practical Approach to the Mind*. Cambridge: Cambridge University Press.
- Block, N. (1990). "Can Mind Change the World?" In G. Boolos, ed. *Meaning and Method: Essays in Honor of Hilary Putnam*. Cambridge: Cambridge University Press.
- (1997). "Anti-Reductionism Slaps Back." *Philosophical Perspectives*, 11, 107-132.
- Burge, T. (1993). "Mind-Body Causation and Explanatory Practice." In Heil and Mele (1993), 97-120.
- Davidson, D. (1993). "Thinking Causes." In Heil and Mele (1993), 3-17.
- Fodor, J. (1989). "Making Mind Matter More." *Philosophical Topics*, 17. Reprinted in Fodor (1990), 137-159.
- (1990). *A Theory of Content and Other Essays*. Cambridge, Mass.: The MIT Press.
- (1997). "Special Sciences: Still Autonomous After All These Years." *Philosophical Perspectives*, 11, 149-163.

- Heil, J. and Mele, A., eds. (1993). *Mental Causation*. Oxford: Oxford University Press.
- Horgan, T. (1989). "Mental Quasation." *Philosophical Perspectives*, 3, 47-76.
- Kim, J. (1976). "Events as Property Exemplifications." In M. Brand and D. Walton, eds. *Action Theory*. Dordrecht, Holland: D. Reidel Publishing Co. Reprinted in Kim (1993), 33-52.
- (1984). "Epiphenomenal and Supervenient Causation." In P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy* 9. Minneapolis: University of Minnesota Press. Reprinted in Kim (1993a), 92-108.
- (1988). "Explanatory Realism, Causal Realism, and Explanatory Exclusion." In P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy* 12. Minneapolis: University of Minnesota Press.
- (1989). "Mechanism, Purpose and Explanatory Exclusion." In J. Tomberlin, ed., *Philosophical Perspectives 3: Philosophy of Mind and Action Theory*. Reprinted in Kim (1993a), 237-264.
- (1992). "Multiple Realization and the Metaphysics of Reduction." *Philosophy and Phenomenological Research*, 52. Reprinted in Kim (1993a), 309-335.
- (1993a). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- (1993b). "Postscripts on Mental Causation." In Kim (1993a), 358-367.
- (1993c). "Can Supervenience and 'Non-Strict Laws' Save Anomalous Monism?" In Heil and Mele (1993), 19-26.
- (1993d). "The Nonreductivist's Troubles With Mental Causation." Reprinted in Kim (1993a), 336-357.
- (1994). "Explanatory Knowledge and Metaphysical Dependence." In E. Villanueva, ed., *Philosophical Issues 5: Truth and Rationality*. Atascadero: Ridgeview Publishing Company, 51-69.
- (1998). *Mind in a Physical World : An Essay on the Mind-body Problem and Mental Causation*. Cambridge, Mass. : MIT Press.
- Lewis, D. (1986a). *Philosophical Papers*, vol. 2. Oxford: Oxford University Press.
- (1986b). "Causal Explanation." In Lewis (1986a), 214-240.
- Loewer, B. and LePore, E. (1989). "More on Making Mind Matter." *Philosophical Topics*, 17, 175-191.
- Loewer, B. (1996). "Humean Supervenience." *Philosophical Topics*, vol. 24, 1, 101-127.
- McLaughlin, B. (1993). "On Davidson's Response to the Charge of Epiphenomenalism." In Heil and Mele (1993), 27-40.
- Ruben, D. (1994). "A Counterfactual Theory of Causal Explanation." *Nous*, 28, 4, 465-481.

Schiffer, S. (1991). "Ceteris Paribus Laws." *Mind*, 100, 1-17.

Endnotes

1. See Kim 1998, 1994, 1993b, 1993c, 1993d, 1992, 1989, 1988 and 1984.

2. 1998, especially chapter 4.

3. See Block 1990 and Fodor 1989.

4. See especially Baker 1993 and Burge 1993.

5. Kim 1998, 60.

6. Kim 1993c, 22, emphasis in the original.

7. Kim 1993c, 21, emphasis added.

8. Davidson 1993, 13.

9. In his early writings (1976) Kim defends a fine-grained manner of event individuation which is incompatible with the idea that a single event typically has several properties, some of which are causally efficacious and some of which are not. But in his more recent writings Kim implicitly assumes a coarse-grained manner of event individuation. Additionally, he has recently claimed that the choice between these distinct ways of individuating events (and, indeed, choice between distinct ontological schemes generally) is "by and large optional" (1993a, ix). In what follows, therefore, I treat Kim as embracing the coarse-grained style of event individuation.

10. McLaughlin 1993, 32.

11. See, for example, Kim 1993d, 355.

12. See his discussion at 1993b, 361.

13. Fodor (1989) advocates such an understanding of causal efficacy.

14. The example is taken from Schiffer 1991, 11, where he uses it in a discussion of causal explanation. But the point applies equally to causation, something that's not surprising given the close connection between the two (see below).

15. See Loewer 1996 for an illuminating discussion of these two characterizations of lawhood.

16. There's another, related, problem. If there are laws covering mental properties then they will be special-science laws, and special-science laws, if there are such things, are ceteris paribus hedged. But it's open to question whether the notion of a ceteris paribus-hedged realist law is coherent, for ceteris paribus laws are dependant on particular matters of fact insofar as their ceteris paribus conditions must typically be satisfied, and yet realist laws are supposed to be independent of particular matters of fact. This tension points in the direction of a dilemma for the advocate of the nomological construal of causal

efficacy, who can have either ceteris paribus-hedged psychological laws or laws that do justice to Kim's metaphors, but not both.

17.Kim 1989, 256.

18.Kim 1989, 256.

19.Kim 1989, 256.

20.Kim 1989, 256.

21.Lewis 1986b.

22.Many have characterized and defended versions of what I am calling MER. See, for example, Baker (1995), Horgan (1989), Loewer and LePore (1989), Ruben (1994) and Schiffer (1991).

23.Baker (1995), Horgan (1989), Loewer and LePore (1989) and Ruben (1994) appear to be A-and-B-type-MER advocates; Schiffer (1991) an A-type-MER advocate.

24.See Kim 1988 and 1998.

25.If properties don't function to "ground" or "generate" instances of causation, one might wonder what role they do play in our metaphysics. My suspicion is that they function less as entities that somehow fuel instances of causation and more as things that facilitate our epistemic access to the world. This view is in opposition to what Kim calls "Alexander's dictum," a doctrine according to which "to be real is to have causal powers" (1993d, 348).

26.My thanks to two anonymous referees for this journal who drew my attention to an error in an earlier version of this paper.