

WILEY

Difference-Making in Epistemology

Author(s): Juan Comesaña and Carolina Sartorio

Source: *Noûs*, June 2014, Vol. 48, No. 2 (June 2014), pp. 368-387

Published by: Wiley

Stable URL: <https://www.jstor.org/stable/43828877>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



Wiley is collaborating with JSTOR to digitize, preserve and extend access to *Noûs*

JSTOR

*Difference-Making in Epistemology**

JUAN COMESAÑA
University of Arizona

CAROLINA SARTORIO
University of Arizona

1 Introduction

Difference-making is thought to play an important role in epistemology. It is quite common to suggest, for example, that for a belief to amount to knowledge the world has to make the relevant kind of difference to what is believed. It is also quite common to suggest that for a belief to be justified on the basis of some evidence the evidence must make the relevant kind of difference to what is believed. In this paper we put forth a novel difference-making constraint on evidence and justification—and therefore, given that knowledge entails justification, a constraint on knowledge as well. We motivate such a constraint by means of a parallel with the suggestion that causation is a difference-making relation. In particular, we argue that a recent account of how causes make a difference to their effects can be adapted to explain how evidence makes a difference to justified beliefs. We also show that the proposed difference-making constraint can shed new light on the problem of “easy knowledge.”

2 Difference-Making as Counterfactual Dependence

When does a fact make a difference to another fact? What does it mean to say that a relation that holds between two facts is difference-making? One tempting answer employs the notion of counterfactual dependence:

DM1: R is a difference-making relation if and only if, whenever R holds between some facts F and G, G wouldn't have obtained if F hadn't obtained.

That is to say, F makes a difference to G given the relation R that holds between them when G counterfactually depends on F.

It is often thought that causes make a difference to their effects, and that the difference-making nature of causation can be spelled out in terms of counterfactual dependence (that is, in terms of DM1). On a first pass, at least, this view of causation entails:

Causal DM1: *c* caused *e* only if, had *c* not occurred, *e* would not have occurred either.¹

In other words, *c* caused *e* only if the fact that *e* occurs counterfactually depends on the fact that *c* occurs. For example, the rain caused the flood in that the flood wouldn't have occurred if the rain hadn't occurred. But there are notorious problems with Causal DM1. A main threat is posed by scenarios of causal preemption and overdetermination, where effects don't counterfactually depend on their causes. For example, an assassin can cause the death of his victim in circumstances where the existence of a backup assassin breaks the dependence between his act of shooting and the victim's death (even if he hadn't shot, the victim's death would still have occurred as a result of the backup assassin's shooting). Proponents of the counterfactual theory of causation have attempted to deal with this problem by tinkering with Causal DM1 in different ways, but the difficulties persist.²

So much for the application of DM1 to causation. One could also try applying it to epistemology—and in particular to the question about the conditions that have to obtain for a subject's belief to amount to knowledge. What results is the thought that knowledge requires that the world make a difference to the belief in that, if the world had been different in certain ways, the subject wouldn't have had the belief in question. In other words, the resulting condition is one that specifies that, for some fact (or set of facts) of the world *F*:

Epistemic DM1: *S* knows that *P* only if, had *F* not obtained, *S* would not have believed that *P*.

We can obtain different versions of Epistemic DM1 by varying the fact(s) that *F* stands for. For instance, if we let *F* stand for the fact that *P* itself, we get the well-known sensitivity condition:

Sensitivity: *S* knows that *P* only if, if *P* were false, then *S* wouldn't believe that *P*.³

The sensitivity condition is subject to counterexamples (even on its more sophisticated versions). Notably, it seems to rule out the possibility of inductive knowledge, at least in some cases.⁴ Thus, suppose that at noon one day in July in Tucson, I leave a glass with ice outside. Two hours later, I wonder about that glass and come to think (based on my background knowledge of the weather in Tucson in July) that the ice has thawed. Consider now a situation where it is not true that the ice has thawed. Why would that be? One possible explanation is that the temperature has suddenly sunk to below the freezing point. Or maybe someone keeps surreptitiously adding some powerful anti-thawing agent in the glass. Whatever the possible cause, it is a highly unlikely one. And let us suppose that if one of those highly unlikely situations were to arise I would not become aware of it (I am already far from Tucson, say). Thus, if it were not true that the ice has thawed, I would still believe that it has, and based on the same reasons for which I actually believe it. That is to say, my belief that the ice has thawed is not sensitive. But given, in part, the highly unlikely nature of the possible interferences, we do want to grant that I know that

the ice has thawed (if we are happy with the existence of inductive knowledge at all). Therefore, the sensitivity condition fails.

Another option would be to let the fact *F* in Epistemic DM1 be whatever fact in the world the subject bases her belief that *P* on. The most interesting proposal along these lines has it that, if the subject believes that *P* based on *Q*, then the subject knows that *P* only if, had *Q* not obtained, the subject wouldn't have believed that *P* based on *Q*.⁵ But this view is either trivially true or refuted by slight modifications to counterexamples to the previous sensitivity interpretation. Whether the view is trivially true or not depends on whether we think that when *Q* is false a subject can still base his belief on *Q*.⁶ If we say that he can't, then the principle will be trivially true. And if we say that he can, then the ice-thawing case may be suitably modified to create a counterexample, as follows. In the actual case, the subject believes that the ice has thawed based on the fact that he left it outside in Tucson in July. In the counterfactual scenario, the subject didn't leave the ice outside in Tucson in July, but he nevertheless believes that the ice has thawed based on his having left the ice outside in Tucson in July. Why? Because, unbeknownst to the subject, an evil neuroscientist has implanted a device in his brain that guarantees that he will form the belief that he has left the ice outside even if he actually didn't (we may assume, à la Frankfurt (1969), that if the subject shows any sign of being about to doubt that he has left the ice outside, the implant will force that belief upon him by inducing the required connections in his brain). In the actual case, however, the implant never intervenes, and the subject is a perfect epistemic correlate of someone without an implant of that kind. So he knows that the ice has thawed, in spite of the fact that, had he not left the ice outside, he would still have believed that he did (and on the same basis). So Epistemic DM1 fails, even on this revised interpretation.⁷

We have seen that a natural account of difference-making, DM1, fails to capture the sense in which causation is a difference-making relation, and it also fails to capture the sense in which difference-making is relevant in epistemology. But there is another account of what it means for a relation to be difference-making that is worth investigating. We turn to this account in the next section.

3 An Alternative Account of Difference-Making

The alternative account of difference-making is the following:

DM2: *R* is a difference-making relation if and only if, whenever *R* holds between some facts *F* and *G*, *R* wouldn't have related *F*'s absence to *G*, if *F* had been absent.⁸

Let us start by comparing DM2 with DM1. Note, first, that DM2 is a weaker claim than DM1. If *R* is a DM1-relation that relates *F* to *G*, then, if *F* had been absent, *G* would have been absent too, and thus *G* wouldn't have been the relatum of any relation (in particular, *R*). Therefore, any DM1-relation is a DM2-relation. But a relation can be a DM2-relation without being a DM1-relation. All DM2 requires is that, if *F* had not occurred, *F*'s absence would not have borne relation *R* to *G*. This is consistent with *G*'s occurrence, and hence with the failure of DM1.

Although DM2 is a weaker constraint than DM1, it is still a significant constraint. It is only satisfied if the nature of the relation in question is such that something's bearing the relation to another thing requires that its absence would not have borne that same relation to that thing. Clearly not all relations are like that (for a simple counterexample, consider the relation of temporal succession: F's being temporally followed by G doesn't require that the absence of F wouldn't also have been temporally followed by G). As we will see, when a relation satisfies this constraint, there is an important sense in which the first relatum makes a difference to the second relatum. The first relatum makes a difference, not in the sense that the second relatum wouldn't have occurred in its absence (which is what DM1 says), but in the sense that the second relatum wouldn't have borne the same relation *to* its absence.⁹

One relation that is difference-making in the sense of DM2 is the probability-raising relation. Let us say that F raises the probability of G when the conditional probability of G given F (in symbols, $\Pr(G|F)$) is higher than the unconditional probability of G ($\Pr(G)$). It is a theorem of the probability calculus that, if $\Pr(A|B) > \Pr(A)$, then $\Pr(A|\text{not-}B) < \Pr(A)$ (if B raises the probability of A then not-B doesn't raise—in fact, it lowers—the probability of A).¹⁰ That is to say, probability-raising is a DM2-relation.

There have been probabilistic accounts of both causation and the evidential support relation.¹¹ Thus, some authors have thought (as a first approximation) that c causes e if and only if c raises the probability of e. As was the case with counterfactual accounts of causation, preemption cases also raise problems for probabilistic accounts of causation (and, in particular, for the claim that probability-raising is necessary for causation). Thus, suppose that a quite reliable probabilistic process (say, one that makes the occurrence of an outcome 90% likely) has been launched but then a second, much more unreliable process is started, which cuts the reliable process off at the same time that it issues in the outcome. A simple probabilistic account of causation entails, wrongly, that the reliable process caused the outcome and the unreliable process didn't.

Bayesian epistemologists have, in turn, proposed a probabilistic relevance criterion of evidential support, according to which a proposition A is evidence for another proposition B just in case $\Pr(B|A) > \Pr(B)$, where the probability function in question models the subject's credences. But probabilistic accounts of the evidential support relation must also face serious difficulties. Thus, for instance, they entail that subjects are logically omniscient (with respect to the tautologies determined by the language chosen to model the subject), and thus cannot easily account for logical learning.¹² Moreover, the arguments for a probabilistic account of the evidential support relation themselves are subject to serious problems.¹³

If (despite what these problems seem to suggest) the probabilistic accounts of causation and evidential support were correct, then causation and the evidential support relation would automatically be difference-making relations in the sense of DM2. We believe that the standard problems for such accounts have not been adequately dealt with, and so we will not be assuming that the accounts are correct. However, causation and evidential support could still be DM2-relations. In fact, as

we will see, it is quite plausible to believe that they *are* difference-making relations of this kind. If that were the case, then the temptation to give probabilistic accounts of causation and evidential support could be seen as misguided attempts to capture the fact that those relations are DM2-relations. For example, the fact that causes are difference-makers in the sense of DM2 could help explain why it seemed plausible to believe that causation is just probability-raising, when in fact it is not. And similarly for evidential support: the fact that evidential support is a DM2-relation could help explain why it seemed plausible to believe that providing evidential support to a proposition is just raising its probability, when in fact it is not.

Let us start with the claim that causation is a DM2-relation. To say that causation is a DM2-relation is to say that causes make a difference to their effects, not in the sense that their effects would not have occurred in their absence, but in the sense that their effects would not have been caused *by* their absences (even if those effects would still have been present):

Causal DM2: *c* causes *e* only if, had *c* not obtained, the absence of *c* would not have caused *e*.¹⁴

On this view, it is part of the nature of a cause that its absence could not have had the same effects it has. The contribution a cause makes is unparalleled by the contribution its absence would have made (even if the ultimate effect is the same in both cases). If the contribution *c* makes to *e* doesn't "stand out" in this way, then *c* doesn't make the relevant kind of difference, and thus *c* fails to be a cause of *e*.

Note that Causal DM2 (unlike Causal DM1) is a non-reductive condition on causation (the notion of causation itself appears in the stated condition). Partly due to this fact, it is not subject to the same problems that threaten to undermine Causal DM1. In particular, the possibility of preemption doesn't create the same problem it creates for Causal DM1 (and, more generally, for reductive counterfactual views of causation) because it is plausible to argue that preempting causes meet the constraint: a preemptor makes a contribution that its absence would not have made.

On the other hand, consider a scenario where intuitively we would like to say that an event is not a cause of an outcome because it didn't make a difference to the outcome. Imagine, for example, the following scenario. Jack and Jill have a common enemy and they have come up with the following plan to kill him: If I wear a hat on a certain morning, Jack will shoot him; otherwise (if I don't wear a hat that morning) Jill will. Imagine that, as it turns out, I decide to wear a hat that morning. So Jack shoots the enemy, and he dies. Was my wearing a hat that morning a cause of his death? Intuitively, no. Intuitively, my wearing a hat wasn't a cause of his death because it made no difference to it. Now, how can this be, if (as we have noted) there are other scenarios—preemption scenarios—where an event causes an outcome even if there is a sense in which it *doesn't* make a difference to the outcome? The answer suggested by Causal DM2 is: the relevant sense of difference-making, according to which my wearing a hat that morning doesn't make a difference to the death (and hence isn't a cause), is that its contribution is

on a par with the contribution that its absence would have made. If we were to say that my wearing a hat caused the death, then we would also have to say that my *not* wearing a hat would have caused the death, if I had not worn a hat. For example, imagine that we said that my wearing the hat caused the death because it caused Jack's shooting, which in turn caused the death. Then, by similar reasoning, we would also have to say that my not wearing the hat would have caused the death, since it would have caused Jill's shooting, which in turn would have caused the death. In other words, the contribution of the event of my wearing the hat vis-à-vis the death doesn't stand out from the contribution of its absence, which is what Causal DM2 requires. By contrast, in a preemption case the contribution of the preempting cause does stand out from the contribution of its absence. In this way Causal DM2 allows us to say that, whereas preemptors are causes, events like my wearing a hat are not causes because they don't make a difference to the ensuing outcomes, in the relevant sense.¹⁵

We have seen that Causal DM2 is a plausible way of capturing the idea that causes are difference-makers without falling prey to some serious problems that cast doubt on Causal DM1. Could there be a similarly plausible *epistemic* difference-making principle that rests on DM2? We turn to this in the next section.

4 The Epistemic Difference-Making Constraint

Before introducing the new principle we need to make a three-way distinction regarding justification. There is, first, the notion of evidential support, which we take to be a mind-independent relation between propositions. We will assume that evidential support takes place against a background of (other) evidence. Thus, that Joe says that it is raining may evidentially support the proposition that it is raining against one background but support the negation of that proposition, or perhaps neither proposition, against some other background (one background may include evidence for Joe's reliability and the other against it, for example).

Second, there is what has been called "propositional justification," which involves not only the fact that some body of evidence supports a proposition but also that there is a subject who *has* that evidence. It is subjects who have propositional justification for believing, and they do in virtue of having evidence which evidentially supports the propositions in question. We will remain neutral on what it takes for a subject to have a piece of evidence. One option is to say that evidence is (or at least can be represented by) propositions. Of course, not any proposition will be part of a subject's evidence. Perhaps if the subject is justified in believing the proposition then it is part of his evidence. Or perhaps a stronger condition is necessary or a weaker one sufficient. Perhaps, for instance, a proposition may be part of a subject's evidence by virtue of the fact that it is the content of an experience that the subject undergoes. An alternative proposal is that a subject's evidence does not consist of propositions that are the content of the beliefs or experiences of the subject, but it consists of those beliefs and experiences themselves. One could, of course, also hold a hybrid view—for instance, one could hold that propositions we are justified in believing as well as experiences are part of our

evidence. In what follows we mostly talk in propositional terms, as we did when presenting the relation of evidential support, but the arguments that we make are neutral in this respect. If the subject has some evidence that supports a proposition *P* against a background of evidence that the subject also has, then he has *prima facie* justification for believing *P*. In the special case where the background evidence includes all the rest of the evidence that the subject has, then the subject has all things considered propositional justification for believing *P*.

A subject may of course not believe what he has propositional justification for believing, or may believe it for the wrong reasons. In either of those cases, the subject fails to have “doxastic justification,” which involves not only a subject’s having propositional justification but also believing something he is propositionally justified in believing on the basis of what so justifies him. If the subject has all things considered propositional justification for believing that *P* and does believe that *P* on the basis of a subset of his evidence which gives him propositional justification for believing *P*, then the subject is doxastically justified in believing that *P*.

The principle that we wish to put forward is a constraint on evidential support. In arguing for it and applying it, however, we will inevitably have to appeal to cases of subjects who are justified or unjustified, or who know or fail to know, and so we will appeal to the notions of propositional and doxastic justification as well.

The constraint is the following:

Epistemic DM2: *E* evidentially supports *P* (given a background of evidence *B*) only if the absence of *E* does not evidentially support *P* (given the same background *B*).¹⁶

If *S* is doxastically justified in believing in *P* based on evidence *E*, then *E* must propositionally justify *P* for *S*, and therefore *E* must evidentially support *P* relative to *S*’s background. It follows from this and Epistemic DM2 that if *S* is justified in believing that *P* based on evidence *E*, then the absence of *E* does not evidentially support *P* relative to *S*’s background. Since knowledge entails justification, it also follows from Epistemic DM2 that if *S* knows that *P* based on *E*, then the absence of *E* does not evidentially support *P* relative to *S*’s background.¹⁷

According to Epistemic DM2, the condition that has to be satisfied for a subject to be justified in believing that *P* on the basis of *E* is, not that the subject wouldn’t have believed *P* in *E*’s absence, but rather that the belief in *P* wouldn’t have been justified *by* *E*’s absence (even if the belief would still have occurred). The general idea is that whatever support exists between a piece of evidence and my belief could not have existed between the absence of that evidence and that very belief. If the relation that exists between the evidence and my belief is on a par with the relation that would have existed between the absence of that evidence and my belief, then the evidence doesn’t make the relevant kind of difference to my belief, and thus the evidence doesn’t support it—it is not *good* evidence for it.

Notice two important differences between Epistemic DM2 and a probabilistic conception of epistemic support. First, as we said above, a probabilistic account has trouble allowing for logical learning, whereas we will argue that Epistemic DM2 *can* explain how we can gain knowledge of necessary propositions on the

basis of contingent ones (see section 5 below). Second, the probabilistic account entails that it is impossible for logically necessary propositions to be evidence for anything. This is the flip side of the impossibility of logical learning. Whereas the impossibility of logical learning follows from the fact that nothing can raise the probability of a logical truth, the impossibility of learning *from* logic follows from the fact that a logical truth cannot raise the probability of anything (the conditional probability of any proposition on a logical truth is the same as the unconditional probability of that proposition). Epistemic DM2, on the other hand, leaves it open that logical truths can epistemically support propositions—although, of course, it doesn't explain *how* that may happen.

In the next section we argue for Epistemic DM2.

5 In Defense of Epistemic DM2

First of all, notice that Epistemic DM2 succeeds where Epistemic DM1 fails (in the same way that Causal DM2 succeeds where Causal DM1 fails). Consider, again, the Frankfurtian version of the ice-thawing case from section 2 above (which, as we saw, Epistemic DM1 incorrectly rules out as a case of knowledge). The constraint specified by Epistemic DM2 is met in that case. Let E be the propositions on which I base my belief in the actual case: for simplicity, let it be the proposition that I left the ice outside in hot weather. Had I not left the ice outside in hot weather, the neuroscientist would have made sure that I still believe that I did, and as a result I would have still believed that the ice melted, but the proposition that I didn't leave the ice outside doesn't support the proposition that the ice melted. In this case, the proposition wouldn't have justified me simply because it wouldn't even have been part of my evidence (this is in effect what the neuroscientist is doing: preventing that proposition from being part of my evidence). But, of course, even if that proposition had been part of my evidence, it still would not have supported the proposition that the ice melted. Therefore, Epistemic DM2 doesn't deny that the subject has knowledge in that case. Of course, it doesn't entail that the subject has knowledge, but that is because Epistemic DM2 provides only a necessary condition on evidential support—and so, via the evidential support condition on justification and the justification condition on knowledge, only a necessary condition on knowledge.

Not only does Epistemic DM2 deal correctly with that counterexample; suitably modified, the example can also be used to further illustrate its plausibility. Imagine that I find out that scientists have discovered a special kind of ice that melts only under moderate temperatures (not hot temperatures). In that case, had I *not* taken the ice outside, that very fact would have supported the belief that the ice has melted. But, clearly, the proposition that I left the ice outside doesn't support the proposition that the ice has melted *against that background*. Again, we see the asymmetry between what constitutes the evidence in each case and its absence.

A few points of clarification are in order. First, note that Epistemic DM2 is consistent with the claim that a belief can be *justified* for a subject both in the presence

and in the absence of a particular piece of (good) evidence for the proposition. What Epistemic DM2 says is that a piece of evidence and its absence cannot both equally *support* the belief.¹⁸ Relatedly, Epistemic DM2 is consistent with the claim that a subject S could have been justified in believing that P on the basis of *different* evidence (other than E). Epistemic DM2 says that E doesn't support P unless not-E, the negation of that *very same* proposition, doesn't support P. Crucially, we must distinguish between E's negation and some other proposition E* that (in the circumstances) might guarantee that E's negation obtains. Suppose, for example, that there are bells whose sound is so powerful that they temporarily blind whoever is in front of them when they are tolling. A subject can know that those bells are in front of him based on his seeing them and he can also know the same fact on the basis of his hearing them, even if his hearing them entails his not seeing them. Epistemic DM2 is compatible with this fact. When seeing the bells supports the proposition that the bells are there, not seeing them doesn't. Hearing them might have, but not seeing them is different from hearing them, although the latter guarantees the former. This mirrors what we have identified as an important feature of Causal DM2: Causal DM2 is compatible with the claim that, in the absence of c, e would have been caused by something other than c (and different from c's absence), say, d. As we have seen, it is for this reason that Causal DM2 can accommodate the possibility of preemption.

Note, also, that a key difference between Epistemic DM2 and Epistemic DM1 is that Epistemic DM2 is a non-reductive condition: the overtly epistemic notion of *support* is appealed to in the formulation of the condition. In section 3 we noted a similar difference between Causal DM2 and Causal DM1. As it was the case with causation, we think that this is why Epistemic DM2 is able to deal correctly with the counterexamples that arise for DM1-based views. But perhaps Epistemic DM2 has counterexamples of its own? In what remains of this section we lay out what seem to us to be the best candidates for counterexamples, and explain why they are not counterexamples after all.

We have already hinted at one possible kind of potential counterexample: the bells example. To see that this kind of case doesn't threaten Epistemic DM2 it is crucial to distinguish scenarios where the subject bases his belief that P on E in the actual case and on not-E in the counterfactual case from scenarios where the subject bases his belief that P on E in the actual case and on E* (where E* is different from not-E, although E* might guarantee not-E) in the counterfactual scenario. Epistemic DM2 rules out cases of the first kind only. The bells case (as well as all the others that can be built on the same template) is of the second kind.

A second, and trickier, kind of potential counterexample also involves arriving at the same belief through two different routes, but the two routes are such that one of them involves a proposition and the other one involves the negation of that very proposition. Imagine a variation on the hat case discussed above, in section 3. Imagine, this time, that I have arranged things so that if I wear a hat this morning then Jack will kill an enemy of mine, whereas if I don't wear a hat Jill will perform the deed. I wear a hat and some time later come to justifiably believe that my enemy is dead. Now, had I not worn a hat, I would have come to the same belief, and with

the same degree of justification. It seems, therefore, that the fact that I wore a hat supports the belief that my enemy is dead, but also that, had I not worn a hat, my not wearing a hat would have supported the same belief, and thus it seems that this case is a counterexample to Epistemic DM2.

In reply, it is crucial to distinguish the relation of evidential support between propositions (which is what Epistemic DM2 is about) from the relation of justification that accrues to a subject's beliefs. At the level of evidential support, neither the proposition that I am wearing a hat nor its negation support the proposition that my enemy will be killed, because they don't add anything to the support that that proposition already receives from the background propositions (which entail that my enemy will be killed whether or not I wear a hat). Now, the proposition that I am wearing a hat does support the proposition that Jack will kill my enemy, and it may be the case that the proposition that Jack will kill my enemy supports the proposition that my enemy will die to a greater degree than it was already supported by the background. If so, then this case illustrates the fact that evidential support is not transitive. But this is as it should be. For example, that a die has come up on 1 supports the belief that it came up either on 1 or on 2, and that it came up either on 1 or on 2 supports the belief that it came up on 2, but that it came up on 1 definitely doesn't support the belief that it came up on 2—on the contrary, it supports the opposite belief. And it's not just that the epistemic support relation is not transitive, but the kind of case with which we are dealing here is precisely one for which we would expect transitivity to fail. Again, if I know that my wearing a hat and my not wearing a hat would both lead, through different routes, to the same outcome, then the fact that I am wearing a hat doesn't add anything to whatever epistemic support the belief that my enemy will die already has—it only changes my beliefs as to how that death will come about. There is another parallel with causation here. As we saw in section 3, Causal DM2 entails that my wearing the hat that morning doesn't cause my enemy's death. This is so even if it causes something that causes the death. In other words, this scenario illustrates how causation can fail to be transitive.¹⁹

There is an alternative reading of the case according to which it is not an example of transitivity failure. According to that reading, the proposition that Jack will kill my enemy does not support the proposition that my enemy will die, for that proposition is already maximally supported by the background of the case. We grant that the background *entails* that my enemy will die, but we want to leave it open that propositions entailed by the evidence may gain further evidential support. Suppose, however, that we are wrong, and propositions entailed by the evidence cannot gain evidential support. Under this description of the case we again fail to have a counterexample to Epistemic DM2, for there is no proposition such that it and its negation both evidentially support another proposition. Thus, even though there is a reading of the case according to which there is no failure of transitivity, that alternative reading still provides no challenge to Epistemic DM2.

Finally, notice that, at the level of justification, if my belief that Jack will kill my enemy is based exclusively on my belief that I am wearing a hat (and not, for example, on my seeing Jack firing), then that Jack will kill my enemy is not part

of my evidence, and so even if it supported the proposition that my enemy will die, it could not *justify me* in believing that proposition. What justifies me, in this case, is my background evidence. We conclude, then, that the hat case is not a counterexample to Epistemic DM2.

A final kind of potential counterexample has to do with necessary truths. Take the proposition that *P or not-P* (where P is a contingent truth). Given the logical rule of addition (X or Y logically follows from X), doesn't P epistemically support that proposition? And doesn't not-P also support it in exactly the same way? If so, this is a simple counterexample to Epistemic DM2.

What is undeniable is that both P and not-P entail *P or not-P*. But the evidential support relation doesn't always mirror the entailment relation. For instance, to use an example by Dretkse (1970) (although for a different purpose than the one he put it to), that *that* animal in front of me is a zebra entails that it isn't a mule cleverly disguised to look like a zebra. But, against normal backgrounds, that it is a zebra will not support that it is not a cleverly disguised mule (if anything, in this case, the support relation goes the other way). Hence, our view is that, although necessary truths are *entailed* by any proposition whatsoever (and so, in particular, *P or not-P* is entailed by P and by its negation not-P), they are not *supported* by any proposition whatsoever.

To clarify, we are not denying that contingent truths can evidentially support necessary truths. Testimony is clearly one way to achieve this: that a mathematician tells me that P is a theorem supports P. This is no threat to Epistemic DM2, because her not telling me that it is a theorem doesn't similarly support P. But no one would accept that *any* contingent proposition supports *any* necessary proposition. Thus, everyone thinks that there must be constraints on which contingent propositions support which necessary propositions. We think that Epistemic DM2 provides one such constraint. This is an important advantage that Epistemic DM2 has over probabilistic accounts of epistemic support, which cannot easily account for contingent propositions supporting necessary ones.

This concludes our defense of Epistemic DM2. In the next section we show how the principle can throw light on a thorny issue in contemporary epistemology.

6 Easy Knowledge

As we have seen, the fact that Epistemic DM2 gives a non-reductive condition on evidential support (and thus on knowledge and justification) is in great part responsible for its plausibility. But, of course, a non-reductive principle cannot be put to the same use as a reductive one—in particular, Epistemic DM2 cannot be used in the traditional project of defining propositional knowledge in wholly non-epistemic terms. This doesn't mean that it cannot have interesting applications, however. In this section we show that Epistemic DM2 helps to shed light on the problem of easy knowledge.

The problem of easy knowledge has been forcefully presented by Stewart Cohen in Cohen (2002).²⁰ Adopting the terminology of this paper, Cohen's problem can be presented as follows. Suppose that an epistemological theory allows that a subject

can know a proposition on the basis of certain evidence without knowing that the evidence and the proposition co-vary (that is, without knowing that either the evidence obtains and the proposition is true or the evidence doesn't obtain and the proposition is false). More generally, suppose that the theory allows for the mere exercise of a faculty or competence to give rise to knowledge, even in the absence of independent evidence for the reliability of that faculty. Suppose, for instance, that the theory allows that a subject can know the colors of objects on the basis of how they look to her without knowing that objects have the colors they visually seem to have—let us call this the “matching” proposition. Any such theory, Cohen argues, will be prey to the easy knowledge problem. For take a subject, Roxanne, who doesn't know whether the matching proposition is true. An epistemological theory of the kind we are envisaging would allow Roxanne to easily amass evidence for the matching proposition. For Roxanne could engage in any number of instances of the following argument:

1. The object in front of me looks *x*.

Therefore,

2. The object in front of me is *x*. (From 1)

Therefore,

3. My color vision didn't deceive me this time. (From 1 and 2)

Suppose that an object looks red to Roxanne. Then an instance of 1 is true for her. According to the kind of epistemological theory that we are considering, the corresponding instance of 2 will then be justified for Roxanne—and, if everything goes right, she will be able to know the instance of 2 on the basis of the instance of 1.²¹ The movement from 1 to 2 is not, therefore, *an inference* that Roxanne engages in, for she need not even believe 1 for the transition to be justified for her. Rather, the very fact that 1 is true justifies Roxanne in believing 2. But even though Roxanne *need not* believe 1 to be justified in believing 2 on its basis, she may of course find out that 1 is true, and thus believe it. Indeed, it should take only a little bit of reflection for Roxanne to figure out (and know) that 1 is true. Thus, Roxanne may well know both 1 and 2 (the latter on the basis of the former). Moreover, she can easily see that 3 follows from 1 and 2. Therefore, by closure, Roxanne should also be able to know 3.²² It doesn't follow from closure alone that Roxanne knows 3 *on the basis* of 1 and 2. But this follows from the construction of the case, since we assumed that Roxanne *didn't* know 3 before knowing 1. Therefore, if she comes to know 3 *after* knowing 1, that must be because she knows 3 on the basis of 1 (and 2). The alternative is that knowledge of 3 materializes *ex-nihilo every time* a subject knows 1 (and 2) and notices that they imply 3. Theories that allow that subjects can know instances of 1 without knowing instances of 3, then, are committed to saying that subjects can know instances of 3 on the basis of instances of 1 (and 2). Most epistemologists agree that this consequence is unpalatable. Of course, every non-skeptic should agree that we do know many instances of 3, but almost no one is comfortable allowing that we know it on the basis of 1 alone (which gives us 2).

It is crucial to notice that there would be nothing wrong with an argument of that form *if 2 had independent support*. Suppose, for instance, that Roxanne is wondering about her eyesight, and so sets up an experiment where different colors flash on a computer screen. She takes notice of what color the screen looks like to her and then asks a friend with reliable color vision what color the screen is. In *this* way, Roxanne can of course amass lots of legitimate evidence for the reliability of her own color vision. This is *hard* knowledge of one's own reliability. The *easy* (too easy for comfort) knowledge is obtained when 2 receives all of its epistemic support from 1 itself.

Epistemic DM2 can shed light on the problem of easy knowledge because it can explain *what exactly is wrong* with cases of easy knowledge, or with a theory that allows that we can come to know 3 just on the basis of 1. Here is why. Suppose that the object in front of Roxanne doesn't look red to her. It seems that she could now mount an argument for the non-deceiving nature of her color perception that has approximately the same plausibility as the previous argument:

1*. The object in front of me doesn't look red.

Therefore,

2*. The object in front of me is not red. (From 1*)

Therefore,

3. My color vision didn't deceive me this time. (From 1* and 2*)

Here 1* is (the relevant instance of) 1's negation, and 2* is (the relevant instance of) 2's negation. And here is the crucial part of our argument: any theory that says that 1 supports 2 all by itself will also have to say that 1* supports 2* all by itself. For, if the proposition that the object looks red (in normal perceptual conditions) supports the proposition that the object is red, then the proposition that the object doesn't look red (in similarly normal perceptual conditions) also supports the proposition that the object is not red. We are taking "normal perceptual conditions" to include the fact that Roxanne is looking at the object in front of her with her well-functioning eyes open, in a well-lit environment, in the absence of factors that would make the object look a different color than it is, etc. Perhaps any such theory would *also* have to say that the proposition that the object looks blue supports the proposition that it is not red. But this doesn't contradict the claim that the same proposition is also supported by the proposition that the object doesn't look red. Perhaps both its not looking red and its looking blue support the proposition that it is not red.

Now, some philosophers might find a difference between the support that 1 confers on 2 and the support that 1* confers on 2*.²³ According to those philosophers, 1 supports 2 *truly by itself, even against an empty background*, whereas 1* supports 2* only against the background that I am looking at the object. (After all, the object in front of me might not look red to me because my eyes are closed.) As a result, the support that 1 confers on 2 may be *defeated*; in contrast, 1* doesn't even lend *prima facie* support to 2* unless some further conditions obtain.

The view that there is this asymmetry must face some serious difficulties, but we don't want to pursue them here. Even granting that such asymmetry exists, it is still true that theories that allow for easy knowledge will have to grant that 1* supports 2* "by itself" in the sense of "without independent evidence for the relevant matching proposition." As we have pointed out, the views for which the easy knowledge problem arises are views according to which the mere exercise of a faculty (perception) can justify a belief even in the absence of prior or independent evidence for the reliability of that faculty. The thought that these views are trying to capture is that you don't need to have any independent evidence that your faculties are reliable to conclude that the world is like the faculties you are exercising suggest it is. Thus, if something looks red to you when you're looking at it, you are justified in believing it's red even if you don't have any prior or independent evidence for the reliability of your eyesight. But, similarly, then, if something doesn't look red when you're looking at it, it seems that you would be justified in believing that it's not red even if you don't have any prior or independent evidence for the reliability of your sight. In this case, you have to be looking at the object to be justified in believing it's not red when it doesn't look red, but all this means is that you have to be exercising your faculties of perception. Again, the view is that, without prior evidence of the reliability of your faculties, you can be justified in believing the world is like your faculties indicate it is on an occasion where your faculties are indeed being exercised (not when you are not exercising them). Thus, views according to which the object's looking red support its being red will have to grant that, had the object not looked red *in the same circumstances* (that is, while still exercising my perceptual faculties), that in itself would have supported the proposition that the object isn't red.

Now, given that any theory that says that 1 supports 2 also has to say that 1* supports 2*, any theory that says that 1 supports 3 on the basis that it supports 2 will equally have to say that 1* (the negation of 1) supports 3, on the basis that it supports 2*. But this is precisely what Epistemic DM2 prohibits. According to Epistemic DM2, 1 cannot support 3 when the negation of 1 supports 3. For this clashes with the idea that in order for our evidence to support what we believe it has to make a difference to what we believe. In particular, it clashes with the idea that facts about how things appear to us cannot support beliefs about the reliability of our faculties unless the absence of those facts wouldn't similarly have supported those beliefs.

It is important to note that Epistemic DM2 allows cases of *hard* knowledge to go through, as expected. For in a case of hard knowledge, 3 is supported not on the basis of 1 alone, but on the basis of 1 and some independent support for 2 (say, the friend's testimony that the object is red). In that case, the negation of 1 doesn't support 3. Given the friend's testimony that the object is red, that the object doesn't look red to Roxanne doesn't support the proposition that the object isn't red. A crucial difference between easy and hard cases of knowledge is that in easy cases 2 is supported only by 1 itself, whereas in hard cases it is supported independently. In easy cases, if 1 were false then 2 wouldn't be part of Roxanne's evidence, even if true—for a fact like 2 can only be part of a subject's evidence by being the content of a (justified) belief the subject has. In hard cases, if 1 were false 2 would still

be part of Roxanne's basis, for belief in 2 would still be supported by her friend's testimony.

Cases of easy knowledge are similar in some respects to the hat case that we examined above, but they are also crucially different. In both cases, there is a fact (1, or the subject's wearing a hat) that fails to support a conclusion (3, or that the subject's enemy will die), even though it supports *another* belief (2, or that Jack will kill the subject's enemy) that in turn supports the conclusion (in the case of 2, together with 1). We said, however, that the subject may well know in the hat case, whereas there is no easy knowledge. How can this be? The crucial difference is that in the hat case part of the subject's background evidence is that if he wears a hat Jack will kill his enemy, whereas if he doesn't Jill will. That is, it is part of the subject's background evidence that, regardless of whether or not he wears a hat, his enemy will die. The hat case is a case of hard knowledge. It is *not* part of the subject's background evidence in a case of easy knowledge that, regardless of whether or not 1 is true, 3 is true.²⁴

We have seen that Epistemic DM2 disallows cases of easy knowledge (while allowing for cases of hard knowledge). It disallows cases of easy knowledge on the basis of the fact that the evidence doesn't make the right kind of difference to the belief in those cases. Now, is this the best explanation of what goes wrong in these cases? In the remainder of this section we consider two existing explanations in the literature on the problem and we argue that Epistemic DM2 is better than both of them.

We start with a proposal by Michael Titelbaum.²⁵ According to Titelbaum, allowing for easy knowledge is a problem because allowing for easy knowledge is allowing for the legitimacy of "no-lose investigations." Titelbaum defines a no-lose investigation as an investigation that satisfies the following conditions:

- 1) The subject is not justified in believing P when the investigation begins.
- 2) The subject knows when the investigation begins that he will not be justified in believing not-P when the investigation ends.
- 3) The subject knows when the investigation begins that if P is true then he will be justified in believing P when the investigation ends.

We want to make three related points about Titelbaum's alternative explanation of what goes wrong with cases of easy knowledge. First, our explanation is much simpler (according to natural, if perhaps hard to capture precisely, standards of simplicity). According to us, Roxanne doesn't know the matching proposition because the reasons for thinking that her evidence supports that proposition are also reasons for thinking that the negation of that evidence supports the same proposition. According to Titelbaum, Roxanne doesn't know the matching proposition because she is not justified in believing it when her investigation begins, she knows at the beginning of the investigation that she will not be justified in believing the negation of the matching proposition at the end of the investigation, and she knows at the beginning of the investigation that if the matching proposition is true then she will be justified in believing it at the end of the investigation.

Second, Titelbaum's explanation is not just less simple than ours: it seems to have the wrong shape to capture what is wrong about cases of easy knowledge. Notice that two of Titelbaum's conditions contain knowledge operators. At a first glance, what goes wrong in a case of easy knowledge has nothing to do with that the subject knows about how an investigation will proceed. Of course, that first glance can be corrected by theorizing. If the best available explanation of what goes wrong in cases of easy knowledge posits such conditions with knowledge operators, then that is a reason to believe that the fact that those conditions are violated captures what goes wrong in cases of easy knowledge. But there is an explanation available that posits no such conditions: our own. From the vantage point of Epistemic DM2, Titelbaum's explanation seems to contain superfluous appeals to knowledge.

Finally, there are cases that Titelbaum argues would count as no-lose investigations according to some epistemic theories but are not ruled out by Epistemic DM2. One might think that this is a point in favor of Titelbaum's approach. On the contrary, we think that it is a point against it. Here is the description of one such case, by Titelbaum himself:

The Court Jester: Noblemen from Italy have arrived in the King of England's court, bringing with them the jester Giacomo. The King has heard a rumor that Giacomo is quite the ladies' man, but the King knows the rumor's source is unreliable and so lacks justification to believe it.

To settle the matter, the King orders the jester to regale the court with tales of his amorous conquests. The King's instructions are very precise: If Giacomo is indeed a ladies' man, the tales are to be true; if not, the jester is to make up false tales that sound convincingly real. (...)

As the King expects, the jester spends a long evening describing broken hearts left littering the landscape. In fact, Giacomo is a ladies' man and all his tales are true. At the end of the evening, is the King justified in believing this?²⁶

Clearly, the answer to Titelbaum's (purely rhetorical) question is that the King isn't justified. More precisely, it is obvious that Giacomo's testimony doesn't justify the King in believing that the jester is a ladies' man. Titelbaum makes two claims about the case. First, there are theories according to which, after hearing Giacomo's tales, the King is justified in believing that Giacomo is a ladies' man. Second, such theories are wrong because they entail that the King's investigation amounts to a no-lose investigation. Notice, by contrast, that Epistemic DM2 doesn't have anything to say about the case. Epistemic DM2 doesn't explain why the testimony fails to justify the King's belief. Although the testimony fails to justify the King's belief, this *isn't* because any reasons for thinking that the jester's testimony supports the proposition that he is a ladies' man would also be reasons for thinking that the absence of the testimony supports the same proposition.

We believe that this is the right result, however. Why does Giacomo's testimony not justify the King in believing that Giacomo is a ladies' man? The obvious answer is: because the King has a defeater for that testimony—namely, the fact that Giacomo's testimony has been forced by the King. The King knows that Giacomo

will testify to his amorous adventurous whether or not those tales are true. This has nothing to do with the reason why Roxanne is not justified in believing the matching proposition—she doesn't have defeaters. And, as Titelbaum himself notes, paradigmatic reliabilists such as Goldman (1976) add a no-defeaters condition to their theories of epistemic justification. If the King has a defeater, then those theories won't entail that the King is justified. Therefore, the fact that Epistemic DM2 is silent about the court jester's case is not an embarrassment for the condition or a reason to think that Titelbaum's account is a better account. In fact, it is a reason to think that Epistemic DM2 is right on target (and Titelbaum's approach is not). To be clear, we are not saying that anything in Titelbaum's paper is false. Titelbaum might well be right that any theory that allows for no-lose investigations is incorrect. What we are saying is that our account fares better with respect to the goal of identifying the particular epistemic failure in cases of easy knowledge—but this may well not be a goal that Titelbaum set for himself in the paper.

Jonathan Weisberg has suggested that what goes wrong in cases of reasoning such as Roxanne's is that they violate what he calls a "no-feedback" principle.²⁷ The details of the principle won't matter—applied to Roxanne's case, the idea is that 1 doesn't support 3 *by itself*; it only supports 3 with the help of 2, which is itself derived from 1. Given that 2 is derived from 1, whatever support 2 lends to 3 would have to be a support that 1 *already* lends to 3. So 2 cannot legitimately be used to derive 3, unless 1 could be used to derive 3 by itself. Without committing himself to a probabilistic view of justification, Weisberg points out that there is a probabilistic result behind his no-feedback principle: probabilistic support is not transitive—even though 1 raises the probability of 2 and 1 and 2 together raise the probability of 3, 1 doesn't by itself raise the probability of 3.

Now, note that this argument doesn't really explain *why* 1 doesn't support 3. (The probabilistic result associated with the no-feedback principle would, if one assumed a probabilistic view of justification; but, as we have pointed out, this is not Weisberg's suggestion.) Weisberg's proposal is, in a nutshell, that 2 cannot be used to support 3 *unless* 1 already supports 3 by itself. We, together with Weisberg, think that 1 *doesn't* support 3 by itself, but the question is *why* we think this. (More specifically: Why doesn't 1 support 3, if 3 can be derived from 1 in conjunction with other things that can be derived from 1?) Weisberg's no-feedback principle doesn't provide an answer to these questions. But Epistemic DM2 does. The answer is, again, that if we said that 1 supports 3, then we would have to say that the negation of 1 supports 3 as well, but this clashes with the idea that in order for our evidence to support what we believe it has to make a difference to what we believe.

We conclude that Epistemic DM2 provides a better explanation of what goes wrong in cases of easy knowledge than both Titelbaum's and Weisberg's. The thought that the evidence must make a difference to the belief it supports is a very simple, intuitive thought. And Epistemic DM2 is a principle that captures that simple thought in simple and elegant terms, in accordance with an independently motivated, general conception of difference-making (DM2). Hence the fact that the difference-making constraint is violated in cases of easy knowledge arguably

constitutes the best explanation of why knowledge is absent in those cases. Knowledge is *hard*, among other things, in that it requires difference-making.

Notes

*For helpful comments and suggestions on earlier drafts of this paper we are grateful to a referee for *Nous*, audiences at the University of Buenos Aires and the National Autonomous University of Mexico, Stewart Cohen, and, especially, Michael Titelbaum.

¹ Almost no one currently defends this simple claim about causation (although see Coady (2004) for an exception). See Lewis (1986) for a more sophisticated counterfactual account.

² For discussion of the main proposals and challenges, see Collins, Hall, and Paul (2004b).

³ A more sophisticated statement of sensitivity also includes a same-basis condition: If P were false, then S wouldn't believe that P on the same basis on which he actually does. See Dretske (1970) and Nozick (1981).

⁴ The kind of counterexample that follows is taken from Vogel (1987).

⁵ Why not, more simply: The subject knows that P on the basis of Q only if, had Q not obtained, the subject wouldn't have believed that P? Because this is obviously false: a subject could know that P on the basis of Q although, had Q not obtained, he would have known that P on some other basis—see, for instance, the tolling bells example in section 5.

⁶ There is an obvious ambiguity in saying that a subject bases his belief that P on Q—does the subject base his belief on the fact itself or on his belief that the fact obtains? Depending on how we answer this question we might have different opinions on which disambiguation of the principle is better. As we point out later in the text, we will remain neutral on the question about what can constitute a basis for a belief.

⁷ Sosa's safety condition can be seen as a variant on the same theme—see Sosa (1996) and Sosa (1999). Counterexamples to safety also abound—see, for instance, Comesaña (2005). Sosa (2010) arguably abandons safety—see Comesaña (forthcoming).

⁸ See Sartorio (2005).

⁹ Note that the interesting cases will be those where the constraint is not vacuously satisfied. For instance, if causation by absences is not possible and if the first relatum in a causal relation is an event, then the absence of the first relatum would not have borne the causal relation to the second relatum simply because it's an absence. In that case the constraint will be trivially met. But it will be a substantial constraint otherwise.

¹⁰ As can be seen in the total probability theorem, according to which $\Pr(A) = \Pr(A|B)\Pr(B) + \Pr(A|\text{not-}B)\Pr(\text{not-}B)$, $\Pr(A)$ is a weighted average of $\Pr(A|B)$ and $\Pr(A|\text{not-}B)$, from where the result in the text follows.

¹¹ For a survey of probabilistic accounts of causation (and their problems), see Hitchcock (2011). Bayesian accounts of evidential support (which are legion) are of course probabilistic accounts of evidential support.

¹² Which is not to say that there aren't attempts to model logical learning within a probabilistic framework. One of the better known attempts is that of Garber (1983)—but Garber's approach allows for logical learning while *still* having the consequence that subjects are logically omniscient with respect to the tautologies determined by the chosen language.

¹³ Dutch-book arguments seem to confuse practical with epistemic rationality—but cf. Christensen (2004). More recent "epistemic utility" arguments (see, for instance, Joyce (1998)), although designed to get around this problem, arguably fall prey to it nevertheless—but this is not the place to argue this point.

¹⁴ For a defense of this view, see Sartorio (2005). As anticipated in n. 9 above, for this view to have substance one must take seriously the possibility of causation by absences—otherwise, Causal DM2 will be true, but only vacuously so. Also, as explained in Sartorio (2005), for the view to be plausible it needs to rely on a sufficiently fine-grained view of absences (for example, a view according to which an absence obtains whenever no event of a certain type obtains, where the type in question can be quite specific).

¹⁵ Note that embracing Causal DM2 naturally results in the rejection of the transitivity of causation (for further discussion see Sartorio (2005)). As we point out later, the epistemic difference-making principle based on DM2 that we will propose may also result in a failure of transitivity of the relation of evidential support.

¹⁶ Notice that, as anticipated, we are leaving open what kind of entity the first relatum of Epistemic DM2 is. Notice also that, unlike DM2 and Causal DM2, Epistemic DM2 is not formulated in counterfactual terms. This is because whether a proposition evidentially supports another is independent of whether those propositions are true, whereas an event can cause another event only if they both obtain. The counterfactual is there when applied to doxastic justification, though, because we are then worried about whether the subject *would* have been justified had he based his belief on not-E—and our principle entails that the answer is “No” if E evidentially supports P.

¹⁷ How about cases of “ungrounded” knowledge or justification—cases where a proposition is known or justified but not in virtue of being evidentially supported by anything else? Some philosophers would argue that cogito-like propositions (such as *I think*, or *I exist*) are like this. We are not convinced that these are genuine examples of ungrounded knowledge/justification—one could hold, for instance, that they are grounded in the fact that our own thought (or our own existence) is revealed to us in introspection, just as our knowledge of the color of a ball may be grounded in the fact that the ball looks to be of that color to us. But we will not argue here that all knowledge/justification is grounded. Instead, we choose to present our epistemic difference-making principle as a principle governing evidential support.

¹⁸ In this connection, see our discussion of the hat case below in this section.

¹⁹ See n. 15 above.

²⁰ See also Vogel (2000).

²¹ In what follows we sometimes talk about 1 and 2 (and later 1* and 2*) when we mean their instances. We trust that the context will make the disambiguation obvious.

²² We cannot here comment on the many issues that arise about the proper formulation of a closure principle for knowledge. But we do want to point out that, even if it were *impossible* to formulate a correct closure principle, that wouldn't mean that subjects can never expand their knowledge by deduction. Roxanne's case is a paradigmatic case of this unproblematic kind of expansion, and so we put little faith in denying closure as a way of getting out of the easy knowledge problem.

²³ We suspect that Pryor (2000) and Pollock and Cruz (1999) are three such philosophers.

²⁴ Although, as White (2006) has argued, this is something that the subject could know *a priori*, if it were true that 1 alone justifies 3. This result is problematic for theories according to which 1 alone can justify 3 for two reasons: it is an implausible result, and it signals a certain tension within those theories (knowledge of 3 is not needed for 1 to justify 2, but it will nevertheless always be available).

²⁵ Titelbaum (2010).

²⁶ Titelbaum (2010), p. 125.

²⁷ Weisberg (2010).

References

- Christensen, D. (2009), *Putting Logic in Its Place* (Oxford University Press).
- Coady, D. (2004), “Preemptive Preemption,” in *Collins, Hall and Paul* (2004a), pp. 325–39.
- Cohen, S. (2002), “Basic Knowledge and the Problem of Easy Knowledge,” *Philosophy and Phenomenological Research*, LXV, 2, pp. 309–29.
- Collins, J., N. Hall, and L. A. Paul (eds.) (2004a), *Causation and Counterfactuals* (MIT Press).
- Collins, J., N. Hall, and L. A. Paul (2004b), “Counterfactuals and Causation: History, Problems, and Prospects,” in *Collins, Hall, and Paul* (2004a), pp. 1–57.
- Comesaña, J. (2005), “Unsafe Knowledge,” *Synthese* 146, pp. 395–404.
- Comesaña, J. (forthcoming), “Sosa on Safety and Epistemic Frankfurt Cases,” forthcoming in John Turri (ed.), *Virtuous Thoughts* (Springer).
- Dretske, F. (1970), “Epistemic Operators,” *Journal of Philosophy*, 67, pp. 1007–23.
- Frankfurt, H. (1969), “Alternate Possibilities and Moral Responsibility,” *Journal of Philosophy*, 66, 23, pp. 829–39.

- Garber, D. (1983), "Old Evidence and Logical Omniscience in Bayesian Confirmation Theory", in J. Earman (ed.), *Testing Scientific Theories, Midwest Studies in the Philosophy of Science*, Vol. X, pp. 99–131.
- Goldman, A. (1976), "What is Justified Belief?," in G. S. Pappas (ed.), *Justification and Knowledge* (D. Reidel), pp. 1–23.
- Hitchcock, C. (2011), "Probabilistic Causation," *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/causation-probabilistic/>.
- Joyce, J. (1998), "A Nonpragmatic Vindication of Probabilism," *Philosophy of Science* 65, pp. 575–603.
- Lewis, D. (1986), "Causation," in *Philosophical Papers II* (Oxford University Press), pp. 159–213.
- Nozick, R. (1981), *Philosophical Explanations* (Harvard University Press).
- Pollock, J. and J. Cruz (1999), *Contemporary Theories of Knowledge* (Rowman and Littlefield).
- Pryor, J. (2000), "The Skeptic and the Dogmatist," *Noûs* 34, pp. 517–49.
- Sartorio, C. (2005), "Causes as Difference-Makers," *Philosophical Studies* 123, 1–2, pp. 71–96.
- Sosa, E. (1996), "Postscript to 'Proper Functionalism and Virtue Epistemology'," in J. Kvanvig (ed.), *Warrant in Contemporary Epistemology* (Rowman and Littlefield), pp. 271–81.
- Sosa, E. (1999), "How Must Knowledge be Modally Related to What Is Known?," *Philosophical Topics*, 26, pp. 373–84.
- Sosa, E. (2010), *A Virtue Epistemology* (Oxford University Press).
- Titelbaum, M. (2010), "Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology," *Philosophical Studies* 149, pp. 119–34.
- Vogel, J. (1987), "Tracking, Closure, and Inductive Knowledge," in Luper-Foy (ed.), *The Possibility of Knowledge* (Rowman and Littlefield), pp. 197–215.
- Vogel, J. (2000), "Reliabilism Leveled," *Journal of Philosophy*, 97, 11, pp. 602–23.
- Weisberg, J. (2010), "Bootstrapping in General," *Philosophy and Phenomenological Research*, LXXXI, 3, pp. 525–48.
- White, R. (2006), "Problems for Dogmatism," *Philosophical Studies*, 131, 3, pp. 525–57.