Philos Stud (2022) 179:1205–1220
https://doi.org/10.1007/s11098-021-01703-3

How to modify the strength of a reason

Andrew Kernohan[1]

Accepted: 6 July 2021 / Published online: 13 July 2021

Abstract Kearns and Star have previously recommended that we measure the degree to which a reason supports a conclusion, either about how to act or what to believe, as the conditional probability of the conclusion given the reason. I show how to properly formulate this recommendation to allow for dependencies and conditional dependencies among the considerations being aggregated. This formulation allows us to account for how considerations, which do not themselves favour a specific conclusion, can modify the strength of a reason for that conclusion, and thus to explain the intensifiers and attenuators described by Dancy. The formulation also accounts for the workings of partial undercutters in epistemology. I then show how my account avoids the counterexamples that Brunero levied against probability-based theories of the strengths of reasons. My account supports the theory, suggested by Kearns and Star, that the strengths of reasons are measured by conditional probabilities. If my account is successful, then it will count in favour of the idea that the strengths of reasons are measured on the same scale as are conditional probabilities.

Keywords Metaethics · Reasons · Probability · Modifiers · Attenuators · Intensifiers · Evidence · Undercutters · Jonathan Dancy · Stephen Kearns · Daniel Star · John Brunero

✉ Andrew Kernohan
kernohan@dal.ca

1 Department of Philosophy, Dalhousie University, Halifax, NS B3H 4R2, Canada

1 Introduction

Suppose that we use the conditional probability of a conclusion, given a reason, to measure the degree to which the reason supports the conclusion. I show that this supposition accounts for how a consideration, which does not directly favor a conclusion about how to act, can modify the strength of a reason for that conclusion in the way that Jonathan Dancy described intensifiers and attenuators as working (2004). This account allows the probability-raising theory of the strength of reasons put forward by Stephen Kearns and Daniel Star (2008, 2009) to evade the counterexamples that John Brunero (2009, 2018) levied against it. These two successes count in favor of the original supposition that we should use conditional probabilities to measure the strengths of reasons.

2 Epistemic probabilities and the strengths of reasons

Who would entertain the supposition that the strengths of reasons are measured on a probability scale? Not Timothy Scanlon, for one. He thinks that reasons have only an ordinal measure. “The strength of a reason is an essentially comparative notion, understood only in relation to other particular reasons.” (2014, p. 111) Reasons are sufficient if they make an action rational, and reasons are conclusive if they make an action required. He uses the notion of a sufficient reason to order reasons in terms of one being stronger than another (2014, p. 108).

Not John Broome, for another. He thinks that reasons have a measurement scale analogous to physical weights (2013, pp 52–53). Combining the strengths of reasons is like putting reasons on a metaphorical balance scale and acting according to the weightier pan of reasons. According to Broome’s metaphor, and contrary to Scanlon’s view, there is no such thing as a conclusive reason since there is no largest weight for a reason to have. Unlike probabilities, physical weights have no upper bound. The physical weight metaphor suggests the fallacious view that that the strength of a reason is something that is constant in all contexts (Dancy, 2004). Broome notes that the view that there is something like the weight of a reason that is constant across all contexts is an idealization because “the weights of reasons will often be influenced by their context, and in particular by so-called ‘organic’ interactions between different reasons.” (2013, p 52)

Possibly W. D. Ross, were he alive today. Ross once suggested that moral reasoning is fallible and hinted that it is probabilistic. Reflection on *prima facie* duties will increase the likelihood of doing the right thing, but not guarantee it. Moral reasoning resembles prudential decision making under uncertainty.

> [I]t is certain that we are more likely in general to secure our advantage if we estimate to the best of our ability the probable tendencies of our actions in this respect, than if we act on caprice. And similarly we are more likely to do our duty if we reflect to the best of our ability on the *prima facie* rightness or wrongness of various possible acts in virtue of the characteristics we perceive

> them to have, than if we act without reflection. With this greater likelihood we must be content. (Ross and Stratton-Lake 2002, pp 31–32)

Ross might have been open to the supposition that the amount of *prima facie* rightness is measurable on a probability scale.

Perhaps Ittay Nissan-Rozen, who argued that, to account for comparative moral judgments, moral cognitivism must make use of credence instead of simple belief and be a form of Moral Bayesianism (2017). Perhaps Mark Schroeder, who has recently advanced the position that an objective reason can be a true proposition expressing a probability (2018). Perhaps, too, Sarah Moss, who has argued that the semantic contents of belief, knowledge, and assertion are not propositions at all, but are, instead, probabilistic contents (2018). Her views might be extended to give an account of the probabilistic contents of normative reasons.

Certainly, Stephen Kearns and Daniel Star. They argued that a fact is a reason for an agent to (form a belief or) perform an action if, and only if, the fact is evidence that the agent ought to perform that action. One argument that they offer for their theory is that it implies a useful way of measuring the degree to which a reason favours the performance of an action.

> The strength of evidence of the truth of a proposition, then, can be accounted for in terms of the probability of the proposition being true given this evidence. [new paragraph] This idea translates very nicely into an account of the strength of a reason. The strength of a reason to Φ, R, depends on the degree to which R increases the probability that one ought to Φ. The more probable it is that one ought to φ given R, the stronger reason to Φ R is. R is a stronger reason to Φ than R* is if and only if R makes the proposition that one ought to Φ more probable than R* makes it. R outweighs R* if and only if R is a reason to Φ, R* is a reason not to φ, and R makes the proposition that one ought to Φ more probable than R* makes the proposition that one ought not to Φ. Two reasons R and R* can combine to create a stronger reason to Φ if the probability that one ought to Φ given the conjunction of R and R* is greater than the probability that one ought to Φ given R and the probability of the same proposition given R*. The strength of a reason to perform an action can be accounted for in exactly the same ways as the strength of evidence can. (2009, 232; compare with Nair, 2016, pp 60–63).

In a footnote, Kearns and Star explain that when talking of probability, they have in mind the sort of evidential or epistemic probability defended by Timothy Williamson in chapter 10 of *Knowledge and its Limits* (2009, 232, n. 10). Williamson develops a form of objective Bayesianism that assumes an initial probability distribution measuring the intrinsic plausibility of a hypothesis prior to investigation instead of measuring a subject's prior credence in the hypothesis (2002, p 211; for critical discussion see Hawthorne & Lasonen-Aarnio, 2009). Thus, when Kearns and Star write about the "probability of the proposition being true given the evidence," they are not employing the notion of a personal, subjective probability or degree of belief. Nor are they employing the notion of an objective chance or physical probability for, as Williamson observes, the objective chance that a natural law is true may be one, whereas the epistemic probability of the law being true given our present scientific evidence may be much less than that (2002, 209). Instead, Kearns and Star employ the notion of an epistemic or evidential probability.

Kearns and Star suggest a probability-raising account of reasons in which a reason in favour of Φ-ing raises the probability that one ought to Φ, so that $\Pr(\Phi|R) > \Pr(\Phi)$. They worry that if the total evidence favouring Φ-ing is already extremely high, so that the reason for Φ-ing is already almost conclusive, then the reason, R, will increase the probability that one ought to Φ by very little (2002, 232, n. 10). R will become what Brunero calls a "lightweight" reason (Brunero, 2009, p. 540). To avoid this conclusion, they suggest that the reasons relative to which the probability of Φ-ing is measured should be a "salient relevant subset" of the total evidence for Φ-ing. Brunero complains that they do not give an adequate account of what evidence is salient, which makes it difficult to assess their suggestion (2009, p. 543).

Instead of trying to develop and defend their solution of the lightweight reason problem, I think they should embrace the conclusion, defended by Dancy, that the strength of reasons is not invariant, but instead changes with context (Dancy, 2000, 2004, chapter 5). Suppose that, by simply pressing a button, I can save a little girl from drowning in a pond. The fact that she is drowning in the pond gives me an extremely strong, almost conclusive, reason to press the button. Then I realize that the same button-press will save a little boy who is also drowning in the pond. Now I have an even stronger reason for pressing the button. But it is only a slightly stronger reason since the reason that I already had is already almost conclusive. The analogy between the strength of reasons and the physical weights of reasons breaks down here. Physical weights, at least on the surface of the Earth, are almost independent of position and thus constant within a small margin of error. By contrast, the strengths of reasons change with context. If the little girl were not drowning in the pond, then the fact that the little boy is drowning would be an extremely strong, almost conclusive, reason for me to press the button. The content of this reason—that the little boy is drowning—stays the same whether he is alone or with the little girl, but the strength of the reason varies with the background context, that is, with whether the little girl is also drowning in the pond (Kagan, 1988; Dancy, 2000, p. 132).

The strength of empirical evidence behaves in the same way as the strength of reasons. Intuitively, the third black raven that ornithologists ever observed was stronger evidence for ravens being black than was the thousandth raven that they observed. Even though the content of the two observations was equivalent, the strength of the observations, as evidence for a hypothesis, changed with context. It would be better for Kearns and Star to leave their account of the strength of reasons in terms of raising probabilities relative to one's entire body of evidence unmodified by an unnecessary and potentially misleading theory of salient, relevant evidence. Soon after Kearns and Star suggested using epistemic probabilities to measure the strengths of reasons, John Brunero offered counterexamples to their claim that reasons are probability-raising (Brunero, 2009, 2018). Brunero's counterexamples employ Jonathan Dancy's account of how intensifiers and attenuators modify the strengths of reasons without themselves favouring the conclusion in question (2004, pp. 41–42). Star has recently acknowledged that Brunero's criticisms require them "to give up using that account of evidence to explain the weight of reasons, or to say something more complicated that we have said in the past about how the weight of reasons is to be explained in terms of epistemic probability." (Star, 2018, 8 fn. 8) This retraction is premature. In what follows, I will say "something more complicated" about the calculation of conditional probabilities. First, I will describe a probabilistic account of how to modify the strengths of reasons in the situations that Dancy describes. Then, I will show how the claim that this more complicated account of how to measure the strengths of reasons evades Brunero's counterexamples to the probabilism of Kearns and Star.

3 Intensifiers, attenuators, and undercutters

A modifier is a consideration that changes the support given by a reason to a conclusion, but which is not, by itself, a reason in favor of the conclusion (Lord & Maguire, 2016, pp. 10–12). If a modifier increases the strength of a reason, then it is an intensifier, and if it decreases the strength of a reason, then it is an attenuator. Dancy gives the following examples of considerations that intensify or attenuate the support given by other premises.

> 1. She is in trouble and needs help.
> 2. I am the only other person around.
> 3. So: I help her.
>
> That she is in trouble and needs help is a consideration that favours my helping her. That I am the only other person around does not seem to be another reason, on top of the first one. It is not as if, even if she were not in trouble, that I am the only other person around would still favour my helping her. The reverse, however, is true; even if there were others around, I would still have a reason to help her, a reason given by the trouble she is in. But my being the only other person around does make a rational difference, all the same. I suggest that what it does is to intensify the reason given me by her need for help. Instead of two reasons, what we have here is one reason and an intensifier. Presumably, there is the opposite of an intensifier, an attenuator. An example might be this:
>
> 1. She is in trouble and needs help.
> 2. It is all her own fault, and she got in this situation through trying to spite someone else.
> 3. But still: I help her.
>
> One might think that there is a reason to help her even though it is all her own fault, etc., but less reason than there would otherwise be. (Dancy 2004, pp. 41–42)

Dancy also describes the related notions of disablers and enablers. A disabler is a premise that, without itself undermining a moral conclusion, can deactivate another premise. For example, the consideration that I ought to achieve world peace all by my myself is a reason for the conclusion that I ought to spend my life in pursuit of world peace, but this premise is disabled by the fact that I am powerless to bring about world peace on my own.

Attenuators in metaethics correspond to partial undercutters in epistemology. Testimonial evidence nicely illustrates how one piece of evidence can undercut another. Suppose Tanya tells me that the meeting is at three. Tanya's testimony functions as a premise in my reasoning to the conclusion that the meeting is at three and raises the probability that the meeting is at three. Suppose I meet Ruth who tells me that the meeting is not at three. Ruth's testimony lowers the probability of my conclusion that the meeting is at three by being evidence that the meeting is not at three. Ruth's testimony is a rebutting partial defeater to Tanya's testimony.

Suppose that instead of Ruth I meet Ursula, and Ursula tells me that Tanya is very unreliable about meeting times. Ursula's testimony does not directly affect the probability of the meeting being at three; in the absence of Tanya's testimony, Ursula's testimony would have had no effect on the probability that the meeting is at three. Instead, Ursula's testimony indirectly lowers the probability that the meeting is at three by affecting how Tanya's testimony should alter the probability of the meeting being at three. There is a direct epistemic relationship between Tanya's testimony and Ursula's testimony, but no direct epistemic relationship between Ursula's testimony and whether the meeting is at three. Ursula's testimony is an undercutting partial defeater of Tanya's testimony (Pollock 1984, Pollock and Cruz 1999; Schroeder 2011; Nair and Horty 2018).

Now, suppose that instead of meeting Ursula, I meet Barbara, who testifies that Tanya is very precise and reliable about meeting times. Barbara's testimony buttresses Tanya's testimony that the meeting is at three even though, considered on its own, it is irrelevant to the time of the meeting. Barbara's testimony slightly raises the probability that the meeting is at three. Intensifiers in metaethics would correspond, in epistemology, to considerations that buttress other considerations supporting a conclusion without themselves supporting the conclusion.

4 A probabilistic model of modifiers

I now offer a probabilistic account of how to aggregate reasons that is more complicated than the original account of Kearns and Star. They made the following suggestion about how the strengths of reasons combine with one another.

> Two reasons R and R^* can combine to create a stronger reason to Φ if the probability that one ought to Φ given the conjunction of R and R^* is greater than the probability that one ought to Φ given R and the probability of the same proposition given R^*. (2009, p. 232)

Put into symbols, they say that two reasons combine into a stronger reason if $\Pr(\Phi|R\&R^*) > \Pr(\Phi|R)$ and $\Pr(\Phi|R\&R^*) > \Pr(\Phi|R^*)$. To apply their suggestion, we need be able to calculate $\Pr(\Phi|R\&R^*)$. It is at this point that their idea that reasons are evidence has the potential to become misleading.

Whether or not Kearns and Star meant it to do so, thinking of reasons as being evidence encourages the impression that moral reasoning always involves the sequential updating of the conclusion's strength as new evidence is brought into play. In thinking of new evidence, we assume that pieces of evidence are independent of one another, that the evidence is fixed and certain, and that a rational thinker will update her credence in light of each new piece of evidence. To do so, a rational thinker will use the Bayesian Principle of Conditionalization, which relates the posterior credence of a hypothesis after new evidence becomes certain to its prior credence at an earlier time before the new evidence was obtained. The Bayesian Principle of Conditionalization tells us how to update credence assignments at time-two based on a prior credence assignment at time-one and the assumption that the reason R is true.

$$(BPC) \quad cr_{t=2}(\Phi) = \frac{cr_{t=1}(R|\Phi)}{cr_{t=1}(R)} \times cr_{t=1}(\Phi)$$

The Bayesian Principle of Conditionalization (BPC) is based on Bayes' Theorem (BT), which is provable from the probability axioms and the ratio formula for conditional probabilities.

$$(BT) \quad \Pr(\Phi|R) = \frac{\Pr(R|\Phi)}{\Pr(R)} \times \Pr(\Phi)$$

Treating reasons as evidence and updating according to the Bayesian Principle of Conditionalization becomes misleading when we have a second moral reason, R^*, and wish to update our new posterior probability for Φ at time-three. For it suggests that a simple formula for the probability of Φ given R and R^* follows from applying (BT) a second time to get an iterated corollary of Bayes' Theorem for two premises.

$$(\#BT2) \quad \Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|\Phi)}{\Pr(R^*)} \times \frac{\Pr(R|\Phi)}{\Pr(R)} \times \Pr(\Phi)$$

The formula, (#BT2), is a misleading simplification for the inductive logic of the situation because it hides all relationships of probabilistic dependence and conditional dependence between the premises of the argument. By assuming that the probability of each piece of evidence remains fixed at one, it prevents consideration of dependence relations between the pieces of evidence that might change their probabilities.

The actual two-premise corollary to Bayes Theorem is different. In an appendix, I give a quick derivation for this general theorem, (G), because I rely heavily on it below, and the derivation is not readily available. (G) is the full formula for the conditional probability of a conclusion that is based on two considerations.

$$(G) \quad \Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi)}{\Pr(R^*|R)} \times \Pr(\Phi|R)$$

(G) is the "something more complicated" about the calculation of conditional probabilities that produces a plausible account of how to modify the strengths of reasons and evades Brunero's counterexamples.

The two-premise corollary of Bayes' theorem, (G), is a general formula because it allows for the dependence of one reason on another. The term, $\Pr(R^*|R)$, in the denominator of the first fraction on the right-hand side, allows for cases where the truth of one reason, R, can change the probability of the other reason, R^*. Intuitively, if R is relevant to the truth of R^*, then the probability of R^* will depend on the probability of R. If reason R^* is not relevant to R, then $\Pr(R^*|R) = \Pr(R^*)$ and the truth of R will have no effect on the truth of R^*. For example, when we flip a coin, we assume that each flip is independent of the last one. To assume otherwise is to commit the gambler's fallacy; a bad gambler will assume that a long run of heads, say, is more likely to be followed on the next flip by tails than it is heads.

It follows that if R^* were dependent on R, so that $\Pr(R^*|R) > \Pr(R^*)$, then R^* would make a smaller contribution to the magnitude of $\Pr(\Phi|R\&R^*)$ than it would have done had it been independent of R. We are familiar with how, when pieces of evidence are dependent on one another, this affects the degree of support that they give to a hypothesis. For example, a child's age is evidence of whether she can read, and a child's shoe size is evidence of whether she can read, but because age and shoe size are dependent on one another, the contribution to the probability that she can read that comes from learning her age, after already learning her shoe size, is lower than it would be if age and shoe size were independent.

As well, the term, $\Pr(R^*|R\&\Phi)$, in the numerator of the first fraction on the right-hand side of (G) expresses the conditional dependence of R^* on R given the conclusion, Φ. Reason R is relevant to the strength of the reason R^* when we assume that Φ is true. Intuitively, two propositions are conditionally dependent relative to a third if, and only if, holding the first reason, R, to be true while accepting Φ, changes the probability of the other reason, R^*. Two reasons are conditionally independent relative to a conclusion if, and only if, holding the first proposition to be true while accepting the conclusion does not change the probability of the second. For example, reading ability is dependent on shoe size, but reading ability is conditionally independent of shoe size if we assume the child is of a certain age. For children of a given age, reading ability is not correlated in any way with shoe size. In the language of probabilistic approaches to causality, the age of the child is a common cause of both the child's shoe size and the child's reading ability. The age of the child screens off the spurious dependence of reading ability on shoe size. My suggestion will be that the conditional dependence of modifiers gives an account of considerations that modify other reasons.

The formula, (G), is a general one because it assumes neither the independence nor the conditional independence of the two premises. We can simplify formula (G) if we assume that R^* is conditionally independent of R given C, so that $\Pr(R^*|R\&\Phi) = \Pr(R^*|\Phi)$. We can further simplify the formula if we assume that R and R^* are independent, so that $\Pr(R^*|R) = \Pr(R^*)$. Putting these two assumptions into formula (G) and rearranging the terms, we get a simple ("naïve") formula for the two-premise corollary of Bayes' Theorem (N), which resembles the misleading formula (#BT2):

$$(N) \quad \Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|\Phi)}{\Pr(R^*)} \times \Pr(\Phi|R)$$

(N) is a naïve version of the two-premise corollary of Bayes' Theorem because it neglects the possible dependence and conditional dependence of R and R^* (Anderson, 2007). It holds, however, in important engineering applications such as calculating the reliability of a network where the probabilities of failure for the various components are independent of one another. Bayesian epistemologists assume it to hold for repeated updating of credence based on new evidence through Bayesian conditionalization. Neglecting the complicated dependencies that occur in the general formula, (G), can make otherwise daunting calculations tractable. However, neglecting these probabilistic dependencies can also lead to missing important ways in which reasoning is defeasible and probabilities are subject to modification. (G) is a mathematical theorem provable from the probability axioms and the definition of a conditional probability, whereas (N) is a derived from (G) under two simplifying assumptions: (1) R^* is conditionally independent of R given C, and (2) R^* and R are independent.

We can use (G) to model Dancy's examples of intensifiers and attenuators. It turns out that these considerations change the degree of support offered by reasons to conclusions because the modifiers, the reasons, and the conclusions are not conditionally independent of one another. M is a modifier in the context of a piece of reasoning involving reason R and conclusion C just in case,

(a) $\Pr(C|M) = \Pr(C)$, and
(b) $\Pr(C|R\&M) \neq \Pr(C|R)$

According to clause (a), M does not change the probability of the conclusion, C. In other words, C is independent of M. However, according to clause (b), M together with R increases (intensifies) or decreases (attenuates) the strength of our belief C from what it would have been if we had only premise R. This happens, I will argue, because M is conditionally dependent on R given C.

Recall Dancy's first example from above. In this argument, the first premise, P_1, is a reason to help her, and the second premise,P_2, has the role of an intensifying modifier:

P_1 = She is in in trouble and needs help.
P_2 = I am the only other person around.
C = I should help her.

P_2 is a modifying premise because, by itself, P_2 does not change the probability of the conclusion, C. Thus, C is independent of P_2, and condition (a) holds, $\Pr(C|P_2) = \Pr(C)$. Knowing P_2, that I am the only other person around, by itself gives me no reason to help her or not to help her. To show that P_2 is a modifier, we need to show that $\Pr(C|P_1\&P_2) \neq \Pr(C|P_1)$. We will do this by showing that $\Pr(C|P_1\&P_2) > \Pr(C|P_1)$, despite it being the case that, by itself, P_2 does not favour C. To show that P_2 is an intensifier, we substitute our premises into the general formula for combining reasons, (G).

$$(G) \quad \Pr(C|P_1\&P_2) = \frac{\Pr(P_2|P_1\&C)}{\Pr(P_2|P_1)} \times \Pr(C|P_1)$$

By inspecting (G), we can see that P_2 will raise the probability of C if the fraction that is the first term on the right-hand side of (G) is greater than one. This fraction will be greater than one if and only if its numerator is greater than its denominator, $\Pr(P_2|P_1\&C) > \Pr(P_2|P_1)$.

In the denominator, whether she is in trouble is independent of whether I am around, so $Pr(P_2|P_1) = Pr(P_2)$. That she needs help, by itself, does not raise the probability that I am the only person around.

In the numerator, given that she is in trouble and needs help, and given that that I (and not someone else) should help her, then the probability that I am the only person around will be higher than it would otherwise be. To see this, suppose, contrary to the assumption that I am alone, that there are other people around better placed to help. On this supposition, I would have less reason to be the one who helps her. Knowing that she is in trouble (P_1) and that I ought to help her (C) does increase the probability that I am the only person around (P_2).

Putting this together, we see that the consideration that I am the only person around raises the probability of the numerator and leaves the probability of the denominator unaffected. Thus, the fraction in the right-hand side of (G) will have a value greater than one, and formula (G) will show that the reason and the modifier together will raise the probability that I ought to help her. The premise, P_2, that I am the only person around functions to raise the overall conditional probability that I should help her. The consideration that I am the only person around combines with the reason that she needs help to intensify the strength of my reason to help. Conditional probabilities allow us to model intensifiers.

This probabilistic model also gives the correct answer when applied to Dancy's example of an attenuator. Dancy claims that the consideration that it is all her own fault, and that she got in this situation through trying to spite someone else, is not by itself a reason not to help her. However, it does weaken the normative force of the consideration that she is in trouble and needs help, which is a reason to help her.

P_3 = She is in in trouble and needs help.
P_4 = It is all her own fault, and she got in this situation through trying to spite someone else.
C = So, I ought to help her.

Knowing P_4, that it is all her own fault, by itself, neither supports nor undermines the conclusion that I ought to help her. Thus, C is independent of P_4, and $\Pr(P_4|C) = \Pr(P_4)$. Knowing that she is in trouble tells me nothing about the probability that it is all her own fault, so $Pr(P_4|P_3) = Pr(P_4)$. Knowing that she is in trouble (P_3) and that I ought to help her (C) decreases the probability that it is all her own fault (P_4), so $\Pr(P_4|P_3\&C) < \Pr(P_4)$. By an analogous argument to the one above, the left-most fraction on the right-hand side of (G) is less than one,

$$\frac{\Pr(P_4|P_3\&C)}{\Pr(P_4|P_3)} < 1$$

and the premise, P_4, that it is all her own fault will lower the overall probability that I should help her. So, conditional probabilities also allow us to model attenuators.

Attenuators in the theory of moral reasons function like partial undercutters in the theory of epistemic reasons. For example, Ursula's testimony that Tanya is unreliable attenuates or partially undercuts Tanya's testimony as a reason for me to believe that the meeting is at three. To see how Ursula's testimony lowers the probability that the meeting is at three, symbolize the argument as follows.

T = Tanya testifies the meeting is at three.
U = Ursula testifies that Tanya is unreliable about meeting times.
C = The meeting is at three.

The probability that Ursula's testimony that Tanya is unreliable, U, given the information that the meeting really is at three, C, and that Tanya says it is at three, T, $\Pr(U|T\&C)$, will be lower than the probability of Ursula's testimony without Tanya's testimony. So, Ursula's testimony will lower the value of the numerator of the above ratio. However, without the knowledge that the meeting really is at three, the probability of Ursula's testimony that Tanya is unreliable is independent of Tanya's testimony, so $\Pr(U|T) = \Pr(U)$ in the denominator, and the value of the denominator remains unchanged. Thus, the ratio in the leftmost term of the right-hand side of the general, two-premise, corollary of Bayes' theorem (G) is less than one.

$$\frac{\Pr(U|T\&M)}{\Pr(U|T)} < 1$$

Ursula's testimony partially undercuts Tanya's testimony by lowering the probability that the meeting is at three.

5 Brunero's Counterexamples to Kearns and Star

Brunero used modifiers to challenge the probability raising account of the strength of reasons given by Kearns and Star. Brunero's first counterexample goes like this (2009, 540). On its own, the fact, "Dad would be happy were I to get Mom some specific gift he found featured in the *Sears Catalog*," is a reason for Brunero to buy this gift since his father's increased happiness counts in favour of his buying the gift. However, in the context of the background information, "Whenever Dad would be happy with Mom getting some gift, there is always some competing, weightier reason(s) against getting that gift for Mom," the first consideration lowers the probability that Brunero ought to buy the gift. The second consideration attenuates or partially undercuts the first consideration as evidence for the conclusion. Brunero suggests that this is a counterexample to Kearns and Star's proposal that, for a fact to be a reason, it is necessary that it always raises the probability of the conclusion.

In his second counterexample, Brunero employs Dancy's notion of an enabler, a consideration that is not a favouring reason, but that permits another reason to have its full strength (Brunero, 2009, 544–545; 2018, 329–330). In Dancy's theory, enablers are limiting case of intensifiers, considerations that have the opposite effect of attenuators on the strengths of reasons (Dancy, 2004, 41–42). I will argue that, properly understood, conditional probabilities give an account of attenuators and intensifiers that avoids Brunero's challenges.

Reasoning that employs undercutting premises features in Brunero's first counterexample. His argument has two premises (Brunero, 2009, 540).

P_1 is the acquired information, "Dad would be happy were I to get Mom some specific gift he found featured in the *Sears Catalog*."
P_2 is the background information, "Whenever Dad would be happy with Mom getting some gift, there is always some competing, weightier reason(s) against getting that gift for Mom."
C is the conclusion that Brunero ought to buy the gift.

On its own, P_1 appears to be a reason for Brunero to buy the gift since his father's increased happiness counts in favour of his buying the gift. However, in the context of the background information contained in premise, P_2, P_1 will lower the probability that Brunero should buy the gift. The probability of C will go down in response to both premises so that $\Pr(C|P_1\&P_2) < \Pr(C)$. Brunero suggests that this is a counterexample to the proposal of Kearns and Star that for a fact to be a reason it is necessary that it raises the probability of the conclusion.

Brunero is correct that this is a counterexample if Kearns and Star are presuming iterated Bayesian updating. Iterated Bayesian updating has no mechanism to account for one premise attenuating or undercutting another. However, if Kearns and Star were to use a more general inductive logic and were to employ the two-reason corollary to Bayes' Theorem, (G), that allows for probabilistic dependencies among premises, then they could offer an account of this example.

Premise P_2 undercuts premise P_1. To see this, we need to go back to formula (G).

$$(G) \quad \Pr(C|P_1\&P_2) = \frac{\Pr(P_2|P_1\&C)}{\Pr(P_2|P_1)} \times \Pr(C|P_1)$$

The denominator of the first fraction on the right-hand side gives Brunero's conditional credence (P_2) that his Dad is always wrong about what gift to get his Mom given only the information, (P_1), that his Dad would be happy if Brunero bought this specific gift. The information that Brunero's Dad is always wrong about what gift to get his Mom is independent of the information that his Dad would be happy if Brunero bought this specific gift, so the probability of P_2 will be unchanged by his acquired information, P_1. That is to say, $Pr(P_2|P_1) = Pr(P_2)$.

The numerator on the right-hand side, $\Pr(P_2|P_1\&C)$, gives the conditional probability that his Dad is always wrong about what gift to get his Mom (P_2) given that his Dad would be happy if Brunero got a specific gift (P_1) and that Brunero ought to get this specific gift (C). The additional information in the numerator, that Brunero ought to get this specific gift, when it is added to the information that his Dad would be happy if Brunero bought his Mom this gift, will lower the probability that his Dad is always wrong about what gift to get his Mom. Under these assumptions, his Dad was not wrong. So, the probability of P_2, given both P_1 and C, will go down. Hence, the ratio that is the first term on the right-hand side of (G) will be less than one, the two premises together will lower the probability that Brunero ought to buy the gift, and $\Pr(C|P_1\&P_2) < \Pr(C|P_1)$. Thus, the effect of the background information is to lower Brunero's confidence that he ought to buy the gift, which is what we are aiming to explain.

Brunero also offered a counterexample to the contention of Kearns and Star that if a premise raises the probability of a conclusion, then it is a stronger reason for the conclusion. This time the argument goes like this (Brunero, 2009, 544).

P_1 is the premise, "I have promised to Φ."
P_2 is the premise, "There is no reason for me not to Φ."
C is the conclusion that Brunero ought to Φ.

Premise P_1 is a reason for Brunero to Φ. Premise P_2 does not count in favour of Brunero Φ-ing, and is not, itself, a reason for him to Φ. However, it does raise the probability that he should Φ, and so appears to be a counterexample to the claim of Kearns and Star that raising the probability of a conclusion is sufficient for a premise to be a stronger reason for a conclusion.

The situation here is the opposite of the earlier case. Instead of undercutting or attenuating the first premise, the second premise buttresses or intensifies the first premise. However, we can give the same sort of explanation of intensifying as we gave of attenuating by using the general two-premise corollary of Bayes' Theorem, which allows for probabilistic dependencies between premises, instead of using the Bayesian Principle of Conditionalization, which does not allow for such dependencies.

Brunero asserts that the second premise is what Dancy would call an enabler rather than an intensifier. However, Dancy also suggests that an analysis of the concept of an intensifying consideration is effectively an analysis of an enabling consideration.

> If we are happy with the idea that a reason can be attenuated by a consideration which is not itself a reason, why should we be shy of supposing that it can be reduced to nothing? And if we accept this last possibility, why should we not suppose that the reduction can be achieved all at once rather than by degrees? – which is after all exactly what a disabler does. And if we allow the idea of a disabler, we will have to allow that of an enabler as well. (2004, p. 42)

A conditional probability model should be shy of assigning a probability of zero to the support given by a premise to a conclusion. Doing so may lead to a situation involving division by zero or to a situation where a premise is powerless to raise the probability of a conclusion because the probability of the conclusion has become zero. However, the support given by a premise to a conclusion can be attenuated to as close to zero as we might like, and this will provide a good enough model of a disabler, and a good enough reply to Brunero's second counterexample.

It is important to focus here on the strength of reasons rather than on the nature of reasons. Kearns and Star do not say that facts are reasons if, and only if, they are probability-raising. What they do say is that the strength of a reason for a conclusion depends on the degree to which it raises the conditional probability of the conclusion. "The more probable it is that one ought to φ given R, the stronger reason to φ R is." (2009, p. 232) They claim that understanding reasons to be evidence is the best way to explain how to measure the strengths of reasons. The above analysis, which says something more complicated about probabilities, is a defence of their probabilism about the strengths of reasons. The suggestion that the strengths of reason are conditional probabilities, when it is suitably refined, does explain how modifiers work and can avoid Brunero's counterexamples.

The supposition that conditional probabilities are a guide to reasoning both about how to act and about what to believe accounts for the role of modifiers. So, the supposition that conditional probabilities measure the strengths of reasons, and that reasoning about actions and beliefs is a form of inductive reasoning, is worth further investigation alongside such alternatives as weighing explanations (Broome 2013) and default reasoning (Horty 2007, 2012). If my probabilistic account of how to modify the strength of reasons is successful, then it is a reason in favor of the original supposition that we ought to use conditional probabilities to measure the strengths of reasons.

Appendix: How to derive formula (G)

Suppose we have a piece of moral reasoning where two reasons, R and R^*, each either supports or counts against a conclusion, Φ. To calculate the conditional credence of Φ based on the conjunction of R and R^*, we start with the ratio-formula definition of conditional credence, $Pr(A|B) = Pr(A\&B)/Pr(B)$ and substitute Φ for A and $R\&R^*$ for B. By commuting the conjunctions in both the numerator and the denominator, we get:

$$\Pr(\Phi|R\&R^*) = \frac{\Pr(\Phi\&(R\&R^*))}{\Pr(R\&R^*)} = \frac{\Pr(R^*\&(R\&\Phi))}{\Pr(R^*\&R)} \tag{1}$$

From the ratio formula for conditional probability, we also get the product rule for the probability of a conjunction, $Pr(A\&B) = Pr(A|B) \times Pr(B)$. When we apply the product rule to the numerator of (1) by substituting R^* for A and $(R\&\Phi)$ for B, we get:

$$\Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi) \times \Pr(R\&\Phi)}{\Pr(R^*\&R)} \tag{2}$$

When we again apply the product rule to $\Pr(R\&\Phi)$ in the numerator of (2), we get:

$$\Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi) \times \Pr(R|\Phi) \times p(\Phi)}{\Pr(R^*\&R)} \tag{3}$$

When we apply the product rule to the denominator of (3), we get:

$$\Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi) \times \Pr(R|\Phi) \times \Pr(\Phi)}{\Pr(R^*|R) \times \Pr(R)} \tag{4}$$

By rearranging the terms in both the numerator and denominator of (4), we get:

$$\Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi)}{\Pr(R^*|R)} \times \frac{\Pr(R|\Phi)}{\Pr(R)} \times \Pr(\Phi) \tag{5}$$

Looking back into the main text, we can see that the last two terms on the right-hand side are the right-hand side of Bayes Theorem (BT).

$$(BT) \quad \Pr(\Phi|R) = \frac{\Pr(R|\Phi)}{\Pr(R)} \times \Pr(\Phi)$$

Substituting this into (5) we get a general mathematical theorem (G) for the case of two reasons.

$$(G) \quad \Pr(\Phi|R\&R^*) = \frac{\Pr(R^*|R\&\Phi)}{\Pr(R^*|R)} \times \Pr(\Phi|R)$$

Notice that because all the operations involved in the proof are commutative, the proof will be symmetric in R and R*. We could equally well have derived the inverted formula, (G*).

$$(G^*) \quad \Pr(\Phi|R\&R^*) = \frac{\Pr(R|R^*\&\Phi)}{\Pr(R|R^*)} \times \Pr(\Phi|R^*)$$

We can extend the above derivation to produce corollaries of Bayes' Theorem for inductive arguments that involve as many premises as we may wish, though the corollaries will rapidly become very complicated. (G) is a useful formula for combining the support offered by two moral premises into an overall degree of support for a conclusion. It shows right away that the degrees of support given to a conclusion by each of two considerations combine multiplicatively, not additively (Kagan 1998).

Acknowledgements I would like to thank David White for editorial assistance in the preparation of this paper, and thank Erik Nelson, Sheldon Wein, Bill Seager, and members of the audience at a Dalhousie Philosophy Colloquium for helpful comments. A referee for this journal gave me excellent, critical, but much appreciated advice.

References

Anderson, S. (2007). Combining evidence using bayes' rule. <cs.wellesley.edu/~anderson/writing/naive-bayes.pdf>

Broome, J. (2008). Reply to Southwood, Kearns and Star, and Cullity. *Ethics, 119*, 96–108.

Broome, J. (2013). *Rationality through Reasoning*. Wiley-Blackwell.

Brunero, J. (2009). Reasons and evidence one ought. *Ethics, 119*, 538–545.

Brunero, J. (2018). Reasons, evidence, and explanations. In D. Star (Ed.), *The Oxford handbook of reasons and normativity* (pp. 321–341). Oxford University Press.

Dancy, J. (2000). The particularist's progress. In B. Hooker & M. Little (Eds.), *Moral particularism*. Oxford University Press.

Dancy, J. (2004). *Ethics without principles*. Clarendon Press.

Hawthorne, J., & Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In P. Greenough & D. Prichard (Eds.), *Williamson on knowledge*. Oxford University Press.
Horty, J. (2007). Reasons as defaults. *Philosophers' Imprint, 7*, 1–28.
Horty, J. (2012). *Reasons as defaults*. Oxford University Press.
Kagan, S. (1988). The Additive Fallacy. *Ethics, 99*, 5–31.
Kearns, S., & Star, D. (2008). Reasons: Explanations or evidence? *Ethics, 118*, 31–56.
Kearns, S., & Star, D. (2009). Reasons as evidence. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics (4)* (pp. 215–242). Oxford University Press.
Kearns, S., & Star, D. (2013). Weighing reasons. *Journal of Moral Philosophy, 10*, 70–86.
Lord, E., & Maguire, B. (2016). An opinionated guide to the weight of reasons. In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 3–24). Oxford University Press.
Moss, S. (2018). *Probabilistic knowledge*. Oxford University Press.
Nair, S. (2016). How do reasons accrue? In E. Lord & B. Maguire (Eds.), *Weighing reasons.* Oxford University Press.
Nissan-Rozen, I. (2017). Reasoning with comparative moral judgements: An argument for Moral Bayesianism. In R. Urbaniak & G. Payette (Eds.), *Applications of formal philosophy: the road less travelled* (pp. 113–136). Springer.
Pollock, J. (1984). Reliability and justified belief. *Canadian Journal of Philosophy, 14*, 103–114.
Pollock, J., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Rowman & Littlefield.
Schroeder, M. (2011). Holism, weight, and undercutting. *Nous, 45*, 328–344.
Schroeder, M. (2018). Getting perspective on objective reasons. *Ethics, 128*, 289–319.
Star, D. (2018). Introduction. In D. Star (Ed.), *The oxford handbook of reasons and normativity* (pp. 1–22). Oxford University Press.
Way, J. (2017). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly, 98*, 251–270.
Williamson, T. (2002). *Knowledge and its limits*. Oxford University Press.