

How Explanation Guides Confirmation¹

Nevin Climenhaga
Department of Philosophy
University of Notre Dame
nclimenh@nd.edu

Abstract

Where E is the proposition that [If H and O were true, H would explain O], William Roche and Elliot Sober have argued that $P(H|O\&E) = P(H|O)$. In this paper I argue that not only is this equality not generally true, it is false in the very kinds of cases that Roche and Sober focus on, involving frequency data. In fact, in such cases O raises the probability of H *only* given that there is an explanatory connection between them.

¹ I am grateful to Robert Audi, Daniel Immerman, Ted Poston, and two reviewers for *Philosophy of Science* for very helpful comments on earlier drafts of this paper.

In two recent essays, Roche and Sober (2013, 2014) (hereafter R&S) argue that the proposition that a hypothesis H would explain an observation O is *evidentially irrelevant* to H . Where E says that, were H and O true, H would explain O , R&S's thesis is that

$$(1) P(H|O\&E) = P(H|O).$$

Once we know O , R&S claim, E gives us no further evidence that H is true. In other words, O *screens off* E from H . Call this claim the *Screening-Off Thesis* (SOT) (R&S 2014: 193).

In endorsing SOT, R&S are presumably not making a claim about subjective probabilities, for an agent could virtually always assign coherent subjective probabilities on which the above equality is false. I will instead read them as making a claim about epistemic probabilities, which we can understand as rationally constraining subjective probabilities.

Theses similar to SOT are endorsed by other Bayesians skeptical of inference to the best explanation. For example, van Fraassen (1989: 166) famously denies that the claim that a hypothesis is explanatory can give it any probabilistic "bonus." Often, however, such skeptics do not make clear in precisely what way they think that explanation is irrelevant to confirmation. Van Fraassen does consider a precise version of inference to the best explanation, but it is an uncharitable one, on which inference to the best explanation is understood as a non-Bayesian updating rule on which good explanations get higher probabilities than Bayesian conditionalization would give them.² R&S are thus to be commended for stating a precise anti-explanationist thesis which does not mischaracterize their opponents' position. If SOT is true, there is a clear sense in which explanation is not relevant to confirmation.

That said, we do need to clarify the scope of SOT before we can evaluate its significance and plausibility. It is widely acknowledged by Bayesians that all confirmation is relative to a context. In other words, we always have some background knowledge K , which may be left implicit but is always guiding our judgments of probability. While subjective Bayesians often think of this background as being part of the probability function $P(\cdot)$ itself, for epistemic probabilities, it is preferable to make K explicit as a conjunct of the proposition being conditioned on.³ So we can rewrite R&S's claim as

² Some explanationists have defended this rule against van Fraassen's criticisms of it (see, e.g., Douven 2013). However, whether or not van Fraassen's original arguments against non-Bayesian explanationist updating rules work, I show in my 2016a that such rules lead to synchronic probabilistic incoherence. The argument of this paper suggests that it would be better to think of the "probabilistic bonus" that explanation gives to a hypothesis H as the degree to which the proposition that H is explanatory confirms H .

³ Making K part of the probability function itself makes it impossible to "bring out" any part of K in the way one brings out the evidence in Bayes' Theorem – i.e., where X is a conjunct of K and Y is an arbitrary proposition, the probability of X given Y will always equal 1. This leads to a version of the old evidence problem. By contrast, if K is one of the propositions

$$(2) P(H|O\&E\&K) = P(H|O\&K).$$

Here H, O, and K are variables that could be filled in with various propositions or background knowledge, with the caveat that H is a hypothesis and O an observational statement. We can then ask: for which H, O, and K do R&S take (2) to be true?

The most straightforward interpretation of SOT is as a universal claim: for all O, H, and K, $P(H|O\&E\&K) = P(H|O\&K)$. However, this interpretation is uncharitable because it is trivially false. For example, suppose K includes the material conditional, $[(O\&E) \supset H]$, but does not include the material conditional $[O \supset H]$. (Perhaps an oracle who knows whether O, H, and E are true has told you that if O and E are true, so is H.) Then $P(H|O\&E\&K) = 1$, but $P(H|O\&K) < 1$. Hence,

$$(3) P(H|O\&E\&K) > P(H|O\&K).^4$$

Although the above example shows SOT to not be universally true, $[(O\&E) \supset H]$ is not the kind of information we would ordinarily have as part of our background knowledge. In endorsing SOT, then, R&S presumably mean to hold that (2) is true in ordinary or paradigm cases, in particular, the kinds of cases in which defenders of inference to the best explanation wish to claim that explanatoriness *is* evidentially relevant. I will now argue that even so restricted, SOT is false. I will do this by considering one of the paradigmatic statistical cases R&S use to argue for SOT, involving smoking and cancer.

conditioned on, and not part of the probability function itself, then X can be brought out in such a way that the probability of X given Y does not necessarily equal 1. (For more on this point, see section III of Climenhaga 2016b.)

⁴ R&S consider the possibility of counterexamples to SOT if one drops the standard Bayesian assumption of logical omniscience:

Let I be the proposition that H logically implies O. Then, plausibly, there can be cases where $\Pr(H | O \& I) > \Pr(H | O)$ and $\Pr(H | O \& I \& E) > \Pr(H | O)$. Perhaps there can even be cases where $\Pr(H | O \& E) > \Pr(H | O)$. This would be especially plausible if E were in some way indicative of I. But then the point would be that $\Pr(H | O \& I \& E) = \Pr(H | O \& I)$. Explanatoriness has no confirmational significance, once purely logical and mathematical facts are taken into account (R&S 2014: 195).

However, these comments do not undermine the counterexample in the text. First, that counterexample involves knowledge of a material implication, not a logical implication. ($[(O\&E) \supset H]$ could equivalently be stated as $[\sim(O\&E) \vee H]$.) Thus, the assumption that there are some contexts in which one does not know $[(O\&E) \supset H]$ does not violate logical omniscience. Second, R&S's example involves I describing an implication relationship between O and H on their own. This is what lets them say that $P(H | O\&E\&I) = P(H | O\&I)$. In the counterexample in the text K says that O&E, but not O, implies H. Hence, in the circumstances I have described, $P(H|O\&E\&K) > P(H|O\&K)$.

Suppose that K includes statistical data on which smoking and cancer are correlated as well as the base rate of cancer in the population. R&S and I agree (2013: 660) that

$$(4) P([S \text{ gets cancer}] \mid [S \text{ smokes}] \& K) > P([S \text{ gets cancer}] \mid K),$$

which implies

$$(5) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K) > P([S \text{ smokes}] \mid K).$$

In a response to R&S's 2013 paper, McCain and Poston (2014: 150) (hereafter M&P) briefly argue that the frequency data in K make (4) true only because they support the existence of some causal connection between smoking and cancer:⁵

[T]he data indicated that there was some causal process—albeit unknown at that time—that explains the correlation between smoking and lung cancer. ... Exactly this feature—a justified belief in an unknown explanatory story—plays a crucial role in using the data from observation to get justified beliefs about the relevant frequencies. Apart from a general justified belief in some explanatory story accounting for that data, the observational data would not justify beliefs about the relevant frequencies.

Although M&P are talking about objective frequencies and not epistemic probabilities in this quote, it is similarly true that $[S \text{ smokes}]$ supports $[S \text{ gets cancer}]$ relative to K only because K supports the existence of an explanatory connection between $[S \text{ smokes}]$ and $[S \text{ gets cancer}]$. As I go on to argue, this means that SOT is false.

M&P do not formalize this point, and perhaps for this reason do not recognize this implication of it. Instead they write that it shows that “even if one grants Roche and Sober’s claim that ‘ O screens-off E from H ’ this doesn’t show that explanatory considerations are irrelevant to confirmation.”⁶ But it does not show *this*. Rather, it shows that explanatory

⁵ M&P’s main claim in that paper is that even if SOT is true, explanatoriness can still play an evidential role by increasing the “resiliency” of probabilities. In my view this is based on a mistaken view about epistemic probabilities. As I understand it, the epistemic probability of H given $O \& K$, $P(H|O \& K)$, is a relation between the propositions H and $O \& K$, such that, if $P(H|O \& K) = n$, then someone with $O \& K$ as their evidence ought to be confident in H to degree n . Following Keynes (1921), I take this relationship to be metaphysically necessary and knowable a priori, like the laws of logic or mathematics. Learning new empirical information, like E , does not affect the value or resilience of $P(H|O \& K)$. The value of this probability does not change, just as whether $A \& B$ entails A does not change. Rather, learning E simply makes a *new* probability relevant to what we should believe, namely $P(H|O \& E \& K)$, because now $O \& E \& K$ describes our total evidence.

⁶ Immediately after this, M&P say, ‘Explanatory considerations are already at work in setting

considerations are relevant to confirmation *because* R&S's Screening-Off Thesis (or a near-cousin, as I clarify below) is false. (M&P's confusion on this point may be part of the reason that R&S do not reply to it in their 2014 response to M&P.)

R&S (2013: 661) hold that, as SOT implies in this case,

$$(6) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& K) = P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K),$$

where E is the proposition, [If (S smokes) and (S gets cancer) were true, (S smokes) would explain (S gets cancer)]. R&S claim that "a good estimate of the probability on the right [of (6)] is furnished by frequency data; the same estimate is a good one for the probability on the left" (2013: 663).

I claim that (6) is not in general true. One reason this is hard to see immediately is that (6) involves the counterfactual, [If (S smokes) and (S gets cancer) were true, (S smokes) would explain (S gets cancer)]. In ordinary circumstances we would only get evidence for a counterfactual about a particular case like this by getting evidence for broader explanatory claims like [In general, smoking causes cancer]. As such, it will be helpful to start by considering claims of this form.

Say that there is an explanatory connection between two phenomena X and Y iff, at least sometimes, X causes Y, Y causes X, or X and Y have a common cause.⁷ Let C_1 be the hypothesis that smoking causes cancer, C_2 the hypothesis that cancer causes smoking, and C_3 the hypothesis that they have a common cause. $\sim[C_1 \vee C_2 \vee C_3]$, then, is the claim that there is no explanatory connection between smoking and cancer.

Let us now consider a revised Screening-Off Thesis, SOT*, which says that the existence of *general* explanatory connections is evidentially irrelevant. In this context, SOT* says that, for all K,

$$\text{Where } C_1 \text{ says that sometimes, smoking causes cancer, } P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& C_1 \& K) = P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K).$$

$\Pr(H|O)$ —having E provide additional confirmation for H would be akin to double-counting the information about objective chances.' This suggests that they may be thinking of ' $\Pr(H|O)$ ' as a frequency. But R&S's SOT is not about frequencies (although R&S may invite confusion on this point by moving back and forth between discussing epistemic probabilities and frequencies themselves). Or perhaps M&P are suggesting that the existence of an explanatory connection between smoking and cancer is part of the background information K. (In this case $P(H|O \& E \& K)$ would equal $P(H|O \& K)$ because K and E&K would be logically equivalent.) However, the existence of an explanatory connection is neither directly observed nor entailed by facts which are directly observed. Hence it should not be included in K.

⁷ I here ignore non-causal forms of explanation.

I will now show that SOT* is false in this context.

First, note that

$$(7) P([S \text{ gets cancer}] \mid [S \text{ smokes}] \& \sim [C_1 \vee C_2 \vee C_3] \& K) = P([S \text{ gets cancer}] \mid \sim [C_1 \vee C_2 \vee C_3] \& K).$$

This is because on the (extremely unlikely) hypothesis that there is no explanatory connection between smoking and cancer, the observed frequency data is a huge fluke. But we shouldn't expect huge flukes to continue. If the observed association of smoking and cancer is merely coincidental, then we should expect future smokers that we observe to have cancer at the same rate as the rest of the population. So if we know that $\sim [C_1 \vee C_2 \vee C_3]$, learning that S smokes doesn't raise the probability that S gets cancer above the probability given by the base rate of cancer in the population.

An anonymous reviewer suggests the following objection to this argument. (7) is an instance of the more general principle

$$P(A \mid B \& [\text{there is no explanatory connection between } A \& B] \& K) = P(A \mid [\text{there is no explanatory connection between } A \& B] \& K).$$

However, this principle is subject to counterexample.⁸ Sober (2001) observes that although there is no explanatory connection between the price of bread in Britain and the height of the sea in Venice, they are nevertheless correlated: they both tend to increase over time. As such, if K reports the bread price in Britain and the sea level in Venice historically, B says that the sea level in Venice is x at some unspecified future time t , and A says that the bread price in Britain is y at t , then this equality is false. For example, learning that the sea level in Venice is much higher than at present raises the probability that the bread price is also much higher than at present.

Assuming that time is not a common cause of A and B, then I accept the counterexample to the above principle, but I still hold that the principle is true in the current case. For, as Steel (2003: 313) observes, "British bread prices provide information about Venetian tides (and vice versa) *only in virtue of telling us something about the time*" (emphasis his). Knowledge of the time t screens off the bread prices from the sea levels. The above principle is plausibly true when any relevant temporal information is built into our background K, which we can stipulate is the case in the smoking and cancer example.

Even if this is not right, Sober (2001: 342-43) agrees that separate cause explanations "often" do not predict correlations, and I think he should accept that (7) is such a case. According to Sober, inference to a common cause is often rational because it is frequently the case that a common cause explanation predicts a correlation where a separate cause explanation does not. (Sober is considering cases where it is obvious that neither A nor B

⁸ This counterexample was originally applied to an analogous principle about frequencies.

causes the other, so a common cause is the only explanatory relation available.) However, it is clearly rational to infer a causal relationship between smoking and cancer from the frequency data in K. Hence, Sober's reasoning would suggest that (7) is a case in which the above principle is true.

Here is (7) again:

$$(7) P([S \text{ gets cancer}] \mid [S \text{ smokes}] \& \sim [C_1 \vee C_2 \vee C_3] \& K) = P([S \text{ gets cancer}] \mid \sim [C_1 \vee C_2 \vee C_3] \& K).$$

It follows from (7) that

$$(8) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& \sim [C_1 \vee C_2 \vee C_3] \& K) = P([S \text{ smokes}] \mid \sim [C_1 \vee C_2 \vee C_3] \& K).$$

Presumably learning $\sim [C_1 \vee C_2 \vee C_3]$ on its own does not affect the probability of [S smokes] relative only to our background knowledge K: without any information about whether or not S has cancer, these two propositions are irrelevant to each other. Hence,

$$(9) P([S \text{ smokes}] \mid \sim [C_1 \vee C_2 \vee C_3] \& K) = P([S \text{ smokes}] \mid K).$$

From (8) and (9) it follows that

$$(10) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& \sim [C_1 \vee C_2 \vee C_3] \& K) = P([S \text{ smokes}] \mid K).$$

In other words, learning that there is no explanatory connection between smoking and cancer *and* that S has cancer does not raise the probability that S smokes. However, according to (5), that S gets cancer raises the probability that S smokes. For (5) to be true, it must then be the case that learning that S smokes and that there *is* an explanatory connection between smoking and cancer raises the probability that S gets cancer. That is, from (5) and (10) it follows that

$$(11) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& [C_1 \vee C_2 \vee C_3] \& K) > P([S \text{ smokes}] \mid K).$$

From (10) and (11) it follows that

$$(12) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& [C_1 \vee C_2 \vee C_3] \& K) > P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& \sim [C_1 \vee C_2 \vee C_3] \& K),$$

which implies that

$$(13) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& [C_1 \vee C_2 \vee C_3] \& K) > P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K).$$

Presumably any one of C_1 , C_2 , and C_3 also licenses extrapolation from our frequency data. Consequently, we can replace $C_1 \vee C_2 \vee C_3$ with any one of C_1 , C_2 , and C_3 , and (11)-(13) will remain true. In particular, it will be true that

$$(14) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& C_1 \& K) > P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K).$$

(14) contradicts SOT*. So SOT* is false.⁹

This example shows two other things. First, it shows that R&S are wrong to claim (2013: 662) that the asymmetry of explanation suggests that explanatory facts like C_1 cannot be confirmatory. As they observe, confirmation is symmetric whereas explanation is not. If X explains Y, Y does not explain X, but if $P(X|Y \& K) > P(X|Y)$, $P(Y|X \& K) > P(Y|X)$. But this does not mean that [X explains Y] should not make a difference to the degree to which Y confirms X. In the above example, C_1 raises the probability of [S gets cancer] not by ruling out C_2 and C_3 , but by ruling out $\sim[C_1 \vee C_2 \vee C_3]$. So [X explains Y] can support Y not by ruling out [Y explains X] but by ruling out [There is no explanatory connection between X and Y].¹⁰

Second, note that (10) and (11) together say that [S gets cancer] raises the probability that [S smokes] when it is conjoined with the claim that there is an explanatory connection between smoking and cancer, but *not* when it is conjoined with the claim that there is no

⁹ An anonymous reviewer suggests the following argument against (14). $P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K)$ should be equal to the frequency of cancer among smokers given by K. But then (14) implies that $P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& C_1 \& K)$ is greater than the frequency of cancer among smokers, which seems wrong. My response to this argument is to reject its first premise. K includes data on the correlations between cancer and smoking in previously observed cases. We are extrapolating from this data to a new case, and in general we should not follow the “straight rule” in so extrapolating (compare R&S’s 2013: 663 coin toss example). This is clear in extreme cases: if all the people with cancer we have observed so far have been smokers, we should still not be 100% confident that the next person with cancer we observe will be a smoker. It is true that, as data accumulates, our new probabilities should tend to approach observed frequencies. But this is compatible with (14). The probability on the left-hand side of (14) is closer to the observed frequency of smoking among people with cancer than the probability on the right-hand side, but both values approach this frequency as the number of samples in K increases.

¹⁰ It is compatible with this that the *order* of explanation is evidentially irrelevant, in that X confirms Y to the same degree regardless of what the explanatory relationship between them is. But even this does not follow from R&S’s observation that confirmation is symmetric. For confirmation is only *qualitatively* symmetric: X confirms Y iff Y confirms X. It is not plausibly *quantitatively* symmetric: in general, X does not confirm Y to the same *degree* that Y confirms X (Eells and Fitelson 2001). So, for all R&S have said, X might well confirm Y more (or less) if Y explains X than if X explains Y.

explanatory connection between them. In other words, the existence of an explanatory connection between cancer and smoking is precisely what licenses the inference from S's smoking to S's cancer. Moreover, (8) says that $\sim[C_1 \vee C_2 \vee C_3]$ screens off [S gets cancer] from [S smokes]. The situation is thus almost the *opposite* of what R&S claim: not only does our observation not screen off our explanatory claim from our hypothesis, the negation of our explanatory claim ($\sim[C_1 \vee C_2 \vee C_3]$) screens off our observation from our hypothesis. Our explanatory claim thus *mediates* the move from observation to hypothesis.¹¹

(14) shows SOT* to be false. What about SOT? According to SOT,

$$(6) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& K) = P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K),$$

where E says that [If (S smokes) and (S gets cancer) were true, (S smokes) would explain (S gets cancer)]. If, by contrast, E is positively relevant to [S smokes], then

$$(15) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& K) > P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K).$$

We can break down the left-hand and right-hand side of (6) and (15) as follows:

$$(16) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& K) = \\ P(C_1 \mid [S \text{ gets cancer}] \& E \& K) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& C_1 \& K) + \\ P(\sim C_1 \mid [S \text{ gets cancer}] \& E \& K) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& \sim C_1 \& K),$$

$$(17) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K) = \\ P(C_1 \mid [S \text{ gets cancer}] \& K) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& C_1 \& K) + \\ P(\sim C_1 \mid [S \text{ gets cancer}] \& K) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& \sim C_1 \& K).$$

Plausibly,

$$(18) P(C_1 \mid [S \text{ gets cancer}] \& E \& K) > P(C_1 \mid [S \text{ gets cancer}] \& K)$$

and

$$(19) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& C_1 \& K) \geq P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& C_1 \& K).$$

(18) says that, given that S gets cancer, [(S smokes) would explain (S gets cancer) if (S smokes) were true] makes it more likely that there exists at least one instance of someone

¹¹ In my 2016b, I formalize the idea that explanatory connections mediate confirmation in terms of Bayesian networks.

getting cancer because of smoking.¹² (19) says that, if we know that [S gets cancer]&C₁&K, E does not make [S smokes] *less* likely. It follows from (18) and (19) that the first summand in (16) is greater than the first summand in (17).

However, this does not yet show that (15) is true. This is because

$$(20) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& \sim C_1 \& K) = 0,$$

and so the second summand in (16) equals 0, and hence is less than the second summand in (17). (20) is true because, if smoking *never* causes cancer, and S gets cancer, then it cannot be the case that S's smoking causes S's cancer. However, if it's the case that, were S to smoke and get cancer, S's smoking would be the cause of S's cancer, *and* it's the case that S gets cancer, then if S smokes, S's smoking must cause S's cancer. Hence, the only way for [S gets cancer]&E&~C₁ to be true is for it to be the case that S does not smoke.

Nevertheless, for (6) to be true the second summand in (17) would need to exactly equal the difference between the first summand in (16) and the first summand in (17). While this could be the case, there is no reason to expect it a priori. Hence, far from being a general truth, if (6) is true in this case it is only by fortuitous coincidence.

More importantly, the negative influence of E on [S smokes] sketched above is not the kind of influence that either proponents or opponents of inference to the best explanation have had in mind when disagreeing about whether explanation is relevant to confirmation. And if we build into K information that screens off this influence, then (15) is true. For example, imagine that we know that nothing besides smoking will give S cancer (and that S will not get cancer for no reason). In this case $P(\sim C_1 \mid [S \text{ gets cancer}] \& K) = 0$ – if the only way for S to get cancer is from smoking, then if S gets cancer it is because of S's smoking, and so C₁ is true. Hence the second summand in both (16) and (17) is 0, and the dominance of (16)'s first summand over (17)'s first summand is sufficient for it to be the case that

$$(15) P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& E \& K) > P([S \text{ smokes}] \mid [S \text{ gets cancer}] \& K).$$

I have argued in this paper that R&S's thesis that the explanatory hypothesis [if H and O were true, H would explain O] is irrelevant to confirmation is not true in the kind of case they discuss, at least once we fix our background knowledge so as to screen off irrelevant

¹² In an earlier version of this paper, I claimed that E entails C₁. However, an anonymous reviewer pointed out to me that this is not true, even holding fixed the above background knowledge. For example, suppose that smoking never has and never will cause cancer, so that C₁ is false. S, for his part, does not smoke. Nevertheless, because of S's unique physiology and the chemical properties of tobacco, it is true that *were* S to smoke, his smoking would cause him to have cancer. In this case E is true, but C₁ remains false. However, inasmuch as a scenario like this in which E is true and C₁ is false is incredibly unlikely, it remains extremely plausible that E *confirms* C₁, even if it does not entail it.

information. I have also shown that when we move to more tractable propositions describing explanatory connections, such as ones about general causal links between smoking and cancer, not only are these relevant to confirmation, they actually mediate the connection between observations and theories: the observation confirms the theory (and vice-versa) only insofar as we have evidence that the described explanatory connection exists. Explanation not only *adds* confirmation; it *guides* confirmation.

References

- Climenhaga, Nevin. 2016a. "Inference to the Best Explanation Made Incoherent." Unpublished manuscript.
- Climenhaga, Nevin. 2016b. "The Structure of Epistemic Probabilities." Unpublished manuscript.
- Douven, Igor. 2013. "Inference to the Best Explanation, Dutch books, and Inaccuracy Minimisation." *The Philosophical Quarterly* 63: 428-444.
- Eells, Ellery, & Branden Fitelson. 2002. "Symmetries and Asymmetries in Evidential Support." *Philosophical Studies* 107: 129–142.
- Keynes, John Maynard. 1921. *A Treatise on Probability*. London: Macmillan and Co.
- McCain, Kevin and Ted Poston. 2014. "Why Explanatoriness Is Evidentially Relevant." *Thought* 3: 145–53.
- Roche, William and Elliot Sober. 2013. "Explanatoriness is Evidentially Irrelevant, or Inference to the Best Explanation meets Bayesian Confirmation Theory." *Analysis* 73: 659-668.
- Roche, William and Elliot Sober. 2014. "Explanatoriness and Evidence: A Reply to McCain and Poston." *Thought* 3: 193-199.
- Sober, Elliott. 2001. "Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause." *British Journal for the Philosophy of Science* 52: 331-346.
- Steel, Daniel. 2003. "Making Time Stand Still: A Response to Sober's Counter-Example to the Principle of the Common Cause." *British Journal for the Philosophy of Science* 54: 309-317.
- van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.