Evidence and the Openness of Knowledge*

Assaf Sharon & Levi Spectre

1 Introduction

This article is driven by a simple idea: in the analysis of knowledge, the logic of evidence must have a pivotal role. A proper account of knowledge, in other words, must be compatible with basic facts about the relation of evidential support. Undeniable as this idea may seem, even among contemporary epistemologists who address evidence in their theories little attention has been given to the actual workings of evidence. Founding the theory of knowledge upon the proper analysis of evidence, we argue, has ramifications for epistemology that are wide-ranging as they are fundamental. Specifically, we argue that, since the relation of evidential support is not closed under known entailment, empirical knowledge is also open.[1]

Our argument proceeds in the following form. We inspect the most promising argument in favor of epistemic closure and argue that, in face of a proper understanding of empirical knowledge and its relation to evidence, it fails. Reflecting on this failure and on the logic of evidence to which it is traced, we present an argument for epistemic openness. In contrast to common opinion, we argue, it is not an externalist, "belief-sensitivity" view

* Versions of this paper were presented at the 2007-8 NYU/Columbia Graduate Student Philosophy Conference, The 12th Conference of the New Israeli Philosophical Association at the Open University of Israel, and The Istanbul Bogazici University "Knowledge, Evidence and Probability" Conference; at the Philosophy Departments' colloquia of Bar-Ilan University, The Hebrew University, Ben Gurion University, and The Central European University; and at Stockholm University's "Logic and Language", Stanford's 2008 advanced modal logic "Logic and Rational Agency", and The University of Arizona's philosophy seminars. We wish to thank the participants in those forums for very helpful comments, in particular: Per Martin-Löf, Dag Prawitz, Johan Van Benthem and Timothy Williamson. For many discussions, comments, and suggestions, thanks to: Hagit Benbaji, David Enoch, Dan Halliday, Mattias Högsröm, Mikael Janvid, Maria Lasonen-Aarnio, Krista Lawlor, Gilad Liberman, Ofra Magidor, Itamar Pitowsky, Ruth Weintraub, Åsa Wikforss, and Jonathan Yaari. Special thanks to Stewart Cohen, John Hawthorne, Karl Karlander, Avishai Margalit, Jim Pryor, and Peter Pagin for many engaging conversations that significantly influenced our thinking about the issues of this paper.

[1] The term "open knowledge" was first coined (as far as we know) by Nozick (1981: 208) and refers to the view to which he had subscribed, namely, that knowledge is not closed under known entailment.

that is most congenial to epistemic openness, rather it is the dependence of knowledge on evidence that best motivates this position.

Without attempting to provide a full-fledged theory of evidence, we show that on the modest assumption that evidence cannot support both a proposition and its negation, or, alternatively, that information that reduces the probability of a proposition cannot constitute evidence for its truth, the relation of evidential support is not closed under known entailment. Regardless of whether the proper account of evidence is probabilistic or not, the evidence-for relation is deductively open. We then turn to argue that given a minimal dependence of knowledge of empirical truths on evidence, there is compelling reason to reject a number of intuitively appealing epistemic principles, including not only the principle of epistemic closure, but also other, weaker principles. We present a number of significant benefits of this position, namely, offering a unified solution to a range of central epistemological puzzles as well as an account of their force and resilience to other attempted solutions.

Another way of stating the objective of this article is to set a challenge for epistemic closure: if the openness of evidence can be established (probabilistically as well as non-probabilistically), and some kind of dependence of empirical knowledge on evidence is unavoidable, as we argue, how can knowledge be closed?

2 Closure: Deniers and Defenders

You look at your watch and see that it reads “3:00”. Assuming that the time actually is 3:00 o’clock and that all other things are normal, you now know that the time is 3:00. By trivial reflection you also know that *if the time is 3:00 o’clock, then if your watch reads “3:00”, it is showing the correct time*. Do you know that *if your watch reads “3:00”, it is* *showing the correct time*? Do you know, just by looking at it, that *even if the watch has stopped, it is showing the correct time*?

Intuitively, it does not seem that you do. Perhaps you already knew beforehand – relying on other sources – that your watch is now accurate. But if you don't, it does not seem like the kind of thing that can be known on the basis of the fact that the watch shows "3:00". And yet, epistemological orthodoxy says that you do know this. Since knowledge is closed under known entailment, the claim goes, a belief properly derived from a known proposition is itself known. Having derived the belief that *if my watch reads "3:00", it is showing the correct time*, from your knowledge that *the time is 3:00 o'clock*, you know this conditional is true. Knowing that your watch shows "3:00", you can derive the consequent of the conditional and hence know that *your watch is showing the correct time*. Roomer

Why hold fast to this counter-intuitive conclusion? The answer, as in many similar cases,[2] is the principle of epistemic closure. This widely accepted principle, in one of its better formulations, states that:

(CP) Necessarily, if S knows *p*, competently deduces *q* [from *p*], and thereby comes to believe *q*, while retaining knowledge of *p* throughout, then S knows *q*.

Hawthorne (2004a: 34)

The watch reading example, however, brings out not only the counter-intuitive consequences of CP, but also theoretical reasons for thinking that it fails in cases of this sort. To put it succinctly, the reason we tend to deny the status of knowledge to the conclusions of such inferences is that they lack evidential support. In what follows we wish to elaborate

[2] See Vogel (1990), (2000), (2007) and Cohen (2002), (2005).

and support this claim by analyzing this example in detail. This analysis will, in turn, serve our more ambitious attempt to motivate knowledge openness and lay bare its benefits.

2.1 Evidence and Probabilities

Your reading of the watch provides you with evidence in virtue of which your belief that it is 3:00 o'clock counts as knowledge. But what does it mean that reading "3:00" off the watch is evidence for your belief?[3] One common rendering of this relation is in terms of conditional probability. The probability that the time is 3:00 (*p*) given the evidence that the watch reads "3:00" (*e*) is greater than the probability of *p* without this evidence.[4]

(1) $Pr(p|e) > Pr(p)$

Even if it cannot be accepted as a definition of evidential support, it seems that any account of evidence should grant the following criterion

(EC) Necessarily, if *e* evidentially supports *p*, then the probability of *p* given *e* is not lower than the prior probability of *p*.[5,6]

---

[3] If you think the evidence in such cases is different, e.g. the evidence is not the appearance of the watch but that the time is three, simply adjust the example. As we show later, the argument relies on purely formal features of the relation between evidence and that which it supports. In fact, even if the evidence is that the time is three, since you acquire this evidence by looking at your watch it still seems odd that you could learn on the basis of such evidence that your watch hasn't stopped a half hour ago.

[4] Together with *a priori* propositions, we do treat necessary contingent propositions, e.g. "I exist" and, if Williamson is right, "there is at least one believer" (Williamson 1986) etc., as having probability 1. We also assume that evidence has a probability 1 and that some knowledge has less than probability 1, i.e. that some knowledge is not evidence. In particular, although in some cases it seems plausible to associate probability 1 to known propositions, we do not accept that this is the case across the board. Some of the arguments proposed below can be reformulated as a challenge to those who, following Williamson (2000: 184-237), view knowledge as always having probability 1. Besides our problems with regard to the way Williamson characterizes prior probabilities, there are also epistemological problems with this account (see fn. 52 below for some more details and our MS).

[5] Note that this criterion does not require raising of probabilities.

[6] Some may be worried that not all evidence is propositional, that experiences, for instance, such as the experience of a blue patch in one's visual field, may be evidence for one that there is something blue in the vicinity. If you have such worries, take as the relata figuring in EC (and the other evidence principles below) the proposition that *S is experiencing a blue patch in his field of vision*. We propose this measure only in order to sidestep this thorny issue.

$E(e,p) \Rightarrow (Pr(p|e) \geq Pr(p))$ [$E(e,p)$ = *e* evidentially supports *p*][7]

How does your situation vis-à-vis the accuracy of your watch fare with respect to this criterion? From *p* it follows that *if the watch shows "3:00", then the watch is showing the correct time*. The antecedent of this conditional is just *e*, which trivially implies the consequent (call it "*c*" which is just the conjunction of *e* and *p*) given that the time is indeed 3:00. In other words, it follows from *p*, that *if e, then p-and-e*. Hence one can know a priori, by mere reflection, that:

(2) $p \Rightarrow (e \supset c)$

It follows by closure that you can know

(3) $e \supset c$

But do you have evidence for (3)? Presumably, if you do, it must be the evidence that facilitated knowledge of *p* in the first place, namely, *e*. But if it is a necessary condition on evidence that it not decrease the probability of that for which it is evidence, then *e* does not provide (3) with evidential support. This is because the conditional probability of (3) on *e* is not greater than the probability of (3). In fact, since *e* verifies the antecedent of (3), it *lowers* the probability that this implication is true.[8] The truth of *e* excludes all the cases in which (3) is true in virtue of the falsity of its antecedent. So in fact,

---

[7] "$\Rightarrow$" denotes strict implication (usually either logical implication or some other sort of *a priori* implication). Epistemologists tend to be quite relaxed in their usage of the terms "implies" and "entails," usually not taking much care to distinguish strict implication, or entailment, from material implication. Following this usage, let us note however that all implications referred to in this article are necessary, or *a priori* knowable, strict implications. Similar remarks are in order with respect to equivalence by which we mean not logical equivalence, but conceptual, or *a priori*, equivalence, symbolized by "$\Leftrightarrow$". We use "$\supset$" for material implication. Other symbols are standard unless explicitly defined.

[8] Let us show that $Pr(e \supset c|e) < Pr(e \supset c)$. First, $Pr(\sim(e \supset c)|e) \geq Pr(\sim(e \supset c))$, since:

1. $Pr(\sim (e \supset c)|e) = Pr((e \wedge \sim c) \wedge e)/Pr(e) = Pr(e \wedge \sim c)/Pr(e)$
2. $Pr(\sim(e \supset c)) = Pr(e \wedge \sim c)$

Assuming that $0 < Pr(e) < 1$, $(Pr(e \wedge \sim c)/Pr(e)) > Pr(e \wedge \sim c)$, so $Pr(e \supset c|e) \leq Pr(e \supset c)$.

Second, assuming as we are throughout that the probability of $e \wedge \sim c$ is not zero (i.e. that c is known fallibly and is not a necessary truth), the right side of the inequation is greater than the left. Thus: $Pr(e \supset c|e) < Pr(e \supset c)$.

(4) $\Pr(e \supset c | e) < \Pr(e \supset c)$

From EC and (4) it follows, as several theorists have observed,[9] that *e* is not evidence for $e \supset c$. It is this lack of evidence, we argue, that explains why, although properly derived from known premises, (3) is not known.[10]

We shall consider alternative analyses and possible replies to this argument below (section 3). First, however, let us present the key claim. We believe that the considerations invoked by this argument explain a host of other examples often proposed as challenges to the validity of epistemic closure. It follows from something's being a zebra that it is not a mule disguised to look like a zebra. And yet, seeing a zebra-looking animal in the pen, although providing one with evidence that there is a zebra in the pen, does not provide any evidence that the animal is not a disguised mule. In fact, that there is a zebra-looking animal in the vicinity is, at least to some extent, an indication that there is a zebra-looking disguised mule in the area. Memory of having parked one's car in the driveway ten minutes ago evidentially supports the belief that one's car is in the driveway. It provides no evidential support for the entailed belief that one's car has not been stolen in the last ten minutes. In standard conditions, seeing what appear to be one's hands is evidence that one has hands; it is not evidence that one is not a bodiless brain in a vat. Examples of this sort abound.[11] What is common to all, we suggest, is the failure of the evidence for the originally known proposition to support the inferred proposition. Lacking evidential support, it seems, empirical beliefs of this sort do not qualify as knowledge.

---

[9] See especially White (2006), Cohen (2002), (2005) and Hawthorne (2004) for proofs related to similar cases. Probabilistic proofs that the *evidence for* relation is not transitive date back to Carnap (1950) (at least). See section 3.6 below.

[10] The problem we discus here is similar to the problem discussed in Hawthorne (2004a: 73-8) and in Cohen (2005). As will become apparent, the core issue we believe relates to the failure of evidence closure and differs significantly with respect to the analysis and solution of these problems. We are, nevertheless, indebted to their groundbreaking work on these issues.

[11] The examples are to be found in Dretske (1970) and Vogel (1990), respectively. The last is a variation on Moore's (1959) proof of an external world.

The same idea accounts for a number of other unhappy consequences of epistemic closure. Having proper evidence that *p* is true can allow one to know *p*, but not that the means by which the evidence was acquired are reliable, or that evidence against *p* is misleading.[1] Proper evidence can warrant one in believing that *p*, but does not supply one with reasons for believing that this evidence is not misleading.

Admittedly, denying the status of knowledge to properly inferred beliefs exemplified in these cases has its cost, namely, the rejection of the intuitive and extremely popular principle of epistemic closure. In what follows we shall look at the strongest argument against the rejection of closure, specifically the arguments presented by John Hawthorne (2004, 2005). We will show not only that these arguments fail to protect closure from its deniers, but that, moreover, careful analysis of its premises provides substantive reasons for *rejecting* closure.

2.2 Costs of Closure Rejection

Advocating knowledge openness, the denier of closure stands in opposition to two main kinds of closure endorsers – skeptics and optimists. Skeptics often argue that since one does not know some proposition *q* that is known to follow from some other proposition *p*, one does not know *p*. The optimist claims that both *p* and *q* are known, either simpliciter (e.g. Moore), or with reference to different contexts of ascription (contextualists), or to dif-

---

[12] These are, respectively, versions of what has come to be known as the "easy knowledge" problem (Cohen 2002, 2005) and Kripke's dogmatism puzzle (forthcoming). The following sentence in the text presents an instance of the phenomenon of epistemic ascent. For a focused discussion of all three issues and the how they are related see our 2010 and MS. Notice that our formulation of the problem is general in that it does not rely on the intuition that knowledge is gained too easily (Cohen 2002, 2005) nor on the intuitive oddity of bootstrapping oneself into knowledge of the reliability of one's sources (Vogel 2000, 2007). As will become evident in what follows we rely solely on structural features of evidence and the principles governing the relation of evidential support. Thus, contrary to what some have alleged, the denial of closure is motivated not merely by the desire to avoid Cartesian skepticism. Epistemic closure is implicated in many epistemic puzzles, including, in addition to those already mentioned, the lottery paradox (see Vogel 1990, Hawthorne 2004a), some of the semantic self-knowledge puzzles and probably some other problems that are less central in current writings. There is a further brief discussion of this issue below.

ferent practical environments the subject is in (subject-sensitive invariantists).[13] What both skeptics and optimists agree on is that if *q* is properly derived from a known proposition, *q* is also known. It is this contention that the advocate of the openness of knowledge rejects.

To defend closure against examples advertised by its deniers, Hawthorne argues that, interpreted in the way closure deniers would have us interpret them, these examples conflict with other, more basic, epistemic principles. The advocate of knowledge openness, he claims, is forced to reject these highly compelling principles. In other words, to deny closure on the basis of these examples is tantamount to denying a bunch of weaker principles as well. Thus, if his arguments are cogent, Hawthorne manages to significantly raise the price of knowledge openness.

Nevertheless, closer inspection of Hawthorne's arguments, we believe, shows that this is not in fact the case. We will argue that, although one way of responding to Hawthorne's argument does indeed raise the cost of denying closure, there is a better response that does not involve such costs. The same reasons that motivate the denial of closure tell against the weaker principles Hawthorne puts to task. Thus, anyone rejecting closure for the right reasons will also reject the weaker principles on which Hawthorne's argument relies. Let us look at Hawthorne's arguments.[14]

2.3 Hawthorne's Arguments

Hawthorne offers what is perhaps not only the cleverest but also the strongest argument in defense of epistemic closure. Exposing the deeper connections and further commitments of

---

[13] This is not to say that contextualism or subject sensitive invariantism entail closure of knowledge. Both are compatible with open-knowledge and can be employed to explain certain epistemic phenomena.

[14] Hawthorne (2004a) and (2005). We do not present or attempt to answer all of Hawthorne's arguments in support of closure, only the ones that we take to be most forceful and to pose the greatest challenge for epistemic openness of the kind we are advocating.

closure denial, Hawthorne's argument helps articulate what we take to be the proper grounds for epistemic openness.

The following are Hawthorne's weaker principles, the principles he would have the closure denier hold on to:

Equivalence (EQ):	Necessarily, if S knows that *p*, and S knows that *p* is *a priori* equivalent (or logically equivalent) to *q*, then S knows that *q*.
Addition (AD):	Necessarily, if S knows that *p*, then by competently inferring *p or q* from *p*, S thereby knows *p or q*.[15]
Distribution (DIS):	Necessarily, if S knows that *p and q*, S knows *p* and S knows *q*.

Hawthorne (2004a: 41)

Indeed, all three principles seem highly plausible. To see how they lead to the same conclusion as CP we shall, following Hawthorne, look at Fred Dretske's well known zebra case (1970: 1015-6). Seeing a zebra-looking animal in the pen labeled "Zebra" one knows that the animal in the pen is a zebra (call this proposition *Z*). But, presumably, one does not know that this animal is not a mule disguised to look like a zebra (*~DM* for short). So according to the closure denier (in this case, Dretske) one knows *Z*, and knows that *~DM* follows from *Z*, but does not know *~DM*.

Now Hawthorne's argument runs as follows. We are assuming that S knows that:

(5) Z [assumption]

(6) $Z \Rightarrow \sim DM$ [assumption]

By AD, S can infer:

(7) $Z \vee \sim DM$ [AD,5]

[15] Further clauses can be added to these principles, see Hawthorne (2004a: 39), but for simplicity we omit them here. Nothing in our argument turns on this simplification.

Assuming that S is familiar with basic logical operations, she can know that:[16]

(8) $(Z \vee \sim DM) \Leftrightarrow \sim DM$ [PL,6][17]

By EQ it now follows that S knows that

(9) $\sim DM$ [EQ,7,8]

Thus to avoid the implausible consequences of the example, closure deniers must also deny AD or EQ (or both). Since other counter-examples to closure share the form of this one, the same problem will arise for them as well.[18]

Hawthorne also employs a parallel argument using DIS. Assuming as before that S knows that:

(5) Z [assumption]

(6) $Z \Rightarrow \sim DM$ [assumption]

Again, familiarity with basic logical operations enables S to know that:

(10) $Z \Leftrightarrow (Z \wedge \sim DM)$ [PL,6]

By EQ, S knows:

(11) $Z \wedge \sim DM$ [EQ,5,10]

DIS entails that, knowing (11), S is in a position to know:

(12) $\sim DM$ [DIS,11]

Again, the conclusion closure deniers aim to avoid is reached by principles weaker than closure. If these consequences mandate rejection of closure, they should also warrant rejection of these weaker principles. Hawthorne's argument successfully shows that to avoid

---

[16] Hawthorne (2004a: 41, note 99) notes that, strictly speaking, that a thing is a zebra does not logically imply that it is not a painted mule. Recent reports indicate that zebras may also be mules or at least horses. But let us not allow the facts ruin a good example.

[17] "PL" will stand for basic operations of propositional logic.

[18] The examples we are considering employ only single premise closure. "Multi Premise Closure" is questioned even by Hawthorne though he does maintain that there are *some* prospects for maintaining it (2004a: 186). We show elsewhere why this is extremely problematic (MS).

the undesirable consequence of closure either EQ or both AD and DIS are to be jettisoned. He concludes that closure is pretty much "non-negotiable."[19]

3 The Openness of Evidence

Earlier we saw that it is doubtful that by looking at one's watch one can know that it is accurate. After all, the evidence one has gained counts against this conclusion. Hawthorne, however, contends that the natural way to react to this, namely the rejection of epistemic closure, has significant costs. How are we to manage this tension? In this section we will argue that the same reasons that ought to motivate the rejection of closure, i.e. the fact that evidence is open, also provide reasons to reject Hawthorne's conclusion. This argument will also be used to show that trying to evade the argument we will subsequently present against closure cannot be answered by giving up on a probabilistic understanding of the *evidence for* relation. We will end this section by addressing other possible answers to the argument against closure. We will conclude that since on any account of evidence, probabilistic or otherwise, evidence is open knowledge is open as well.

3.1 Evidential Principles

Imagine you are looking for zebra-look-alike mules. Where would you look? It would be natural for you to do so among zebra-looking animals. Admittedly, zebra-looking mules would be hard to find, but if there is any chance of finding some (at least the ones which are well disguised) you had better search among zebra-looking animals. Seeking them among the elephant-looking animals, or the banana-looking objects holds little promise of

---

[19] Hawthorne, (2004a: 112). Others have simply called the principle "intuitive closure" (Williamson 2000: 117), claimed that rejecting closure is "intuitively bizarre" (DeRose 1995: 201), or "one of the least plausible ideas to gain currency in epistemology in recent years" (Feldman 1999: 95), and that closure is "something like an axiom about knowledge" (Cohen 2005: 312).

success. Your chances of encountering a zebra-looking mule are slim. But the probability that an object encountered is *not* a zebra-looking mule is even lower when a zebra-looking animal (say, a zebra) is visually observed. Although it does not constitute strong evidence, a zebra-looking animal gives some support to the proposition that a given object is a zebra-looking mule. In other words, the presence of a zebra-looking animal raises the probability that a mule disguised to look like a zebra is present.[20]

Now in the normal case, when one sees a zebra-looking animal, one has evidence that the animal is a zebra. But anyone who knows that a zebra is not a mule, must realize that at the same time that one gains evidence for Z in this way, one *loses* evidence for $\sim DM$. Denying this quickly gets one into serious trouble in trying to provide a plausible account of evidential support. The argument below shows why.

The following principle, it seems, must be a part of any plausible theory of evidential confirmation:

Consistency of Evidence (CS): If e evidentially supports h, e does not evidentially support the negation of h.

Let us now examine other seemingly plausible evidential principles analogous to Hawthorne's epistemic ones.[21]

Evidence addition (EAD): If e evidentially confirms h_1, e evidentially supports h_1 *or* h_2.

Evidence equivalence (EEQ): If e evidentially supports h_1, and h_2 is logically (or *a priori*) equivalent to h_1, e evidentially supports h_2.

Like their epistemic counterparts these principles enjoy a high degree of intuitive appeal. The first principle, CS, expresses the simple idea that if something is to count as evidence for some theory, hypothesis, proposition or what have you, it cannot also support its

[20] There is no essential probabilistic point here. All that we are claiming is that some measure of support is given to *DM* by a visual observation of a zebra looking animal.

[21] We will later discuss the relation between these principles and their epistemic analogues.

negation. Or, in other words, that a proposition supporting both a hypothesis and its negation, does not constitute evidence for either.

The principles of EAD and EEQ stem from the idea that the *evidence for* relation is closed under certain logical operations. Addition captures the idea that adding disjuncts to a supported hypothesis does not undermine the degree of support. Equivalence, on the other hand, expresses the idea that "confirmation of a hypothesis is independent of the way in which it is formulated."[22] The truth-values of logically equivalent hypotheses stand or fall together, so equivalent hypotheses must also be confirmed and disconfirmed together.[2] Notice that the justification for these principles is the same as that proposed for their epistemic counterparts.

3.2 Evidence and Underdetermination

Although the principles presented in the previous section all appear plausible, their conjunction with (even a particularly weak version of) the thesis of underdetermination of theory by evidence, leads to a contradiction. Strong underdetermination is the contentious claim that all possible evidence cannot fully determine the choice between (some) mutually incompatible theories. Weak underdetermination (henceforth: UD), however, which is all we shall assume, states merely that, at least insofar as actual evidence goes, there can be two (or more) inconsistent theories supported by the same body of evidence.[24]

---

[22] Hempel (1965: 13). See also Carnap (1950: 474). Notable dissenters are Scheffler and Goodman, see, Scheffler and Goodman (1972: 78), Scheffler (1963: 289), Goodman (1955: 71-2).

[23] In terms of a coarse-grained possible world semantics, we might say that any evidence that the actual world is one of the h_1-worlds (the possible worlds in which h_1 is true) is also evidence that the actual world is an h_2-world, since in those terms the sentences express the same proposition. EEQ is justified by the claim that any evidence that the actual world is one of the h_1-worlds is also evidence that the world is an h_2-world if "h_1" and "h_2" are true in the same worlds.

[24] For the argument below all that is needed is that there are some cases of (inductive) underdetermination.

To fix ideas let us focus on a specific example – the competing interpretations of formulas of quantum mechanics. The two leading interpretations of quantum theory are, apparently, compatible with all (possible) observations. And yet, since one, the Copenhagen interpretation, entails that every particle has a momentum and the other, the Bohmian interpretation, implies that particles have no momentum, the two are mutually incompatible. Presumably, the evidence we have supports both interpretations. Thus, there is evidence supporting the Copenhagen interpretation (CQM). By EAD it follows that this evidence also supports: *the Copenhagen interpretation is true or the Bohmian interpretation is false* (CQM $\vee$ ~BQM). But this is equivalent to: *it is not the case that the Copenhagen interpretation is false and the Bohmian is true* (~(~CQM$\wedge$BQM)), and so the evidence supports this latter proposition as well. Now since the truth of one interpretation entails the falsity of the other, *the Bohmian interpretation is true* (BQM) is equivalent to *the Bohmian interpretation is true and the Copenhagen interpretation is false* (~CQM$\wedge$BQM). Thus by evidence equivalence, EEQ, the evidence supports the claim that *it is not the case that the Bohmian interpretation is true*. It follows from CS that the evidence supports neither the Bohmian interpretation, nor its negation – in contradiction to what we have assumed.[25]

The problem generalizes. UD entails that given a finite set of evidence propositions *e*, this evidence can equally support two incompatible theories, T_1 and T_2. Thus, T_1 implies *not-*T_2, and T_2 implies *not-*T_1. Let us state this more formally as follows:

(13) $E(e, T_1) \wedge E(e, T_2) \wedge (T_1 \Rightarrow \sim T_2)$ [UD][26]

---

[25] In social sciences instances of such underdetermining evidence seem to be even more prevalent and easier to describe. Take, for example, the debate between the “directional” and the “proximity” models of special representation of voting preferences. In 1999 two scholars claimed that “the existing data contain insufficient information with which to distinguish the two theories.” (Lewis and King 1999). The claim was repeated in 2006 (by Van Houweling, Tomz and Sniderman,). This conclusion may certainly be debated, but for the purposes of this argument the possibility of its truth suffices.

[26] The implication here, as well as the equivalence in (16), may be logical, since, presumably, T_1 contains some proposition the negation of which figures in T_2. For our purposes, as noted in note ##, suffice it that the implication and equivalence are *a priori*.

EAD entails the following:

(14) $E(e,T_1) \Rightarrow E(e,T_1 \vee \sim T_2)$ [EAD,13]

which entails:

(15) $E(e, T_1 \vee \sim T_2)$ [MP,13,14]

Now, since T_1 entails the negation of T_2, the following equivalence is true:[27]

(16) $(T_1 \vee \sim T_2) \Leftrightarrow \sim T_2$

It thus follows from (15) that:

(17) $E(e, \sim T_2)$ [EQ,15,16]

But given the principle of consistency CS, this entails that *e* does not evidentially support T_2.

(18) $\sim E(e, T_2)$ [CS,13,17]

and (18) contradicts (13).

We must conclude from this argument that the principles alone, with no appeal to probabilistic interpretation, are inconsistent. Moreover since principles employed in this argument are structurally similar to Hawthorne's principles, it raises the suspicion that something has gone wrong with the latter as with the former. In both arguments seemingly non-contentious premises entail unpalatable conclusions.

3.5 Equivalence, Consistency and Addition

The equivalence of evidence is involved in one of philosophy's notorious paradoxes, namely Hempel's paradox of confirmation. In Hempel's argument the standard conception of evidential confirmation leads to apparently unreasonable results. Specifically, Hempel

[27] The proof is straightforward. First, from left to right: Given that $T_1 \Rightarrow \sim T_2$ and assuming $T_1 \vee \sim T_2$, either T_1 is true, in which case so is $\sim T_2$ (by the implication), or $\sim T_2$ is true. So $T_1 \vee \sim T_2$ implies $\sim T_2$. Now from right to left, $\sim T_2$ clearly implies the disjunction $T_1 \vee \sim T_2$. Hence, $T_1 \vee \sim T_2 \Leftrightarrow \sim T_2$.

showed that coupled with the Nicod principle,[28] the principle of equivalence leads to the conclusion that pink stockings evidentially confirm the claim that all ravens are black. The present argument shows that a black raven equally supports the claim that not all ravens are black and is thus no evidence at all. Given even weak UD, no matter what conception of confirmation it is coupled with, whether Nicod's or some other conception, the EEQ and EAD principles lead to paradox.

It may be suggested at this point that neither EEQ nor EAD are to be identified as the culprit. It is rather UD that is incompatible with CS. To be sure, a similar point was already made by Hempel:

> A finite set of measurements concerning the changes of one physical magnitude, *x*, associated with those of another, *y*, may conform to and thus be said to confirm, several different hypotheses as to the particular mathematical function in terms of which the relationship of *x* and *y* can be expressed; but such hypotheses are incompatible because to at least one value of *x*, they will assign different values for *y*. Hempel (1965: 33)

Unlike Hempel, Carnap was less reluctant to endorse this claim and follow it to its full consequences – the rejection of consistency (1950: 474-6). Thus, it may be suggested, it is UD that is incompatible with consistency, not EEQ.

But the consistency principle that Hempel and Carnap had in mind is significantly stronger than CS:

(CS*) If *e* evidentially confirms h_1, *e* does not evidentially confirm an *inconsistent* hypothesis h_2.

Surely, this principle is not compatible with UD, and given the pervasiveness of under-determination it must be rejected. CS, however, is not as disposable. How can a piece of evidence support some hypothesis if it supports its negation? CS, it seems, is a principle no plausible theory of confirmation can deny, for otherwise what is left of empirical refutation

[28] The Nicod principle: "For any object *a* and any properties *F* and *G*, the proposition that *a* has both *F* and *G* confirms the proposition that every *F* is *G*" Fitelson (2006: 95-6).

of a theory? Indeed Hempel and Carnap both embrace the following definition of "disconfirmation":

(DC) e disconfirms a hypothesis h if it confirms $\sim h$.[29]

Thus, if e confirms both h and its negation, it both confirms and disconfirms h and is thus evidence neither for h nor for *not-h*. Insisting that DC is true and CS false would lead to theoretical nihilism with regard to evidence.

So, although UD arguably entails that CS* is to be rejected, CS is not expendable for any proper theory of the relation between evidence and hypotheses. And yet, as the argument above shows, assuming so much as the weak UD and EAD, the EEQ principle conflicts with CS. Taken together these principles lead to the conclusion that e evidentially supports h and evidentially supports *not-h*. Since UD appears to be an undeniable reality, and since CS must be regarded non-negotiable, it seems that we must give up either EAD or EEQ.

In fact the same conclusion is reached even without CS. Suppose our evidence consists of two atomic proposition a and b, and our hypothesis consists of three independent atomic propositions a, b and c. Thus both c and $\sim c$ are consistent with a and with b. Evidence $a \wedge b$, then, supports both $a \wedge b \wedge c$ and $a \wedge b \wedge \sim c$.[30] So we have the following:

(19) $\mathrm{E}(a \wedge b, a \wedge b \wedge c) \wedge \mathrm{E}(a \wedge b, a \wedge b \wedge \sim c)$

By EAD and the first conjunct of (19), we have,

(20) $\mathrm{E}(a \wedge b,(a \wedge b \wedge c) \vee \sim(a \wedge b \wedge \sim c))$

From (20) we derive by EEQ

[29] Hempel (1965: 37). Cf. Carnap (1950: 479).

[30] This is easy to demonstrate within a probabilistic framework for evidence. Given the assumption that evidence just is the raising of probabilities, both hypotheses' probabilities are raised by $a \wedge b$. (See proof of lemma in note 34 below). But if we have a long conjunction with only the difference of two propositions, it seems safe to say that even without appeal to probabilities the evidence supports two incompatible conjunctions.

(21) $E(a \wedge b, \sim(a \wedge b \wedge \sim c))$.

Another application of equivalence will give us:

(22) $E(a \wedge b, (\sim a \vee \sim b \vee c))$

Since *a, b* and *c* are assumed to be independent atomic propositions, (22) is absurd.[3] We conclude then, that UD, EAD and EEQ, entail unacceptable consequences.

3.4 Equivalence and Distribution

Consider the evidential argument analogous to Hawthorne's second argument in defense of epistemic closure employing the distribution principle. As in the previous case, we start with two underdetermined theories T_1 and T_2:

(23) $E(e,T_1) \wedge E(e,T_2) \wedge T_1 \Rightarrow \sim T_2$ [assumption]

By EEQ, we have:

(24) $E(e,T_1 \wedge \sim T_2)$ [EEQ,23]

But now we need to use a principle which can serve as the evidential counterpart in place of DIS:

(EDIS) If *e* is evidence for *p and q*, then *e* is evidence for *p* and *e* is evidence for *q*.

Given (24), EDIS gives us:

(25) $E(e,\sim T_2)$. [EDIS,24]

Yet having assumed that *e* supports T_2, we end up with the same contradiction we had before. So we need to give up either EDIS, EEQ, CS or (19) (aka UD). The last two we

[31] The absurdity is even more pronounced given the following assumption that one who accept EAD would find hard to deny: if *e* is evidence for *p-or-q*, and *e* is not evidence for *q*, then *e* is evidence for *p*. Since $a \wedge b$ is evidence neither for *c*, nor for $\sim c$ (as we have supposed), this would entail, absurdly, that $a \wedge b$ is evidence for the falsity of atomic *a* or of atomic proposition *b*. Though at this stage the appeal to atomic propositions is not important. We would like to thank an anonymous referee for pointing this out.

have claimed are virtually undeniable (and giving up CS will not get us out of the woods anyway), so either EEQ or EDIS must be discarded.

The arguments from underdetermination show that equivalence cannot be maintained along with evidence addition or with evidence distribution. Thus, regardless of which of the principles should be rejected, since they rely on a conjunction of these principles, the evidential analogues of Hawthorne's arguments must fail. Moreover, since all these principles are weaker than evidence closure, evidence has been shown to be open on non-probabilistic grounds.

3.5 EAD, EDIS and the Logic of Evidence

But which are we to reject, EEQ or EAD and EDIS? Although rejecting EEQ provides a quick way out of the paradox of the ravens,[32] the plausibility of the equivalence of evidence advises against this strategy. Furthermore there are probabilistic considerations supporting the rejection of EAD and EDIS. Let us take them in turn. If T_1 and T_2 are incompatible theories and have the same initial probability, then this probability must be equal to or less than 0.5 (thus: $(T_1 \Rightarrow \sim T_2) \Rightarrow Pr(T_1)+Pr(T_2) \leq 1$). Let us suppose that each theory has an initially probability of 0.2. The probability of $\sim T_2$ is therefore 0.8. Now say we receive evidence e that supports both T_1 and T_2 to an equal degree (for simplicity). Let us assume that $Pr(T_2|e)=0.4$ and likewise for T_1. The initial probability of $(T_1 \vee \sim T_2)=Pr(T_1) + Pr(\sim T_2) - Pr(T_1 \wedge \sim T_2)$ which equals 0.8.[33] Now if e supports T_2 and T_1 equally (as we have assumed), then given e the probability of $T_1 \vee \sim T_2$ decreases to 0.6. The reason is simple,

---

[32] This is not entirely precise. EEQ relates to *a priori* equivalence whereas the paradox of confirmation turns on *logical* equivalence. Thus, although they are very similar and seem to be motivated by the same considerations, one may want to retain the logical equivalence of evidence while rejecting EEQ.

[33] Since the probability of T_2 was stipulated to be 0.2, the probability of $\sim T_2$ is 0.8, and the probability of $T_1 \wedge \sim T_2$ is just the probability of T_1.

since the probability of T_2 rises, the probability of $\sim T_2$ decreases and since from T_1 it follows that $\sim T_2$, the probability of the disjunction $T_1 \vee \sim T_2$ drops. Hence, *e* is not evidence for $T_1 \vee \sim T_2$ even though it is evidence for T_1 and evidence for T_2 (assuming, that is, that if *e* lowers the probability of a given theory *T*, it does not count as evidence in its favor).[34] So EAD fails.

A similar argument holds for EDIS. If we treat the evidence relation as a conditional probability relation, the probability of T_1 given *e* is the probability of *e* and T_1 divided by the probability of *e* ($Pr(T_1|e)=Pr(T_1 \wedge e)/Pr(e)$). Assume that the probability of T_1 is 0.2 and that of *e* is less than 1 (and more than 0). If T_1 entails *e*, then $Pr(T_1|e)>0.2$, and likewise for T_2.[35] This means that the probability of $\sim T_2$ given *e*, is less than 0.8. Now since T_1 entails $\sim T_2$, the prior probability of $T_1 \wedge \sim T_2$ is just the probability of T_1. And, as you may have figured out already, *e* must increase the probability of T_1 to no less a degree than it does that of $T_1 \wedge \sim T_2$. Although *e* raises the probability of $T_1 \wedge \sim T_2$, *e* lowers the probability of $\sim T_2$ (since it raises the probability of T_2). And so, unless we want to claim that *e* provides evidential support for a theory, proposition, or hypothesis, even though it lowers its probability, we must give up EDIS.[36]

This argument employs probabilities, but we have done so merely in order to to determine which principle should be rejected and haven shown non-probabilistically that evidence is non-transitive, to explain how this non-transitivity occurs. The principles we have

[34] This assumption is not to be confused with the stronger claim that evidence just is the increasing of probability. Here we are merely assuming that evidence cannot lower the probability that the proposition it supports is true (see EC above).

[35] This is shown by a simple application of Bayes' theorem.

[36] Starting with a stronger assumption than the one we have been employing, i.e. that evidence is defined by the raising of probabilities, that denial of EEQ conflicts with the Kolmogorov axioms. Let us assume that *p* and *q* are equivalent. Since *p* entails *q*, the probability of *p* cannot be greater than *q*. Likewise, since *q* entails p, *q*'s probability cannot be greater than *p*. Hence, their probability before and after the evidence is taken in must be the same, hence, if *e* is evidence for *p* it is evidence for *q*.

employed are not essentially probabilistic (although they can be given such an interpretation) and it is not difficult to construct a non-probabilistic argument to the effect that EAD and EDIS should be rejected. Let us briefly outline an explanation of what we take to be going on in these cases without employing probabilities. The basic idea is that before evidence comes in, a disjunctive proposition, *p or q*, can already be well supported (the notion of support need not be construed probabilistically). The evidence then introduced can count against the disjunction even if it lends support to one of the disjuncts. For example, since most objects are not disguised mules, the assumption that some (yet unperceived) object *is a zebra or a disguised mule* is highly plausible. Getting evidence that the object looks like a zebra makes it more likely (again – not necessarily in term of probability) that the object is a mule disguised to look like a zebra. It is the neglect of such possibilities that inclines us to accept EAD and EDIS. In the above argument, however, we have shown that on the mere assumption that evidence cannot support both a proposition and its negation, CS, these principles are invalid.

3.6 Carnap's Matrix

The formal considerations of the previous section are exemplified in the following scenario devised by Carnap (1950: 382-5).[37] The table below represents players in a game all of whom have equal chances of winning. The 'M's represent male contestants and the 'F's denote female contestants.

	Local	Out-of-towner
Junior	F F M	M M
Senior	M M	F F F

Assume the following principle about the relation of evidential support:

(ES) Necessarily, if the probability of *p* given *e* is greater than the prior probability of *p*, then *e* evidentially supports *p*. [38]

$\Pr(p|e) > \Pr(p) \Rightarrow E(e,p)$

Let *j* be the proposition that the winner is junior, *s* that the winner is a senior, *l* that the winner is local and *o* that the winner is from out of town. Let *f* represent the proposition

[37] Hempel (1965: 31-33) argues that if evidence is closed under strict implication, every proposition is evidence for any other. We present Carnap's example since, as will become evident, it relates directly to the principles that we have been concerned with. Namely, EAD and EDIS. Hempel's argument proceeds via what he labels the "converse consequence condition." But here is another more general way to proceed: Assuming that a proposition *e* is evidence for a proposition *h* iff the probability of *h* given *e* is higher than the probability of *h* ($E(e,h) =_{def} \Pr(h|e)>\Pr(h)$), let us first establish a lemma:
Lemma. For all empirical propositions *p* and *q*, if *p* entails *q*, then *q* is evidence for *p*.
Recall the definition of conditional probability: $\Pr(p|q)=\Pr(p\wedge q)/\Pr(q)$. Now let us assume (as is plausible if we are considering empirical matters) that $0<\Pr(q)<1$ and that *p* strictly implies *q*. It then follows that:

(1) $(p\Rightarrow q) \Rightarrow [\Pr(p)=\Pr(p\wedge q)]$ [Since p and p∧q are equivalent, Kolmogorov axioms]
(2) $\Pr(p\wedge q)/\Pr(q) > \Pr(p)$ [1,since $0<\Pr(q)<1$ and $\Pr(p\wedge q)=\Pr(p)$]
(3) $\Pr(p|q)>\Pr(p)$ [2,conditional probability]
(4) $E(q,p)$ [3,Evidence def.]

Now, the lemma entails:

(5) $\forall p\forall q E(p,p\wedge q)$ [Lemma]

Assuming for *reductio* that evidence is closed under (known) entailment we have:

(6) $\forall p\forall q\forall r(E(p,q)\wedge(q\Rightarrow r)) \Rightarrow E(p,r))$

But then since *q* follows from $p\wedge q$, we have the triviality result:

(7) $\forall p\forall q E(p,q)$ [5,6]

(7) is surely unacceptable, so one must either reject evidence closure or the proposed definition of evidence (or the Kolmogorov axioms). By relying on the weaker criterion EC rather than on the definition of evidence as the raising of probabilities, I avoid the rejection of the proposed definition of evidence as a reply to the argument against evidence closure. Another version of Hempel's argument – similar to the one presented here – can be found in Kaplan (1996: 45-56).

[38] Notice that this principle is stronger than the one our main argument utilizes, namely, EC. Giving up the probabilistic analysis of evidence expressed by ES, therefore, while useful against Carnap's argument will not resolve non-probabilistic argument nor cases such as the watch case.

that the winner is a female contestant and m that the winner is a male contestant. Now as can be seen in the table above:

$\Pr(s)=0.5$, $\Pr(s|f)=0.6$

$\Pr(o)=0.5$, $\Pr(o|f)=0.6$

So the probabilities of both s and o are raised given the information that the winner is female. Now let us look at the probability of the disjunction $o \vee s$:

Initially, $\Pr(o \vee s)=0.7$. Given f it becomes 0.6.

Thus, the probability that the winner is *either a senior or from out-of-town* decreases given that the winner is a female contestant. So by ES, although f is evidence for s and evidence for o, it is not evidence for *s or o*. If anything, the fact that the winner is a female is counter-evidence to the claim that the winner is either a foreigner or a senior.[39]

Now to EDIS. We have the following initial probability assignments:

$\Pr(l)=0.5$; $\Pr(j)=0.5$; and $\Pr(j \wedge l)=0.3$.

Given the evidence that the winner is a female the probabilities are:

$\Pr(l|f)=0.4$; $\Pr(j|f)=0.4$; and $\Pr(j \wedge l|f)=0.4$.

Thus while f raises the probability of the conjunction it lowers the probability of each conjunct.

Carnap's example shows that given ES, evidence is open. One can have evidence for p and have evidence for q while having no evidence for *either p or q*.

3.7 Open Knowledge

[39] Note that Carnap's example shows more than our argument requires, although by appeal to a stronger principle. It demonstrates that e can raise the probability of each of two propositions in isolation while lowering their disjunction. What we have relied on is merely that e can raise the probability of one of the disjuncts (p) while lowering probability the disjunction $p \vee q$. Notice also how strange it is that one could have evidence for p and evidence for q, yet lack evidence for *either p or q*. Asserting as much in ordinary conversation would seem very strange. This connects directly with DeRose's argument from abominable conjunctions to closed knowledge. See footnote 48 for further detail.

The preceding sections show, then, that the following principle must be rejected:

Closure of Evidence (CE): For all subjects S, evidence *e* and propositions *p* and *q*, if

a. S has evidence *e*,

b. S knows that S has *e*,

c. S knows *e* evidentially supports *p*,[40]

d. S knows that *p a priori* entails *q*,

then *e* evidentially supports *q* for S.

CE conflicts with the indispensable criterion:

(EC) For all evidence *e* and propositions *p*, if *e* is evidence for *p*, then *e* does not lower the probability that *p* is true.

And in addition, given the pervasiveness of underdetermination of theory by available evidence, CE (however weakened) cannot be maintained. It is of course possible to retain CE at the price of losing EC. But first, this must be regarded as a significant cost. It is hard to imagine a theory that captures a workable notion of evidence while violating EC. Second, the examples we have been considering all invoke a strong intuition that, regardless of EC, there is reason to doubt CE. Although one does have evidence that the time is three o'clock (i.e. the watch showing "3:00"), one does not have evidence for the truth of: *even if the watch has stopped, it is showing the correct time*, or that *if the watch shows "3:00", it is showing the correct time.* While one's memory of having parked the car is evidence that the car is in the driveway, one does not have evidence that the car has not been stolen. And likewise for many other cases. Finally, third, one would need to argue that either there are no cases of underdetermination as we have described them, or that CS is false; that is, that an item of evidence can support both a proposition and its negation.[41]

[40] We are here strengthening the antecedent by adding conditions b and c not to beg any question against a would be proponent of evidential closure.

[41] In the final analysis it seems that even CS is not essential to the argument for open evidence as shown by the argument above: (19)-(22).

Moreover, we have explained what the mistake is in endorsing CE, in both probabilistic and non-probabilistic terms. Here is another way of explaining what is going on in these cases this time in terms of possible worlds. Having evidence supporting a proposition *p* may be explicated as having reason to believe that the actual world is one of the *p*-worlds. Or if you prefer, evidence for *p* raises the probability that the actual world is a one of the worlds in which *p* is the case. Now if *q* (*a priori*) follows from *p*, then any world that is a *p*-world is also a *q*-world. So you might think that evidence that the actual world is a *p*-world must also be evidence that the actual world is a *q*-world. But, intuitive as it may be, this last step is wrong. If *p* implies *q*, then surely, a world that is a *p*-world is also a *q*-world. However, whether the evidence supporting the claim that a world is a *p*-world also supports the claim that it is a *q*-world depends on the relation between the purported evidence and *q*. Specifically, it depends on whether the evidence raises or lowers the probability that the world is a *q*-world (or in other words, whether it counts in favor or against the world being a *q*-world). Now, although the probability that the world is a *q*-world cannot be lower than the probability that it is a *p*-world, if the initial probability that the world is a *q*-world is higher than that of *p*, evidence that it is a *p*-world might lower it. This is why, as we have seen, a proposition that increases the probability of *p* can lower the probability of a proposition *q* implied by *p*.[42] This suggests that if, as we have urged, the idea that evidence must not lower the probability of the proposition which it supports is to be preserved, CE must be renounced. Evidence is not closed under entailment, known or unknown.

Now, for knowledge that depends on evidence, the following seems hard to deny for some cases at least:

[42] The same explanation, in essence, can account for cases involving disjunction and conjunction. It cannot be used to explain cases involving equivalence, which is one reason to think we have taken the correct track here in response to the cases we have been considering.

Evidence Dependence (ED): If an agent *S* does not know *p* and is not in a position to know that *p* just by believing that *p* at time t_o, between which time and some later time t_1 the only change in *S*'s evidential state is the addition of information that counts as a whole (together with background information) against *p*, then *S* does not know *p* at t_1.[43,44]

Assuming it is not necessary that one's total evidence entail every proposition one knows, it follows from ED that evidence-dependent knowledge is open. In fact ED need not hold generally. To undermine closure, suffice it that ED is true of one proposition derived from a known proposition but not supported by its evidence. Such propositions abound since for every proposition *p* based on evidence *e*, but not entailed by it, there is at least one proposition *q* deducible from *p* that is not supported by evidence *e*.[45] In fact, there are many such entailed propositions. We have an argument, then, for open knowledge from the dependence of knowledge on evidence. For suppose a subject S comes to gain evidence *e* for *p* (evidence which does not *a priori* entail *p*) on the basis of which S comes to know that *p* is true. Since for all *p* there will be propositions *q* which *a priori* follow

---

[43] The notion of *counting against* in ED can be interpreted probabilistically, or, in light of the forgoing discussion, non-probabilistically. The probabilistic reading of ED can be objectivist or subjectivist with respect to the likelihoods (the probability of the evidence given the hypothesis).

[44] Note that we do not regard ED as a principle. Nevertheless, it becomes evident that this claims holds for many cases (or at least some) when the background assumptions are made explicit, such as: S has not corrected her reasoning, received the kind of evidence that inspires her to realize that she has made a mistake, or remember that she has evidence she completely forgot about, etc.

[45] Proof: let us take as our *q* proposition the proposition $\sim(e\wedge\sim p)$ entailed by *p*. While raising the probability of *p*, *e* lowers the probability of this proposition. On standard assumptions, since *e* does not entail *p*, the probability of $e\wedge\sim p$ is less than 1. By the definition of conditional probability:

(1) $Pr(e\wedge\sim p|e)=Pr(e\wedge\sim p\wedge e)/Pr(e)=Pr(e\wedge\sim p)/Pr(e)$

Assuming that $0<Pr(e)<1$ (as we must), we have:

(2) $Pr(e\wedge\sim p|e) > Pr(e\wedge\sim p)$

Thus:

(3) $Pr(\sim (e\wedge\sim p)|e) < Pr(\sim (e\wedge\sim p))$

By EC, *e* supports *p*, but not a proposition *a priori* known to be entailed by it.

Note that this is not peculiars to *q* propositions the negations of which entail the evidence. As long as the evidence is supported by the $\sim q$, the probability of *q* will be lowered by the evidence. To see this we need only use Bayes' theorem as follows:

(4) $Pr(p|e) = [Pr(e|p)/Pr(e)]Pr(p)$

(5) $Pr(e|p)/Pr(e))Pr(p)>Pr(p) \Leftrightarrow (Pr(e|p)/Pr(e))>1$ [4]

(6) $Pr(e|p)/Pr(e))>1 \Leftrightarrow Pr(e|p)>Pr(e)$ [5]

(7) $Pr(e|p)>Pr(e) \Rightarrow Pr(p)<Pr(p|e)$ [4,5,6]

The examples employed in the text, then (watch, zebra, car, etcetera), are but a small sample of a pervasive phenomenon.

from p but are not supported by e, if S did not know q before the evidence came in, S does not (given ED) know it after. Claiming otherwise is simply to claim that one can systematically come to know propositions by getting evidence that counts against them. Anyone who is committed to the dependence of empirical knowledge on available evidence has, therefore, good reason for thinking that knowledge is not closed under known entailment.

Now the thing to realize is that all the examples we have been considering are instances of this general phenomenon, namely the openness of evidence. Presumably, one does not know at the outset that one's watch is accurate, that the car has not been stolen, or that the animal in the pen is not a disguised mule (we examine views to the contrary in the following section). The evidence one gains – by looking at the watch, recalling where the car was parked, or seeing a zebra-looking animal – counts against the truth of these propositions. Since counter-evidence cannot be the basis on which knowledge is gained, one does not know these propositions – although one can deduce them from what one knows.[46]

Let us take stock of what has been established in the preceding sections. Given the high plausibility, perhaps even the inevitability, of UD and EEQ, EAD and EDIS must be rejected in order to avoid contradictions. We have also seen independent reasons for rejecting EAD and EDIS having to do with the logic of evidential support. This gives us grounds for rejecting epistemic closure, at least for knowledge that depends on evidence in the sense captured by ED. Now, if closure deniers base their position on the grounds we are proposing, Hawthorne's arguments do not put additional pressure on open

[46] Taking p itself as one's new evidence will not essentially effect the argument. Standard conditionalization is unwarranted on p since its probability is less than 1 (doing so would allow unreasonable amplification of probabilities – just think of the special case of $\Pr(p|p)=1$), and using Jeffrey conditionalization will leave things as they stand.

knowledge.[47] His arguments depend on the employment of EQ together with either AD or DIS, but the arguments we have advanced give closure deniers independent stern reasons to think that neither of these pairs is consistent. Proponents of knowledge openness would be well motivated to deny both AD and DIS on the same grounds that support their denial of closure. The reason to deny all three principles is that evidence supporting a known proposition need not carry over through these modes of inference to the propositions inferred.

An additional point should be stressed in this context. We have demonstrated that in order to keep fundamental features of evidence one must give up principles that may at first blush seem undeniable. It is conceivable therefore, that much of the current distaste with closure denial comes from convictions about evidence that are, in any case, misguided.[48]

3.8 Responses to the Open Knowledge Argument

Before proceeding let us address possible responses to our argument. **One** immediate response could be to argue that, while true, our claims apply to inconclusive evidence (i.e. evidence that does not entail the proposition it supports) while known propositions are

[47] Nozick's and Dreske's rejection of closure is a consequence of their theories of knowledge. In doing so they make inconsistent commitments. Nozick rejects DIS and accepts AD (1981: 236) and EQ (1981: 690 note 60). Thus he is vulnerable to Hawthorne's first argument. Dretske on the other hand is vulnerable to Hawthorne's second argument, since he endorses DIS (1970: 1009) and though he is less than explicit about it, implicitly accepts EQ, as he recognizes in (2005).

[48] Perhaps the most influential argument against open knowledge is due to its entailing what DeRose (1995) called "abominable conjunctions" (conjunctions of the form "He knows this is a zebra, but does not know that it's not a disguised mule"). Such assertions indeed sound odd, but we do not regard this as a reason to endorse closure for the following reasons. First, the oddity of such conjunctions meshes with an explanation of the force of closure-based skeptical arguments. There's thus something to be said for simply accepting a measure of unintuitiveness involved in epistemic openness, particularly when it is explained in terms of the underlying evidential structure. Second, as we have shown abominable evidence construction are unavoidable (e.g. "Her evidence supports *p and q*, but does not support *p* nor does it support *q*"; "his evidence supports *p* and supports *q*, but does not support *either p or q*"). Such conjunctions are a problem for everyone and given the connection between knowledge and evidence it should come as no big surprise to find similar constructions with respect to knowledge.

supported by conclusive evidence only. This is to subscribe to an infallibilist conception of knowledge, a conception according to which knowledge is incompatible with epistemic uncertainty (probabilistically this means that what is known has probability 1). This position faces many difficulties. Besides conflicting with the most natural conception of empirical knowledge, in cases where the evidence clearly does not entail the proposition, e.g. propositions about the future, general propositions justified by induction, or testimony, etcetera, it will either have to explain how this apparent gap (between the proposition and its evidence) is closed, or deny the possibility of knowledge altogether (skepticism). The second option is one we must try to avoid almost at any cost and the first, while surely a valid alternative, is not very promising for reasons we elaborate elsewhere.[49] We shall focus, then, on positions that hold that we can and often do know propositions even though our evidence together with our a priori knowledge do not entail them.

A **second** response is simply to deny ED. But while there might be cases in which one can plausibly go from ignorance to knowledge without any new evidence, as we have claimed, this cannot be a systematic way in which knowledge is gained, particularly empirical knowledge, or knowledge of ordinary, contingent truths. After all, as we have shown, *every* fallibly known proposition together with ED provides a sufficient condition for many potential counterexamples to knowledge closure. Describing cases in which intuitively ED fails is not enough to provide a defense of closure from our argument. ED must fail systematically – with regard to every fallibly known proposition.

A **third** response is to endorse both fallibilism and ED and to claim that when one comes to know *p* one's evidential state changes. Specifically, the claim will be that once *p*

[49] We have in mind reasons related to preface paradox-type considerations and difficulties that arise from the requisite distinction between objective and epistemic probabilities of known propositions. See our MS. There we elaborate on several other aspects of the conception of knowledge as having probability 1, and among other things , show that if justification depends on evidence, justification is not deductively closed.

is known it is added to one's evidence and since *p* entails *q*, one does in fact gain new evidence for *q*. But it is difficult to see how this is supposed to work. Presumably the inferred proposition, *q,* is not known prior to its inference from *p*. So prior to acquiring evidence *e*, *q* was not known. But since *e* reduces the probability of *q*, it is apparently not in virtue of the acquisition of *e* that *q* came to be known. If *e* does not provide the justification enabling knowledge of *q*, its role must be in facilitating the inference. Indeed without *e*, *p* would not have been known and *q* could not have been inferred from knowledge. But if *e* cannot justify *q* and if the inference of *q* is not available without *e*, how can the inference provide more than *e* itself could?

The defenders of epistemic closure might suggest that the inference of *q* from *p* is itself part of the evidence. Since the truth of *p* clearly speaks in favor of the truth of *q* (in fact, it guarantees it) and since *p* is known, *p* can be one's evidence for *q*. But if any item of knowledge is allowed to be knowledge-promoting evidence, non-conclusively based knowledge will provide conclusive (infallible) knowledge. In other words, this suggestion contradicts the idea that one's total evidence does not entail one's total knowledge. This is because if *e* is non-conclusive evidence enabling knowledge of *p*, which implies *q*, and *p* can be taken as evidence for *q*, then *q* is conclusively supported by the evidence (in probabilistic terms, given *p*, the probability of *q* is 1).[50] Still worse, all knowledge will in effect be based on conclusive evidence. The reason is that once *p* is known – no matter what evidence it is based on – if it can serve as one's evidence, it will support itself conclusively. In other words, knowing that *p*, I can, according to the present proposal, use *p* as my evidence, and since I know that *p* entails *p*, I can generate for myself conclusive evidence for

[50] The following proves that: For all empirical propositions *p* and *q*, if *p* entails *q* and can serve as evidence for it, then *p* is conclusive evidence for *q*.

(1)	$(p\Rightarrow q)\supset Pr(p)=Pr(p\wedge q)$	[EQ]
(2)	$Pr(p\wedge q)/Pr(p)=1$	[1]
(3)	$Pr(q\vert p)=1$	[2,conditional probability def.]

p. By trivial logical operations, my evidence has been upgraded from fallible evidence to conclusive evidence. Inductively based knowledge turns instantaneously into knowledge having the full support of deduction. Moreover, since I have such conclusive evidence in support of *p*, I can infer (and therefore know) that any evidence counter to *p* is misleading.[51] In simple terms, then, allowing all known propositions to serve as evidence makes knowledge infallible.[52] The suggestion, then, is simply incoherent.

To avoid this consequence it is necessary for adherents of this proposal to admit that a proposition can provide evidential support only to the degree to which it is itself supported. Thus if one's evidence for *p* raises the probability that *p* is true to 0.8, for instance, *p* can provide evidential support no stronger than that. The transmitted evidential support will not be conclusive, but at most 0.8 probability. But this cannot be a mere technical remedy; it must be explained by the proposed theory of evidence. If *p* cannot support *q* to a degree greater than that to which it is supported, it seems, this must be because its epistemic credibility, so to speak, relies on the support *p* itself enjoys, i.e. the support supplied by *e* (together, perhaps, with relevant background evidence and knowledge). Evidence can support an item of belief only to the degree to which it is itself supported. Thus, while it can warrant (or even require) belief in *q*, *p* offers no evidential support of its own, but

---

[51] For an argument that these (Kripke style) dogmatic beliefs are not known, see our (2010).

[52] If a known proposition can count as evidence for other beliefs Multi-Premise Closure is also valid. The main reasons for questioning Multi-Premise closure is that the risks of falsehood accumulate with each premise and can add up to risk that puts the credibility of one's belief beneath the threshold necessary for knowledge (see Hawthorne 2004a: 46-8 and Stanley 2005: 18). But if the evidence is knowledge, then every known proposition is supported by itself so the risk is annulled. Hence, if knowledge is evidence, and one knows mundane empirical truths, Multi-Premise Closure based on such beliefs is as valid as Single-Premise Closure. But this would saddle us once again with the problems faced by infallibilism. See Hawthorne and Lasonen-Aarnio (2009) and our (MS).

merely the support it received from the total evidence on which it is based.[53] Furthermore, this line of argument cannot escape violation of ED. The reason for this is that if *p* gives *q* no more support than *e* provides it, then p will give *q* no more support than it initially had (remember – *q* is initially more likely than the posterior probability of *p*). Thus the new evidential situation which incorporates *p* will provide no new support for proposition *q*. Hence, if ED is correct, then if *q* is not known prior to knowing *p*, *q* is not known after.

Another expression of the implausibility of the proposal that one's knowledge can always serve as evidence is the unreasonable inflation of knowledge it creates. Having received the final confirmation for my invitation to speak at the departmental colloquium next fall, I know I will be presenting a paper in the fall colloquium. This knowledge, as the present proposal would have it, can support my belief that I will not suffer a fatal disease and die between now and next fall. This belief will in turn justify the belief that I will not collect on my life insurance this year. Do I now have evidence warranting cancellation of my insurance policy? Given that I know I will live to present this paper next fall (and given that I have to put more work into it if I am to make a successful presentation), would I be warranted in canceling my physical checkup scheduled for next week? If you think something funny is happening in such cases due to the high stakes involved in them, think of the watch case. Seeing that the watch shows "3:00" I presumably know that it is three o'clock. It follows from this that if my watch is showing "3:00" it is showing the correct time. It follows further that if your watch shows something other than "3:00" your watch is

[53] This is made evident by the following observation. Suppose, given the rate of breakdowns of your watch, the fact that it shows "3:00" raises the probability that it is three o'clock to 0.9. Suppose further that this is enough to know that it is three o'clock and that this knowledge is now your evidence that your watch is accurate. Presumably, since this latter proposition is entailed by what you know, its probability is no less than 0.9. Now if you receive some weak evidence suggesting that the watch is malfunctioning, we do not say that since you have stronger evidence that the watch is accurate you know this. We do not weigh the new evidence against *p*. The belief that the watch is accurate, it seems, requires some independent support in order to count as knowledge. The support of *p* does not aid *q* if the latter is not itself supported by the evidence. But now if, as we have seen is possible, *e* provides no support for *q* (and assuming there's no other source of evidence), how can the mere presence of *p* improve one's evidential situation at this moment?

mistaken. Would it be reasonable of me to instruct you to reset your watch if it does not read “3:00”? Presumably, the answer to all these questions is a resounding “no”.

A **fourth** possible response to our argument is to claim that we already know all propositions that are entailed by what we come to know which are not supported by the evidence. Thus to know that there is a zebra in the pen one must already know that disguised mules are extremely uncommon. To know that the car is in the driveway one must have background knowledge that car-theft is relatively rare in the relevant area. But this again entails an implausible inflation of knowledge. To know empirical truths about zebras and cars, we need to know not only the general claims that have been mentioned, but also claims like *this animal now has not been disguised to look like a zebra* and *my car today has not been stolen and removed from the driveway,* and so forth. The fact that the rate of car-theft is low does not entail that a specific car at a specific time is not on the unfortunate side of the statistic. But the fact that my car is where I left it does. In other words, knowing that it is unlikely that my car has been stolen from the school parking lot is not the same as knowing that my car is not among the few cars stolen from that parking lot. Only the latter, not the former provides a response to our argument (following the line of the current suggestion).[54] Or take a different example. Boarding the plane in Miami, I know I will land in Chicago in a few hours. This entails that my plane will not crash over Orlando. But do I know this? To avoid the closure negating consequence of our argument that I do not know this, the proposal on offer has to claim that I do in virtue of my background evidence supporting the proposition that planes do not ordinarily crash. But, first, it is not enough that I know this general statistical truth, I must also know that I will not be on the less fortunate

[54] It is important to note that ED is not a threshold claim. It concerns the “direction” of support and not its magnitude. Stating the point with regard to the current example, if you do not know beforehand (say by looking at the car) that my car is not one of the cars stolen from a specific vicinity, knowing that I remember parking it there will not allow you to go from ignorance to knowledge that it has not been stolen from that vicinity.

side of this statistic. That is, I must know that mine will not be one of the X% of flights that do crash, a proposition for which I cannot have evidence. The new evidence, e.g. that I'm boarding the plane to Miami, raises the probability that I will crash over Orlando. Still worse, since this must be the case for every time I board a plane, then assuming I will not crash, I am also in the outrageously happy position of knowing I will *never* be in a plane crash. This is why knowledge of general background facts, while in itself plausible, is not enough to relieve the present worry. While it is plausible to claim that based on my prior experience I know that my watch is generally reliable and even that this licenses me to assume that it is working properly now, it is more than odd to claim that I know it has not just stopped, just as knowledge that the rate of car theft is low where I parked my vehicle does not provide knowledge that it wasn't just stolen. Taking all the (relevant) evidence at one's disposal – the rate of auto-theft, the location where the vehicle was parked etc. – it does not entail that one's vehicle has not been stolen. Had it entailed this, one would know the location of one's car conclusively, contra our fallibilist assumptions.

Moreover, a simple probabilistic presentation shows that the proposal we are considering leads to an inflation of a priori knowledge. Since p entails $\sim(e^*\wedge\sim p)$ (when e^* is one's total evidence), and because, by definition, e^* does not support this implication (e^* lowers its probability), it must be known a priori (if the total evidence does not support a known proposition, it must be known a priori). But if this implication is known a priori, then once e^* is acquired, p is known conclusively, that is, its probability is 1. The reason for this is that conditional on one's total evidence and one's a priori knowledge the probability of any known proposition will be 1 ($\Pr(p|\ e^*\wedge\sim(\ e^*\wedge\sim p))= \Pr(p|\ e^*\wedge p)= \Pr(p|p)=1$).[55] The mistake

[55] One might claim that a priori knowledge of this type is not available for conditionalization, but this must be considered a desperate *ad hoc* measure. A priori knowledge, i.e. knowledge that does not require relevant evidence, is exactly the kind of knowledge that we can and should normally be warranted perhaps even required to conditionalize on.

is to believe that one's new evidence together with one's total prior evidence can simultaneously support p and all of its logical implications. But on a probabilistic understanding of the *evidence for* relation, this is a mathematical impossibility unless p is supported conclusively (that is, with probability 1).[56]

The appeal to the entirety of one's evidence is not really necessary. Imagine that I learn from you that you have a car which you parked in a certain area where the rate of car-theft is exceptionally low. My evidence (your report to me) raises the probability that your car is in the area where you say you parked it, but it does not raise the probability that if your car was parked in an area where car-theft is quite rare, then it has not since been stolen from that area. I have not gained any evidence that this proposition is true (in fact the posterior probability is lower than it was before getting the evidence). But since I have no other relevant evidence, according to the current suggestion I would have to have known this proposition before your report to me, that is a priori.[57] In other words, I know thanks to your report where your car is and since I can derive this conditional a priori from this

[56] We have shown (fn. 45) that for every proposition that is not based on conclusive evidence and is known, there are many propositions that follow from it that are not supported by the totality of one's evidence. As long as the known proposition does not have probability 1, the same will hold when we add background a priori knowledge to the totality of a subject's evidence. Stated differently, the argument is this. Since we can prove that for all p and all a priori knowledge (AK), if $\Pr(p|E\wedge AK)<1$ there is a proposition q that follows from p such that $\Pr(q|E\wedge AK)< \Pr(q)$, it follows from the current proposal that q must be known without evidence and is not known a priori (if it is known a priori, it should be part of AK and hence the inequality is false rendering false the assumption that $\Pr(p|E\wedge AK)<1$). Hence it seems we are left with the following choices: Either q-type propositions are known a priori and knowledge is infallible (contrary to our assumption), or they are known a posteriori and ED is false. But disregarding ED while maintaining fallibilism is incoherent as well since, as we have seen probabilistically and non-probabilistically, every proposition supported by a total body of a subject's evidence that does not entail that proposition has consequences that are not supported by the totality of one's evidence. And so unless we want to distinguish knowledge not based on evidence of this kind from a priori knowledge, the only option we are left with is that knowledge is infallible and a priori knowledge is much more widespread than we may have imagined. In any event pending any new suggestion of how it may still cohere, the prospects of the suggestion that fallible knowledge can be combined with closure look dim.

[57] Even if we assume that this defense of knowledge-closure can be made to work, it does not support the claim that knowledge can be extended by proper deductive inference, which is a major driving force of the closure intuition Williamson (2000: 117) (if the extension is known, it is known a priori so inference is superfluous).

knowledge, there are only two alternatives.[58] Either I already knew it beforehand, or not. If I knew that the conditional is true beforehand this knowledge must be a priori (I had no relevant evidence before your report), but then my knowledge that your car is where you parked it will turn out infallible. If I did not know this, the evidence I gained will not help me, since the only evidence I gained counts (if anything) against this conditional. Needless to say, if you are in the same evidential situation as I, you too are faced with the same predicament. You do not know a consequence of something that you know a priori follows from your knowledge. The same argument, as we have shown, is applicable to every proposition you know that is not entailed by the totality of your evidence and your a priori knowledge. Pending an explanation of how our a priori knowledge of contingents could be so widespread and how this suggestion can be made to cohere with fallibilism, the proposal is unfounded.[59]

Perhaps, **fifth**, one might be tempted to suggest that on evidence *e* one comes to know both propositions *p* and *q* at once. On one interpretation this is simply to deny ED, and on another it will not provide a proper response. One way to interpret the idea is as claiming that on the basis of *e* one comes to know *p* and one comes to know that *q* at the same time. This entails a denial of ED since *e,* as we have seen, supports the negation of *q*. Whether

[58] This example can also be used to show that the order of receiving the evidence should not make an epistemic difference, unless, of course, it is claimed, implausibly, that only you know where your car is and I do not. Moreover, the thought that somehow the background knowledge can be in place since the evidence can first raise the probability of the conditional and then have it go down slightly without destroying one's knowledge, will not work on this and many similar examples.

[59] Notice that although there is a failure of warrant transmission in the cases we are inspecting (Wright (2000) and Davies (2000)), warrant transmission failure does not fill the evidential lacuna required for preservation of knowledge closure. In other words, transmission failure cannot answer the challenge posed here to epistemic closure unless it is accompanied by an explanation of how the requisite background knowledge is attained, an explanation, we argue, that is not possible given fallibilist assumptions. If fallibilism is not assumed, then there is no room for transmission failure. We would like to thank an anonymous referee for making us think harder about the relations between these two issues.

or not *p* gets to be known in virtue of *e* makes no difference.[60] A second interpretation is that the conjunction *p-and-q* is known on the basis of *e*. The open knowledge proponent will have no qualms with this claim since this conjunction is a priori equivalent to *p*. She will have a problem, however, with the next step of deducing and coming to know that *q* from the conjunction since this again will violate ED.

In general, any reply to our argument on behalf of closure must take one of two courses. One is to claim that the inferred proposition is known although the only change in one's evidential state (from a previous time in which per hypothesis both propositions were not known) is the addition of evidence counting against it.[61] The other is to claim that in the course of inference knowledge of the known proposition is lost. This despite the fact that one's evidential state remains unchanged.[62] Both options do not take seriously either the *evidence for* relation and its logic or the dependence of knowledge on evidence, that is,

[60] It might be thought that since the probability of *q* cannot be lower than that of *p*, if *p* is known *q* must be known, or at least knowable, as well. But as the case of lotteries shows high probability conditional on the total evidence does not guarantee knowledge. Our argument concerns what one has evidence for, i.e. relative to any state what does one have evidence for given all of one's evidence, and not on the probability of propositions on one's total evidence or the degree of rational credence (which is influenced by initial credence assignments). See section 4.2.

[61] For every body of evidence that does not entail the known proposition *p*, there will be some proposition *q* entailed by *p* but not supported by the evidence supporting *p* as can be proven by taking *e** to be the total evidence and q to be $\sim(\sim p \wedge e^*)$. See footnote 45.

[62] Some might be tempted to offer a contextualist (or subject sensitive invariantists) reply to the above argument. The basic idea is that inferring from a known proposition sometimes changes standards for knowledge ascriptions resulting in loss of prior knowledge rater than gaining knowledge of what is inferred. Knowledge, then, remains deductively closed. We do not deny that the plausibility of such cases, but they do not seem to cover all instances of apparent closure failure. To properly respond to the argument from Evidence Dependence standard-shifts must be shown to occur with systematic congruence with evidential support relations. The features commonly associated with shifting standards – practical environments, salience, etc. – do not characterize many of the problematic cases we have been looking at.. Realizing this, leading proponents of contextualism and subject sensitive invariantism, e.g. Cohen (2002, 2005), have not relied on standard-shifting to handle some of the cases that fall under the ED claim. See Hawthorne 2004 for similar remarks regarding Cohen's easy knowledge problem.

that knowledge cannot be gained without also gaining evidence for what becomes known.[63]

4 Evidential Knowledge

The attempt to provide the advocate of epistemic openness with grounds for his position is in effect complete. Our defense of the idea that knowledge is open from Hawthorne's objections has given rise to reasons for thinking that knowledge is open, namely, that evidence is not closed under known entailment. To complete our argument we now turn to specify some theoretical advantages of knowledge openness. The benefits of epistemic openness, we show, reach beyond the foregoing considerations regarding evidence - which we take to be the primary basis for epistemic openness - and bear on many of the central issues of contemporary epistemology. We also show how our position can accommodate the intuition motivating closure, i.e. the idea that a belief formed on the basis of competent inference from a justified belief is itself justified. We do this by proposing a distinction between two types of justification, one of which is closed under deduction but does not facilitate knowledge, while the other is knowledge-conducive but not closed.

4.1 The Benefits of Epistemic Openness

The openness of evidence, we said, provides the advocate of epistemic openness with a reasonable positive account for his position and a defense against attacks of the sort mounted by Hawthorne (together with an explanation of why his argument seems so compelling). But, as in urban planning, there are other, environmental reasons for preferring

---

[63] High probability conditional on one's total evidence, is influenced decisively by subjective prior probabilities and therefore should not be confused with having evidence in favor of a proposition. Gettier examples, lotteries and skeptical hypotheses demonstrate that high subjective probability is insufficient (on its own) for knowledge. To go from ignorance to knowledge one needs to gain evidence supporting the proposition, regardless of its credence or probability.

openness to closure. In the case of knowledge the relevant environment consists of a host of epistemological problems that have seemed quite resilient to proposed solutions, which are easily solved, or rather dissolved, once epistemic closure is denied (for the right reasons).

Skepticisms of various sorts rely on the validity of closure.[64] These are not merely Cartesian skeptics, i.e. skeptics undermining entire realms of knowledge, but also more mundane skeptics.[65] Skeptics of both brands argue from the admitted lack of knowledge of an inferred proposition to the dismissal of ordinary knowledge claims. It is easy to see that this maneuver cannot get off the ground without closure. Kripkean Epistemic dogmatism is the idea that, since the truth of *p* implies that evidence counter to *p* is misleading, knowing that *p* one can also know by mere reflection that any counter-evidence is misleading and thereby be – absurdly – warranted in disregarding evidence counting against what one believes. Again, if closure is denied, the inference is invalid and the odd knowledge claim is avoided. Similar considerations apply to lottery propositions. Knowing mundane propositions about the future does not commit one to knowledge that one's lottery ticket is a loser or that one will not be one of the unfortunate victims of sudden heart attacks etc. (Hawthorne 2004). Easy knowledge of the reliability of one's faculties (Cohen 2002, 2005) is also blocked once closure is discarded. Likewise with respect to the bootstrapping problem (Vogel 2000, 2007). The correlation between what I believe is true and the deliver-

---

[64] Pace Klein (2004). While we find Klein's arguments problematic, we cannot address them here.

[65] A mundane skeptic is one that does not target entire realms of knowledge in one fell swoop ("there's no knowledge of the external world"), but rather works piecemeal ("how do you know it's a zebra if you don't know it's not a disguised mule?"). She utilizes the gap of fallibility between knowledge and evidence and points out the implications of the proposed knowledge for which one lacks evidence (the gap guarantees that there are such possibilities). Since her opponent has no evidence for such propositions, he is expected to take back his original knowledge claim. By demonstrating that this maneuver can be used for all fallible knowledge, the mundane skeptic gains the upper hand. Her appeal is to a method rather than a hypothesis (as is common with e.g. the skeptical argument from illusion). See Vogel (1990) for an argument of this type.

ances of my faculties does not provide knowledge that my faculties are reliable. All these (and perhaps some other) problems do not so much as arise once closure is given up.[66]

It should be noted that our account of why closure fails is readily applicable to each of these cases. Seeing my hand provides me with evidence that I have a hand but not that I'm not a brain in a vat deluded to believe that I have a hand. Evidence for *p* can support my belief that *p* is true, but does not indicate that evidence against *p* is misleading. My promise to meet you at the movies does not make it more probable that I will not fall on the way and break my leg, or that my folks will not show up for a surprise visit. Equally, experiencing perception of red patches makes it more likely that there are red patches before me, but not that my perceptual faculties are functioning well. A single account that both explains and dissolves a wide range of what were previously considered resilient and detached problems, is surely very attractive and deserving of serious attention.

4.2 Denying Closure: Not as Bad as You Think

Giving up epistemic closure surely has its costs. Strong intuitions support the principle of closure, not least among them is the idea that inference is a good source of justification. Whatever one's theory of knowledge, a belief formed via proper inference should be a candidate for knowledge. Regardless of whether there's evidence, it would seem, anything properly inferred from a known belief is justified. In this sub-section we claim that epistemic openness need not conflict with this idea. By drawing a distinction between belief (or doxastic) justification and knowledge-promoting (or epistemic) justification, knowl-

---

[66] It also avoids the Gettier style problems we raise in our (2010) and one of the problems for compatibilism of semantic externalism with first-person access (see Brown 2004: 239-42) and can explain failures of warrant transmission. This is perhaps the place to note that the watch example represents a type of case not covered by the standard account of transmission failure (even those who think warrant for believing an animal is not a disguised mule is a necessary precondition for knowing that it is a zebra, will, we presume, agree that to know that it's three o'clock one does not need to be already warranted in believing that even if the watch is broken it is showing the right time now). See Wright (2000) and Davies (2000).

edge closure can be denied without thereby undermining the justificatory capacity of inference. The issues pertaining to justification are copious and convoluted and surely cannot be exhausted here. Our aim here is merely to tease out some intuitions and common conceptions about justification that can go some way towards clarifying and supporting the distinction between doxastic and epistemic justification. Given this distinction, we show that epistemic openness need not be as alarming as it appears.

It is widely accepted that knowing *p* requires not just having evidence or justification for believing *p*, but also forming the belief on the basis of this justification. But Gettier cases show that, though necessary, even this does not suffice for knowledge. Russell's example, for instance, of forming a correct belief regarding the time of day on the basis of a stopped clock illustrates that even if the belief is based on one's justification – and is thus properly justified – still, it might not amount to knowledge.[67] Some philosophers believe that different types of belief require different types of justification. Knowledge of a mathematical theorem's truth, according to these philosophers, requires knowing its proof. Believing it, say, on the basis of testimony, although possibility justifying the belief, cannot provide sufficient grounds for knowledge. Even those who dispute such a distinction between types of beliefs tend to agree that reasons to ascribe high probability do not always promote knowledge. Presumably, one knows it is highly probable that a lottery ticket will lose, and is thus justified in believing it will lose, yet we are not inclined to say that one *knows* the ticket is a loser. A belief that is (known to be) highly probable is surely justified. But if justification in the sense of reason-to-ascribe-high-probability could promote knowledge, then at least some lottery propositions would be known. Or take the example of believing there is a sheep in the field based on seeing a sheep-shaped rock behind which

[67] Russell (1948, p. 154). Russell mentions similar Gettieresque worries about knowledge much earlier, see his (1912: 132).

a sheep happens to be grazing. Perception of a sheep-shaped object in the field surely raises the probability (or the subjective credence) that a sheep is in the field, thus making it reasonable to believe it, and yet under the circumstances one would not be said to know as much. Knowing my financial state, it would be reasonable of me to believe that, despite my life-long dream, I will not buy a classic estate in Provence this year. But if my long-lost uncle has just tracked me down and is planning to bequeath me a large sum of money, my belief does not amount to knowledge, even if eventually I do not receive the money.[68]

To gain some clarity, we may distinguish between different notions of justification here. One can be justified in believing *q* on the grounds, for example, that this is what one must believe in order to retain coherence amongst one's beliefs. Thus, we may have reason to believe that there are external objects if we are to maintain coherence without revising a wealth of our beliefs. In this sense one can be said to be justified *in believing q*. But is this evidence telling in favor of *q*'s truth? Not necessarily. The fact that coherence amongst our beliefs requires us to believe that the external world is real does not constitute evidence telling in favor of it being real. Yet, it can justify us in believing that the external world is real.

A second notion of justification – *epistemic justification* – requires evidence.[69] In this sense, a belief is justified when, for instance, it is supported by the evidence or has been formed in the right way (by reliable method or whatever). Thus, if one believes something on the basis of a false (yet epistemically justified) belief, one can be doxastically justified in believing it while the belief itself is not epistemically justified. (We have already argued probabilistically and non-probabilistically that evidence is not deductively closed.) Doxas-

[68] Hawthorne ascribes a similar example to Joseph Raz (Hawthorne 2004a: 65) and Harman (1973) presents similar examples.

[69] For a similar distinction see Engel (1992).

tic justification for believing, then, does not suffice for knowledge. Even those who think justification is a necessary condition for knowledge, will agree that being merely doxastically justified in believing something does not always guarantee knowledge, even if the belief is true. As lottery and other cases show, doxastically justified beliefs may not amount to knowledge.[70]

What about beliefs justified by inference? Surely, the mere fact that a belief is the product of a valid inference does not suffice for it to count as knowledge. The inference has to be from a true and *justified* belief. But then if the justification of the original belief is evidential, and evidence is not deductively closed, what reason is there to think that the inferred belief is evidentially justified? Inference, it seems, is not an independent source of justification, at most, it *transmits* justification from beliefs to inferred beliefs.[71] But, as our argument has shown, at least one type of justification, namely evidential justification, does not transmit across inference. Therefore, to insist that inferring a proposition supplies one with knowledge-promoting justification for its truth is, in the present context, to beg the question.

But if doxastic justification is not enough for knowledge what else is needed? Here is one proposal following our reflections on evidence and a probabilistic conception of justification. The relation of evidential support, we can say, has at least three dimensions. The *degree* of support, i.e. the conditional probability that a proposition is true given the evidence, is just one dimension. A second dimension can be called the *direction* of support,

---

[70] The distinction does not relate to the degree of justification. Few of our beliefs are as justified, probabilistically speaking, as our beliefs in lottery propositions.

[71] This is shown by the following consideration. Suppose S has justification for *p*. Forming the justified belief that *p*, S then infers from it that *p* is true. Surely her inferred belief does not enjoy a greater degree of justification than her original belief. Inference does not itself provide justification; rather it is supposed to be a mechanism of transmitting justification from premises to conclusion. If *p* implies *q*, the truth of *p* guarantees the truth of *q*, and therefore, presumably, whatever justifies the belief in *p* is also reason for believing that *q* is true.

i.e. whether the evidence raises or lowers the probability whether it counts in favor or against a proposition; and a third dimension is the *magnitude* by which the evidence changes (raises or lowers) the proposition's probability assignment[72]. In each of the cases of closure failure we have canvassed, the evidence functions properly only along the first dimension. It is only the first dimension – the posterior probability – that is preserved through inference. If *p* entails *q*, then, necessarily, the probability of *q* is equal to or higher than that of *p*. Empirical knowledge, we suggest, requires that the inferred proposition be supported in the second dimension as well, i.e. that the probability that the proposition is true be raised by one's evidence, or in non probabilistic terms, that one has evidence for this proposition.

The following example is instructive. Suppose a scientist is wondering whether to invest money in an experiment that, if successful, will confirm *p*. Suppose further that the scientist is not interested in *p* but rather in *q* which is entailed by *p*, and the probability of which will be lowered if the experiment is successful (confirming *p*). Now imagine the scientist reasons as follows: "I am well aware that if the experiment produces the results I expect, it will lower the probability that *q* is true. So I know I will get no new confirmation for *q*. Nevertheless if the experiment works out as planed, I will have evidence for *p*, and will then infer *q* from *p* and thus acquire justification for believing *q*. So, granted, I will have no evidence for my desired conclusion, but still, who needs evidence when there's justification?" We take it that such reasoning is untenable and is arguably a confusion between the two different notions of justification.[73]

The example suggests that the point may be more general than the question of whether the evidence raises or lowers the probability of some proposition; that knowledge requires

[72] There are various ways of measuring this dimension.

[73] This example is inspired by Kaplan (1996: 45).

something qualitatively different from what doxastically justifies belief. This is reflected in some of our most entrenched linguistic practices regarding knowledge and belief. While questioning “how do you know?” is perfectly natural and intelligible, the question “how do you believe?” is hardly either of these. Conversely, the question “why do you believe that *p*?” is commonplace, whereas the question “why do you know that *p*?” is very unusual.[7] Notice that both questions pertain to justification. When asked why one believes something one is prompted to provide a justification for one’s belief. When asked how you know something, likewise, you are required to come up with the grounds or justification for your knowledge claim. In both cases the question is what it is that supports one’s belief/knowledge. And yet the question takes on significantly different forms in the context of belief and in the context of knowledge. We use different notions of justification in these respective contexts. When referring to beliefs we ask for one’s reasons for holding them. Referring to knowledge we ask *how it is supported*.

This suggests that knowledge is governed, among other things, by objective external constraints (such as evidence), while belief is primarily sensitive to rational constraints such as reasons and coherence with other attitudes. As the previous reflections suggest, doxastic justification can be based on agent-relative reasons such as coherence and cre-

[74] There can be contexts in which one emphasizes “why do *you* know that *p*?” in which this sentence makes sense, perhaps because *p* was not supposed to be public information.

dence.[75] Being justified in believing something depends on how it relates to the rest of one's attitudes or to one's credences. But this does not always suffice for knowledge. That is why if someone were to ask "how do you know you are not a brain in a vat?" answering "well, it follows from the fact that I have hands" or "it coheres with many of my beliefs" would hardly seem appropriate. When it comes to knowledge, it matters *how* it is justified.[76] Epistemic justification, we might say, is backtracking – it tracks *how* the justification was acquired or based.

To explicate the notion of backtracking consider the parked car scenario. Seeing your car in the driveway justifies your belief that it has not been stolen. Remembering where you parked it justifies the belief that it is where it was parked and this belief in turn justifies (even requires) the belief that it hasn't been stolen. But none of them epistemically justifies this latter belief. For this it matters *how* the justification was received. If – backtracking your justification – we find that your belief is based on your looking at your car, we would not question your knowing that it has not been stolen. But if it was based on memory of parking the car, we do not ascribe to you such knowledge. To doxastically justify q, suffice it that p stand in the appropriate logical relation to q. To justify it epistemically, the way in which p was evidentially established must be taken into account as well.

---

[75] Our use of this notion is akin to Parfit's, despite the obvious difference in context. As Parfit says, agent-relative reasons "are reasons only for the agent…When I call some reason agent-relative, I am not claiming that this reason cannot be a reason for other agents. All that I am claiming is that it may not be" Parfit (1986: 143). The fact that p coheres with my beliefs may be a reason for me to believe it, but might not be a reason for you if your doxastic repertoire is different from mine. It is interesting to note that in the cases we have been discussing whether one's evidence supports p, and thus provides reason for believing q, depends on one's belief states. Since the evidence in each case supports both p and *not-q* (e.g. that I have a hand or that I'm experiencing vat hands), whether it counts as a reason for believing q or not depends on whether one believes that p is true. In general epistemologists neglect the fact that there are those who hold such things as true. Gnostics, for instance, believed that our world is governed by an evil deity while the benevolent God is in exile. Berkeley believed that there are no external material objects. Taking these and other positions more seriously would perhaps facilitate greater appreciation of the kind of justification we are trying to demarcate. While you might be justified in believing that there are material external objects, Berkeley might not have been. But this does not mean you have better evidence then he did.

[76] The same thought, we take it, is behind reliabilism and sensitivity theories of knowledge – it is not enough that one has reason to believe something is true, or that the belief is in itself justified (perhaps it is not even necessary), one must stand in a certain epistemic relation to it.

The point is a simple one. Just as there can be practical reasons for believing something, which provide practical, but not doxastic, justification for one's belief, so too there may be reasons providing doxastic justification, but not epistemic justification (the kind needed for knowledge). Epistemic justification is backtracking – sensitive to the ways in which a belief was formed or acquired. Therefore, when one's belief is based on evidence lowering the probability that it is true, the belief may be doxastically justified (if the probability is high enough), but one does not know it.

Surely, a lot more than we are able to provide here needs to be said about the details of the distinction. What we have tried to show, however, is merely that with the aid of a reasonable distinction between doxastic and epistemic justification – a distinction well-suited to some of our linguistic practices and in line with intuition – the idea that knowledge is open can be sustained, providing its many epistemological benefits, without sacrificing the idea that a belief properly inferred from knowledge is justified. The novelty of this proposal, we might say, is in suggesting an account of epistemic openness while retaining (at least some version of) closure of justification. Surely there might be other ways of capturing this idea. We have attempted here neither a complete theory of justification nor an exhaustive account of its relation to knowledge, but to show that giving up closure does not necessarily require completely abandoning the main intuition behind epistemic closure.[77]

5 Conclusions

The current state of the debate suggests that any position regarding the validity of epistemic closure carries an intuitive cost. We have therefore tried in this article to steer the debate about closure away from the battleground of intuitions and counter-intuitions and

---

[77] Notice that Gettier employs closure of justification, not of knowledge. "[F]or any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q." (Gettier 1963)

into the realm of theoretical considerations. Traditionally, such reasons for rejecting closure were advanced by externalist epistemologies. Philosophers such as Dretske and Nozick are famous (or infamous) for having argued against closure not on the basis of its unintuitive consequences, but rather from their substantive epistemological positions. In contrast to this traditional setting of the debate, our arguments suggest that the dependence of knowledge on evidence provides the most favorable grounds for epistemic openness. Rejections of closure grounded in the subjunctive nature of knowledge do not stand up to Hawthorne's charges of inconsistency. Furthermore, such positions fail to appreciate the dependence of empirical knowledge on evidence and the backtracking structure of epistemic justification. It is these features of knowledge that gives rise to and explain its openness. The position advanced here thus provides a unified account of the failure of various seemingly intuitive epistemic principles and offers a systematic foundation for reaping the numerous theoretical fruits of epistemic openness.

As we have acknowledged above, the denial of closure has its costs. Yet, we think, at least some of its unintuitive consequences are grounded in the unintuitive logic of evidence which all, including those who deny that evidence can be accounted for probabilistically, must accept, and can be (at least partially) accommodated by distinguishing between doxastic and epistemic justification. Since belief is governed by norms of rationality, most prominently coherence, believing that *p* and properly inferring *q* from *p*, one ought to believe *q*. Knowledge, on the other hand, depends on justification and, in the case of empirical knowledge, on evidential justification. If the evidence one has lowers the probability that something is true, one does not know it in virtue of this evidence. This oft-conflated disparity can explain the inclination to dismiss epistemic openness. Whether one ought to believe something depends on its relation to other things one takes to be true and thus on the inferences one makes. But this should not be confused with the question of whether what one has derived enjoys evidential support requisite for the status of knowledge.

The arguments we have presented do not depend on the contentious definition of evidence by purely probabilistic notions. Rather, we have only assumed that evidence does not lower the probability of that which it is evidence for. Even this modest assumption, we have shown, is not needed. By accepting that evidence cannot support both a proposition and its negation and that there are cases of (weak) underdetermination, we are already committed to the rejection of evidence closure, addition (EAD) and distribution (EDIS). What counts as evidence for what and to what degree is an extremely complicated issue, perhaps no less complex than reasoning itself and no less elusive than the ingenuity of our multifarious attempts at reaching truth. This should not deter us from illuminating some aspects of evidential support by identifying and drawing out connections between evidence and principles of probability of which, arguably, we have clearer understanding. The idea we have been following is that without pretending to know what evidence ultimately amounts to, we can show something about the logic of evidence and use it to draw conclusions about knowledge and the principles it is governed by. The evidence-knowledge link provides good ground for being suspicious of principles that do not coalesce with the features of evidence on which, presumably, empirical knowledge normally depends. This suspicion can be formulated as a challenge. If evidence is not deductively closed, how can empirical knowledge be so closed? What allows knowledge to break free from that which it is based on? How can inference provide what the evidence enabling it cannot?

In the course of this argument we have also provided an analysis of why evidence fails to be closed under different logical operations. The basic idea was that although the conditional probability of the implied proposition given the evidence is not low (not lower than that of the proposition supported by the evidence), given high initial probability (relative to the known proposition) the evidence can, and often does, lower the probability that the proposition is true. Thus, the evidence may change what we might call the "direction" of support. Evidence is basically directional, it points in favor of the truth of some proposition or against it. Evidence pointing in favor of one proposition may point against a proposition it entails.

Using this characterization of evidence, we have also claimed that various epistemological issues which are often considered distinct are, at bottom, one and the same phenomenon, namely, the openness of evidence. The puzzle of dogmatism, "lottery propositions", the problem of "easy knowledge", and various kinds of skepticism, are different manifestations of the queer structure of knowledge owing to the openness of evidential support. The implausible ramifications of epistemic closure in the different types of cases discussed in the literature are all one and the same. They all share a common feature, namely, exceeding the scope of the evidence on which the propositions from which they are derived is based.

To conclude, let us propose our point in a more abstract form. Our conception of empirical knowledge includes the following ideas. First, that we have knowledge of ordinary empirical truths. Second that we gain such knowledge by way of evidence that more often than not, is not conclusive (the evidence is compatible with the falsity of what we know). Third, that we do not know certain empirical truths that are implied by what we do know (either because given our epistemic limitations we cannot know them as in skeptical scenarios, or because the grounds we have do not suffice for knowing them, as in the case of ordinary propositions exemplified by the watch, zebra and car cases). And fourth, that knowledge can always be extended by deduction. Combined, these ideas generate a contradiction giving rise to a host of problems and examples that amount to what is perhaps the most pertinent problem of contemporary epistemology.

Various ways have been proposed of how to modify or deny each of the above stated ideas. Skepticism opts for denial of the claim that we have empirical knowledge even of the most mundane sort. Infallibilists deny that knowledge can be had while having inconclusive reasons. Others claim that we have *a priori* knowledge of anti-skeptical propositions and even of mundane implication, or that by having knowledge of ordinary truths we *ipso facto* gain knowledge of their implications. Contextualists accommodate knowledge ascriptions which on an invariantist conception of knowledge would seem bizarre. Subject sensitive invariantism explains behavior and its relation to knowledge by an appeal to practical environments and salience considerations. The costs and shortcomings of each of these proposals are by now familiar. We have tried to show that the variety of problems arising from our ideas about empirical knowledge are owed to the unintuitive features of evidence and that a proper understanding of these features supports the resolution of these problems by rejecting epistemic closure. By sustaining a distinction between doxastic and epistemic justification we were able to account (at least partially) for the intuitive pull of closure – believing that p and that p implies q one is normally justified in believing q. Yet, we maintain, beliefs justified in this way might not amount to knowledge.

Works cited

Brown, Jessica. *Anti-Individualism and Knowledge*, Cambridge: MIT Press, 2004.

Brueckner, Anthony. "Skepticism and Epistemic Closure," *Philosophical Topics* 13, pp. 89–117, 1985.

Brueckner, Anthony. "Strategies for Refuting Closure," *Analysis* 64, pp. 333–335, 2004.

Carnap, R. *Logical Foundations of Probability*, Chicago: Chicago University Press, 1950.

Cohen, Stewart. “Basic Knowledge and the Problem of Easy Knowledge,” *Philosophy and Phenomenological Research*, 65/2 pp. 309-329, 2002.

Cohen, Stewart. “Why Basic Knowledge is Easy Knowledge,” *Philosophy and Phenomenological Research*, 75.2: 417-30, 2005.

Davies, Martin. “Externalism, architecturalism, and epistemic warrant,” In C. Wright, B. Smith & C. Macdonald (eds.), *Knowing Our Own Minds*, Oxford: Oxford University Press, 2000.

DeRose, Keith. ”Solving the Sceptical Puzzle,” *Philosophical Review* 104, pp. 1-52, 1995.

Dretske, Fred. “Epistemic Operators,” *Journal of Philosophy* 67, pp. 1007-1023, 1970.

Dretske, Fred. “Information and Closure,” *Erkenntnis* 64, pp. 409–413, 2006.

Dretske, Fred. “Is Knowledge Closed Under Known Entailment?” in M. Steup & E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, pp. 13-26, 2005.

Dretske, Fred. “Reply to Hawthorne,” in M. Steup & E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, pp. 43–46 2005.

Engel, Mylan. “Personal and Doxastic Justification in Epistemology,” *Philosophical Studies* 67, pp. 133–150, 1992.

Feldman, Richard. “In Defense of Closure,” *Philosophical Quarterly* 45, pp. 487-494, 1995.

Fitelson, Branden. “The Paradox of Confirmation,” *Philosophy Compass* 1/1, pp. 95–113, 2006.

Gettier, Edmund. “Is Justified True Belief Knowledge?,” *Analysis* 23 (6):121-123, 1963.

Goodman, Nelson. *Fact, Fiction, and Forcst*, Cambridge, Mass: Harvard, 1955.

Harman, Gilbert. *Thought*, Princeton: Princeton University Press, 1973.Hawthorne, John. *Knowledge and Lotteries*, Oxford: OUP, 2004a.

Hawthorne, John. “Précis” (pp. 476-481) and “Replies [to Cohen, Harman and Sherman, and Vogel]” (pp. 510-523), Symposium on Knowledge and Lotteries, *Philosophical Issues* 14, 2004b.

Hawthorne, John. “The Case for Closure,” in M. Steup, E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell, pp. 26-43, 2005.

Hawthorne, John and Lasonen-Aarnio, Maria. “Knowledge and Objective Chance,” in Patrick Greenough & Duncan Pritchard [eds.]. *Williamson on Knowledge*, Oxford: Oxford University Press, 2009.

Hempel, C. G. “Studies in the Logic of Confirmation” in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: Free Press, pp. 3-53, 1965.

Kaplan, Mark. *Decision Theory as Philosophy*, Cambridge Mass: Cambridge University Press, 1996.

Klein, Peter. “Closure Matters: Academic Skepticism and Easy Knowledge,” *Philosophical Issues*, 14, Epistemology, pp. 165-184, 2004.

Kolmogorov, A. *Foundations of Probability*, 2nd English ed. Providence, Rhode Island: AMS Chelsea Publishing, 1956.

Kripke, Saul A. *Philosophical Troubles, Collected Papers Vol. I*, Oxford University Press, forthcoming.

Lewis, Jeffrey and King, Gary. “No Evidence on Directional vs. Proximity Voting,” *Political Analysis*, Vol. 8, No. 1, pp. 21-33, 1999.

Moore, G. E., “Proof of an External World,” in *Philosophical Papers*, George Allen & Unwin, London, 1959.

Murphy, Peter. “A Strategy for Assessing Closure,” *Erkenntnis* Volume 65, 3, pp. 365-383, 2006.

Nozick, Robert. *Philosophical Explanations*, Cambridge: Harvard University Press, 1981.

Parfit, Derek. *Reasons and Persons*, Oxford, Oxford, 1986.

Russell, Bertrand. *Human Knowledge Its Scope and Limits*, Simon and Schuster, New York, 1948.

Russell, Bertrand. *The Problems of Philosophy*, Oxford, Oxford, 1912.

Scheffler, Israel and Goodman, Nelson. “Selective Confirmation and the Ravens: A Reply to Foster”, *The Journal of Philosophy,* Vol. 69, No. 3, pp. 78-83, 1972.

Scheffler, Israel. *The Anatomy of Inquiry*, New York: Knopf, 1963.

Sharon, Assaf and Spectre, Levi. “Dogmatism repuzzled,” *Philosophical Studies* 148 (2), 2010.

Sharon, Assaf and Spectre, Levi. “Epistemic closure under deductive inference: what is it and can we afford it?” MS.

Sosa, Ernest and Kim, Jaegwon. (eds.) *Epistemology: An Anthology*, Oxford: Oxford University Press, 2000.

Stanley, Jason. *Knowledge and Practical Interests*, Oxford: Oxford University Press, 2005.

Van Houweling, Tomz and Sniderman “New Evidence on Directional vs. Proximity Voting”, unpublished manuscript.

Vogel, Jonathan. "Are There Counterexamples to the Closure Principle?" in Doubting: Contemporary Perspectives on Skepticism, M. Roth and G. Ross (eds.), Dordrecht: Kluwer Academic Publishers, 1990.

Vogel, Jonathan. "Reliabilism Leveled," *The Journal of Philosophy*, Vol. 97, No. 11. pp. 602-623, 2000.

Vogel, Jonathan. "Subjunctivitis," *Philosophical Studies*, 134, pp. 73–88, 2007.

Warfield, T. A. "When Epistemic Closure Does and Does Not Fail: a Lesson From the History of Epistemology," *Analysis*, 64, pp. 35–41, 2004.

White, Roger. "Problems for dogmatism," *Philosophical Studies* 131 (3):525--57, 2006.

Williamson, Timothy. "The Contingent A Priori: Has It Anything to Do With Indexicals?" *Analysis*, pp. 113-117, 1986.

Williamson, Timothy. *Knowledge and Its Limits*, Oxford: Oxford University Press, 2000.

Williamson, Timothy. "Reply to John Hawthorne and Maria Lasonen-Aarnio," in Patrick Greenough & Duncan Pritchard [eds.]. *Williamson on Knowledge*, Oxford: Oxford University Press, 2009.

Wright, Crispin. "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof," *Philosophical Issues* 10, pp. 140-163, 2000.