

Can Probability Theory Explain Why Closure Is Both Intuitive and Prone to Counterexamples?

Marcello Di Bello

Received: March 3, 2017 / Accepted: July 22, 2017

Abstract Epistemic closure under known implication is the principle that knowledge of φ and knowledge of $\varphi \rightarrow \psi$, together, imply knowledge of ψ . This principle is intuitive, yet several putative counterexamples have been formulated against it. This paper addresses the question, why is epistemic closure both intuitive and prone to counterexamples? In particular, the paper examines whether probability theory can offer an answer to this question based on four strategies. The first probability-based strategy rests on the accumulation of risks. The problem with this strategy is that risk accumulation cannot accommodate certain counterexamples to epistemic closure. The second strategy is based on the idea of evidential support, that is, a piece of evidence supports a proposition whenever it increases the probability of the proposition. This strategy makes progress and can accommodate certain putative counterexamples to closure. However, this strategy also gives rise to a number of counterintuitive results. Finally, there are two broadly probabilistic strategies, one based on the idea of resilient probability and the other on the idea of assumptions that are taken for granted. These strategies are promising but are prone to some of the shortcomings of the second strategy. All in all, I conclude that each strategy fails. Probability theory, then, is unlikely to offer the account we need.

Keywords: Epistemic Closure; Probability Theory; Knowledge; Warrant Transmission; Risk; Confirmation Theory; Assumptions; Resiliency.

1 The Question

Epistemic closure under known implication is the principle that if S knows φ and knows $\varphi \rightarrow \psi$, then S knows ψ .¹ Countless examples illustrate why this is an intuitive principle. Suppose you know, first, that the defendant killed the victim accidentally, and second, that if the defendant killed the victim accidentally, he is not guilty of premeditated homicide. So, you also know the defendant is not guilty of premeditated homicide. Or, suppose you know the trash is collected once a week on Fridays, and also, that if the trash is collected once a week on Fridays, it won't be collected on Mondays. So, you also know the trash won't be collected on Mondays. And so on. Taken at face value, these and many other examples suggest that

typically if S knows φ and knows $\varphi \rightarrow \psi$, then S knows ψ .

This claim is uncontroversial. Not even those who deny epistemic closure will challenge it.² But this is not quite the same as *full* epistemic closure, that is,

always if S knows φ and knows $\varphi \rightarrow \psi$, then S knows ψ .

In fact, the literature is replete with putative counterexamples. Consider

- (A-1) Amaril knows H ; and
- (A-2) Amaril knows $H \rightarrow \neg B$.

where H stands for 'she has hands' and B for 'she is a brain-in-a-vat'. Assuming Amaril has sensory evidence that she has hands, (A-1) intuitively holds, at least based on a fallibilist theory of knowledge. Item (A-2) intuitively holds in virtue of the meaning of the words involved: handed creatures, by definition, are not handless brains. So long as Amaril understands English, she should know the implication. Now, from (A-1) and (A-2), by the full principle of epistemic closure, it follows that

- (A-3) Amaril knows $\neg B$.

But (A-3) is questionable. If Amaril's evidence consists in the deliverances of her senses and her knowledge of the meaning of words, how can this be enough to know the denial of a skeptical hypothesis? Skepticism cannot be so easily refuted. The oddity of (A-3) can be straightforwardly avoided by denying the full principle of epistemic closure. The pair of propositions H and $\neg B$ would then be an exception or a counterexample to full epistemic closure.³

¹ There are many formulations in the literature. One formulation is more explicit about the underlying psychological process. That is, if S knows φ , competently infers ψ from φ and comes to believe ψ while maintaining knowledge of φ throughout, then S knows ψ ; see [Williamson(2000)] and [Hawthorne(2004)]. Another formulation is in terms of being in a position to know. That is, if S is in a position to know φ and in position to know $\varphi \rightarrow \psi$, then S is in a position to know ψ . For simplicity, I will work with the simple formulation in the main text, but the argument of the paper *mutatis mutandis* should also hold for the other formulations.

² '[T]o deny closure is not to say that you can never know (find out, discover, learn) that Q is true by inferring it from a P that you know to be true. It is merely to deny that this can be done for *any* Q ' [Dretske(2005), p. 17].

³ See [Dretske(1970)], [Nozick(1981)] and [Vogel(1990)].

The literature has traditionally been concerned with arguments for or against full epistemic closure.⁴ But, in recent years, a growing body of literature has attempted to carve out a middle way. This literature recognizes the force of certain putative counterexamples such as the one just presented, but also admits that denying full epistemic closure cannot simply amount to denying any principle of closure whatsoever. After all, that would leave us with an epistemological theory that is deficient in significant ways.⁵ Once full epistemic closure is denied, a theory is still needed which can explain why epistemic closure typically succeeds as well as why it fails in a number of cases.⁶

The most well-known proposals in the literature are based on theories that relativize knowledge to subject matter or a set of contrast questions.⁷ The literature also contains suggestions to the effect that probability theory should be able to offer a satisfactory explanation of why epistemic closure sometimes fails and also why it typically succeeds, although these suggestions have not been developed systematically.⁸ To be sure, given the connections between probability, evidence and epistemology more generally, it would be good if probability were the right epistemological theory underlying epistemic closure. Whether probability theory can be successful in this respect is the question of this paper. To put it more precisely, the question is whether combining probability theory with structural principles linking knowledge and probability can offer an explanation of, first, cases in which epistemic closure intuitively succeeds, and second, cases in which epistemic closure intuitively fails, for example, statements such as (A-1), (A-2) and (A-3).

As it turns out, probability theory cannot offer the explanation we need. To substantiate this claim, the paper examines four probability-based strategies: risk accumulation;⁹ evidential support;¹⁰ assumptions;¹¹ and resiliency.¹² I will argue that the first two strategies do not offer a satisfactory account of the

⁴ Closure deniers include [Dretske(1970)] and [Nozick(1981)]. Defenders of closure include [Williamson(2000)], [Hawthorne(2004)] and [Kripke(2011)].

⁵ As [Kripke(2011)] put it, ‘with a mere rejection of the deductive closure of knowledge, anyone who proves anything from known premises could be criticized for the well-known fallacy of giving a valid argument for a conclusion!’ (p. 200).

⁶ By the word ‘explanation’, I mean an independently motivated epistemological theory that can yield the intended predictions about successes and failures of epistemic closure in various circumstances.

⁷ On knowledge relative to subject matter, see [Yablo(2014)]. On knowledge relative to a set of contrast questions, see [Schaffer(2007)]. For an approach that combines subject matter, topic of inquiry and bounded resources, see [Di Bello(2014)]. For other approaches, see [Holliday(2015)] and [Hawke(2016)].

⁸ See, in particular, [Sharon and Spectre(2013)] and [Sharon and Spectre(2017)]. In a different way, see also [Lasonen-Aarnio(2008)].

⁹ This is a well-known strategy for cases of multi-premise epistemic closure. Interestingly, [Lasonen-Aarnio(2008)] extends this strategy to cases of single-premise epistemic closure.

¹⁰ See [Cohen(2005)], [White(2006)], [Silins(2008)], [Weisberg(2010)], and [Sharon and Spectre(2013)].

¹¹ On the role of assumptions in a theory of knowledge, see [Wright(2004)] and [Sherman and Harman(2011)].

¹² On resiliency, see [Skyrms(1980)] and [Leitgeb(2014)].

intuitiveness of epistemic closure and of its putative failures (sections 2 and 3). I will argue, instead, that the other two strategies are more promising but the main explanatory work is not done by probability theory (sections 4 and 5). I will also consider putative counterexamples to epistemic closure that do not involve far-fetched skeptical scenarios such as brains-in-a-vat. I will argue that the four probability-based strategies considered here still face problems with these putative counterexamples (sections 6). The result of this paper, then, is mostly negative: probability theory, supplemented by structural principles linking probability and knowledge, cannot offer the explanation we need.

Before getting started, a few preliminary remarks are in order. The first is about the interpretation of probability.¹³ The principles linking probability and knowledge that will appear in the paper are unintelligible on the frequency-based or propensity-based interpretations. For instance, it is difficult to see how brain-in-vat propositions have well-defined frequencies or propensities. This leaves the subjective interpretation, based on degrees of beliefs, and the evidential interpretation, where the probability of a proposition is the degree to which evidence supports the proposition. I am not committed to one interpretation in particular, and I leave it to the reader to pick the most suitable.

Second, there are two possible responses to the oddity of (A-3). One is to deny full epistemic closure, as suggested above. The other is to deny that warrant always transmits across a known implication.¹⁴ It is not the task here to adjudicate between the two responses, and for the purpose of this paper, the differences are not extremely significant. When it comes to account for the oddity of (A-3), I shall mostly focus on whether probability theory can offer a satisfactory account of why Amaril does not know $\neg B$ *on the basis of just the sensory evidence and her knowledge of the meaning of words*. Insofar as probability theory can offer such an account, it will satisfy the demands of those who expect an account of warrant transmission failure. It will also satisfy the demands of those who expect an account of closure failure provided an implicit assumption is granted. This is the assumption that sensory evidence, along with the knowledge of the meaning of words, is the only relevant evidence available to Amaril about $\neg B$. So, in absence of other relevant evidence, if

¹³ Regardless of the interpretation, I am using Kolmogorov's axioms, that is,

(non-negativity) $0 \leq P(\varphi) \leq 1$, for any formula φ in the language;
(normalization) $P(\top) = 1$, with \top a logical tautology; and
(additivity) $P(\varphi \wedge \psi) = P(\varphi) + P(\psi)$ for logically inconsistent formulas φ and ψ in the language.

Conditional probability $Pr(\varphi|\psi)$ is defined as $P(\varphi \wedge \psi)/P(\psi)$. I am interpreting $\varphi \rightarrow \psi$ as equivalent to $\neg\varphi \vee \psi$. In order to avoid Lewis-style triviality results, $P(\psi|\varphi)$ cannot be equated to $P(\varphi \rightarrow \psi)$.

¹⁴ See [Davies(1998)] and [Pryor(2012)]. On probability theory and warrant transmission failure, see [Pynn(2013)]. I use 'warrant' to mean what is added to true belief to yield knowledge. What is problematic with (A-3), according to this response, is not that Amaril knows $\neg B$, but rather, that she knows $\neg B$ on the basis of just the sensory evidence and her knowledge of the meaning of words. While sensory evidence is a warrant for H , it is not a warrant for $\neg B$, not even together with one's knowledge of the meaning of words. On this reading, Amaril can still know $\neg B$ *by other means*, and if so, full epistemic closure need not fail.

Amaril fails to know $\neg B$ on the basis of said evidence, she must fail to know $\neg B$ *simpliciter*. There is no need, therefore, to treat the two responses separately.

Finally, I am only concerned with epistemic closure about contingent propositions. For the sake of simplicity, tautologies and mathematical theorems will be excluded from the discussion.

2 Risk accumulation

The first strategy appeals to risk accumulation.¹⁵ The basic insight is that whenever risks accumulate significantly, the principle of epistemic closure will fail, and by contrast, so long as risks do not accumulate too significantly, the principle of epistemic closure will succeed. But let us examine this strategy more closely. It consists of two ingredients.

The *first* ingredient is the very idea of risk accumulation. To illustrate, suppose the probabilities of φ and $\varphi \rightarrow \psi$ equal $1 - r$ and $1 - r'$, respectively, where r and r' are the risks associated with each proposition. By the probability calculus, the probability of ψ can be as low as $1 - (r + r')$.¹⁶ If, for example, the risk of error associated with φ and $\varphi \rightarrow \psi$ is .2 for each, the risk of error associated with ψ could be as great as .4. This leads to a potentially dramatic accumulation of risks. Now, to see how risk accumulation affects knowledge, a *second* ingredient is needed, namely a principle that links probability (or risk) and knowledge, such as,

KNOWLEDGE-AS-THRESHOLD: S knows φ *only if* the probability of φ meets an appropriately high probabilistic threshold t .

This is a necessary, not a sufficient condition on knowledge, so the risk-based strategy is not committed to any specific theory of knowledge. It is also a minimal condition that is widely acceptable, unless one holds that we can know improbable propositions.

Now, suppose an epistemic subject S knows φ and $\varphi \rightarrow \psi$. The probabilities of φ and $\varphi \rightarrow \psi$ are $1 - r$ and $1 - r'$, and they barely meet the threshold t . By risk accumulation, the probability of ψ could be as low as $1 - (r + r')$, which would be below t . By KNOWLEDGE-AS-THRESHOLD, S can fail to know ψ , while knowing φ and $\varphi \rightarrow \psi$. Epistemic closure can therefore fail. The strategy based on risk accumulation exploits the observation that different risks accumulate and

¹⁵ The term ‘risk’ has several meanings. To avoid any confusion, the risk associated with a proposition is here understood as the probability that the proposition is false. So, if $P(\varphi) = 1 - r$, the risk associated with φ is r .

¹⁶ Assume $P(\varphi) = 1 - r$ and $P(\varphi \rightarrow \psi) = 1 - r'$. Then we have:

$$P(\neg\varphi \vee \psi) = 1 - r', \text{ since logically equivalent formulas have the same probability.}$$

$$P(\neg\varphi) + P(\psi) - P(\neg\varphi \wedge \psi) = 1 - r', \text{ by the additivity axiom.}$$

$$P(\neg\varphi) + P(\psi) \geq 1 - r', \text{ by algebra.}$$

$$r + P(\psi) \geq 1 - r', \text{ since } P(\varphi) = 1 - r \text{ and } P(\neg\varphi) = 1 - P(\varphi).$$

$$P(\psi) \geq 1 - r' - r, \text{ by algebra.}$$

$$P(\psi) \geq 1 - (r + r'), \text{ by algebra.}$$

derives the conclusion that epistemic closure, in its full force, cannot be true of us.¹⁷ The strategy is also compatible with the claim that whenever risks do not accumulate too significantly and the probability threshold for knowledge is met throughout, knowing an implication and knowing the antecedent puts the epistemic agent in a position to know the consequent. Risk accumulation, then, predicts that closure will fail in some cases and succeed in others.

And yet, this is not enough for our purposes. Consider again items (A-1), (A-2) and (A-3). As said earlier, both (A-1) and (A-2) are plausible insofar as Amaril has sensory evidence and knows English. By contrast, (A-3), which follows from (A-1) and (A-2) by the full principle of epistemic closure, is odd. This is the putative counterexample we expect an account of. Unfortunately, risk accumulation cannot meet this expectation.

To see why, consider a case in which the probability of the implication $\varphi \rightarrow \psi$ equals one. This means that if the probability of φ is $1 - r$, the probability calculus guarantees that the probability of ψ should at least equal $1 - r$. So, whenever the probability of the implication is one, there is no risk accumulation, because the probability of ψ cannot be lower than the probability of φ . Interestingly enough, in our counterexample to epistemic closure, given a reasonable interpretation, the implication has a probability of one. The implication $H \rightarrow \neg B$ holds in virtue of the meaning of the words used: handed creatures, by definition, are not handless brains. The probability of the implication in question, given our understanding of English, must be one.¹⁸

An objection is likely to arise here. The probability of the implication $H \rightarrow \neg B$ cannot be one because we can be wrong about it. This would lead to an accumulation of the risks of error, so the probability of $\neg B$ could be lower than the threshold required for knowledge. If so, we would have a risk-based account of the putative failure of closure.

But there is a problem with this line of reasoning. Recall one of the examples at the beginning of this paper. Suppose you know the defendant killed the victim accidentally (abbreviated, KA). Suppose you also know that if the defendant killed the victim accidentally, he is not guilty of premeditated homicide (abbreviated, $KA \rightarrow \neg G$). So, you know the defendant is not guilty of premeditated homicide (abbreviated, $\neg G$). Now, if the implication $H \rightarrow \neg B$ has a probability below one, other implications which are also true in virtue of the meaning of the words used, such as $KA \rightarrow \neg G$, must have a probability below one. Because of the accumulation of risk, the consequents of these implications, such as $\neg G$, might have a probability that is lower than the threshold

¹⁷ This point is well understood in the case of multi-premise closure. Interestingly, [Lasonen-Aarnio(2008)] makes the same point in the case of single premise closure, roughly, the principle that if S knows φ and competently infers ψ from φ , then S knows ψ by competent inference. Even this principle, she argues, is prone to risk accumulation because of the risk of error underlying the process of inference.

¹⁸ Contrast $H \rightarrow \neg B$ with ‘If there is a party in the neighborhood, Theo will be at the party’. Suppose Theo has not missed one single party in the neighborhood for the past 10 years. Given Theo’s track record, the implication ‘If there is a party in the neighborhood, Theo will be at the party’ is highly probable, but may still be short of one hundred percent probability.

required for knowledge. So, given risk accumulation, the consequents of these implications might not be items of knowledge. But, we do not want to say that despite one's knowledge of KA and $KA \rightarrow \neg G$, one might fail to know $\neg G$. That would be counterintuitive. If we admit that the risk of being mistaken about the implication $H \rightarrow \neg B$ exists, by parity of reasoning, we must also admit the existence of such a risk for alike implications such as $KA \rightarrow \neg G$. The risk will spread, and it will spread where we do not want it to spread.¹⁹

Where does this leave us? Suppose we endorse knowledge-as-threshold and consider cases in which the implication has a probability lower than one. Although both the antecedent and the implication have a high probability, the consequent might have a probability below the required threshold for knowledge. If so, this would account for the putative failure of closure in certain cases. We have seen, however, that there are cases in which the implication itself has a probability of one, and in such cases the probability of the consequent cannot be lower than the probability of the antecedent. No risk accumulation can be invoked here. Interestingly enough, cases in which the implication has a probability of one are precisely those that constitute the putative counterexamples to epistemic closure.

Some will object that the previous analysis, based on risk accumulation, did not explicitly consider the evidence for Amaril's knowledge. Let us remedy that. Suppose Amaril looks at her hands; she inspects and touches them. This body of evidence makes it highly probable, maybe even certain, that she has hands. Next, consider the implication $H \rightarrow \neg B$. As seen earlier, the implication has a probability of one because it is true in virtue of the meaning of the words. So, we have

(A-1-prob) It is highly probable that H given Amaril's sensory evidence.

(A-2-prob) It is one hundred percent probable that $[H \rightarrow \neg B]$ given Amaril's understanding of the meaning of words.

If so, $\neg B$ must have a probability at least as high as the probability of H , so

(A-3-prob) It is highly probable that $\neg B$ given Amaril's sensory evidence and given Amaril's understanding of the meaning of words.²⁰

¹⁹ The same difficulty applies to the account by [Lasonen-Aarnio(2008)]. In her account, it is the process of competent inference or deduction that is subject to risk of error. If so, the same risk must exist in inferring $\neg B$ from H as well as in inferring $\neg G$ from KA . In both cases, after all, the inferential process is based on inspecting the meanings of the words. But if the inferential risk is the same, and if that risk is sufficiently high, there must a closure failure in both cases. But we do not expect, intuitively, a closure failure in the case of KA and $\neg G$.

²⁰ To establish this, two plausible facts should be considered. First, taking into account Amaril's understanding of English should not affect the probability of an ordinary proposition such as H . To deny this would mean that one's linguistic competence can change the probability of empirically evidenced propositions. Second, taking into account Amaril's sensory evidence should not affect the probability of the implication $H \rightarrow \neg B$. So, from (A-1-prob) and (A-2-prob), along with the two plausible facts just mentioned, it follows that

(A-1-prob*) It is highly probable that H given Amaril's sensory evidence and Amaril's understanding of the meaning of words.

Here we encounter the same result as before. Since the implication has a probability of one, if H has a probability on the evidence that is enough to meet the threshold for knowledge, then $\neg B$ must meet the threshold as well. No risk accumulation can be invoked.²¹

3 Evidential support

I now turn to a second strategy. Instead of considering risk accumulation and the overall probability of certain propositions, let us look for an account of whether a piece of evidence supports a certain proposition. We would like to have a probabilistic account of evidential support on which the evidence available to Amaril supports H , to some degree at least, but does not support $\neg B$. Now, since knowledge of a proposition must require some evidential support for that proposition, if the accounts predicts that H is evidentially supported, but $\neg B$ is not, the former will count as an item of knowledge, but not the latter, as desired.

More precisely, the second strategy consists of two ingredients: (1) an account of evidential support; and (2) a principle that links evidential support and knowledge. Let us begin with the first ingredient, that is,

EVIDENTIAL SUPPORT. Evidence e supports proposition φ depending on the extent to which e increases the probability of φ , namely, $P(\varphi|e) > P(\varphi)$. The larger the probability increase, the stronger the evidential support.²²

(A-2-prob*) It is one hundred percent probable that $[H \rightarrow \neg B]$ given Amaril's sensory evidence and Amaril's understanding of the meaning of words.


From (A-1-prob*) and (A-2-prob*), the probability calculus guarantees that, given a body of evidence available, $\neg B$ must have a probability at least as high as the probability of H . So, from (A-1-prob), (A-2-prob), two plausible facts and the probability calculus, (A-3-prob) follows.

²¹ Incidentally, a puzzle arises with probability theory itself, independently of epistemic closure. Consider (A-3-prob). It is an odd conclusion. It is odd to say that $\neg B$ is highly probable, especially considering Amaril's meager evidence. Looking at one's hands and thinking about the meaning of words cannot be enough to conclude that a skeptical hypothesis is highly improbable. Note the parallelism here. It is odd to conclude that Amaril knows she is not a brain-in-a-vat (given meager evidence), and similarly, it is odd to conclude that it is highly probable that Amaril is not a brain-in-a-vat (given meager evidence). Interestingly enough, the argument leading to (A-3-prob) did not appeal to knowledge nor to epistemic closure. It was based on probabilities only. Probability theory gives rise to the same oddity it was supposed to explain. *Tu quoque*, probability! A response here might be that it is not appropriate to consider the *overall probability* of a proposition given the evidence as a way to assess whether the evidence supports the proposition. Rather, it is more appropriate to consider the *changes in probability* of a proposition brought about by the evidence. I deal with this response in detail in the next section.


²² This account, which is standard in confirmation theory, is susceptible to a well-known counterexample by [Achinstein(2001)]. The account is, however, an adequate approximation for the purpose of this paper.

The important feature of this account is that, given two propositions φ and ψ such that the implication $\varphi \rightarrow \psi$ has a probability of one, it does not follow that if some evidence e supports φ , then e also supports ψ .²³

To illustrate, consider H and $\neg B$. Suppose, without the sensory evidence, Amaril does not know whether she has hands, or probabilistically, suppose $P(H) = P(\neg H) = 1/2$. Since H is incompatible with B and thus $P(H \rightarrow \neg B) = 1$, it follows that $P(H \wedge B) = 0$ and $P(H \wedge \neg B) = 1/2$. Instead, $\neg H$ is both compatible with B and $\neg B$, so we can plausibly assume that $P(\neg H \wedge B) = P(\neg H \wedge \neg B) = 1/4$. This means that $P(\neg B) = 3/4$. Pictorially,



Consider now sensory evidence, call it s , showing the presence of hands. Sensory evidence has no bite in the case Amaril is a brain-in-a-vat, B . Note that the hypothesis B , for the purpose of this argument, is understood to mean that Amaril is a brain-in-a-vat that is endowed with simulated hands indistinguishable from real hands. If we were brains-in-a-vat with simulated hands, the sensory evidence would still show the presence of hands. Hence, sensory evidence cannot rule out the hypothesis $\neg H \wedge B$. By contrast, sensory evidence has strong credentials in the case Amaril is not a brain-in-a-vat, $\neg B$. Sensory evidence, then, does rule out $\neg H \wedge \neg B$ but leaves untouched $\neg H \wedge B$. Pictorially,



As we can see, by taking into account the sensory evidence showing the presence of hands, the probability of H increased from $1/2$ to $2/3$, while the probability of $\neg B$ decreased from $3/4$ to $2/3$. This means that the evidence s supports H , but not $\neg B$. This is good news.²⁴

²³ For an extensive discussion of this point in relation to epistemic closure, see [Sharon and Spectre(2013)].

²⁴ [White(2006)] gives an argument to this effect using fake hands instead of brain-in-vat scenarios.

But we are not done yet. Not only do we want an account of why s does not support $\neg B$, but also an account of why Amaril does not know $\neg B$ on the basis of s . To this end, we need a plausible principle that links evidential support and knowledge, such as,

KNOWLEDGE-REQUIRES-SUPPORTING-EVIDENCE: S knows φ on the basis of evidence e *only if* e supports (to a positive degree) φ .

The above principle, together with probability-based evidential support, accounts for why even when Amaril knows she has hands (on the basis of sensory evidence), she need not know she is not a brain-in-a-vat (on the basis of the same sensory evidence). The proposition H is evidentially supported by the sensory evidence, while $\neg B$ is not, for the sensory evidence increases the probability of H without raising the probability of $\neg B$. This accounts for the failure of closure in the counterexample featuring Amaril.


As a sanity check, consider a case in which closure is not expected to fail. Recall the following example from the beginning. Suppose you know the trash is collected once a week on Fridays (abbreviated, F), and also, that if the trash is collected once a week on Fridays, it won't be collected on Mondays (abbreviated, $F \rightarrow \neg M$). So, you also know the trash won't be collected on Mondays (abbreviated, $\neg M$). Now, before acquiring any evidence about when the trash is collected, we can assume that any day from Monday through Friday is equally likely, so $P(F) = P(M) = 1/5$ and $P(\neg M) = 4/5$. Once evidence that the trash is collected only on Fridays is acquired, the probability of F will increase and the probability of M will decrease. We need not determine the exact extent of such an increase. It is clear, however, that if the probability of F increases, the evidence supports F , and if the probability of M decreases, and thus that of $\neg M$ increases, the same evidence supports $\neg M$. Since the probabilities of F and $\neg M$ both increase, the strategy based on evidential support does not predict a failure of closure in this case, as desired.

Things are not so easy, however. There are two problems with this strategy. The first problem has to do with the sensitivity to how the logical space is initially partitioned and how the prior probabilities are assigned to the competing hypotheses. The second problem has to do with conjunctive statements such as $H \wedge \neg B$. Let me consider each in turn.

3.1 First problem

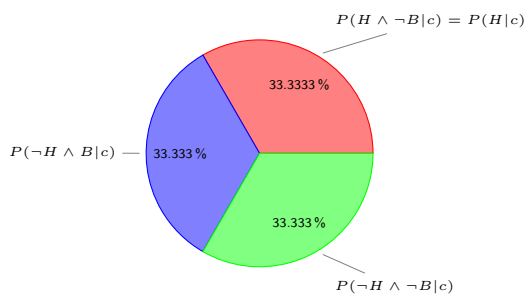
Given a certain partition of the logical space, the strategy based on evidential support accounts for why Amaril does not know she is not a brain-in-a-vat *on the basis of just the sensory evidence*. But, to be precise, the original intuition to be accounted for was why Amaril does not know she is not a brain-in-a-vat on the basis of the sensory evidence *and* her knowledge of the meaning of words. Note that the evidence consisting in the meaning of words tells us that it is impossible to have a world in which one has a hand and at the same time one is a brain-in-a-vat. Call it *conceptual evidence*. So, before considering the

conceptual evidence, four possibilities are still viable. That is, (1) $H \wedge B$; (2) $H \wedge \neg B$; (3) $\neg H \wedge B$; and (4) $\neg H \wedge \neg B$. Assume that without considering the conceptual evidence or the sensory evidence, every possibility is initially equiprobable, that is,




This is a different initial partition of the logical space and a different prior probability distribution than the one considered earlier, but one that is no less plausible.

Now, conceptual evidence, call it c , rules out $H \wedge B$, so we have



Since $P(\neg B) = 1/2$ and $P(\neg B|c) = 2/3$, the conceptual evidence raises the probability of $\neg B$, and thus supports $\neg B$. What happens after adding the sensory evidence s ? By showing the presence of hands, this evidence rules out $\neg H \wedge \neg B$. It cannot rule out the conjunction $\neg H \wedge B$ because, as seen earlier, sensory evidence has no bite if we were brains-in-a-vat with simulated hands indistinguishable from real hands. So,



The upshot is that sensory and conceptual evidence, considered together, leave untouched the probability of H and $\neg B$. The probability was initially $1/2$ and ends up being again $1/2$. On this reading, the two items of evidence do not support $\neg B$, as expected, nor do they support H , contrary to expectations. So, by distinguishing sensory and conceptual evidence, the account of evidential support does not yield the intended predictions.

This is by no means a general argument. But it suggests that a complete defense of the strategy based on evidential approach should explain why some initial partitions of the logical space and some prior probability assignments are not acceptable and while others are. Absent such an explanation, the proposal would be incomplete.

3.2 Second problem

There is, in addition, a deeper problem with the strategy based on evidential support, a problem that would persist even if the previous problem could be satisfactorily addressed. This second problem has to do with the conjunction $H \wedge \neg B$, asserting that one has hands and also that one is not a brain-in-a-vat. Discussions of such conjunctions are common in the literature on epistemic closure, especially in relation to the point that $H \wedge \neg B$ is a priori equivalent to H .²⁵ I shall rehearse some of these discussions for the purpose of the argument here, namely, to raise a problem for the strategy based on evidential support.

Now, if $H \wedge \neg B$ is a priori equivalent to H , the following two hold:

$$(/) P(H) = P(H \wedge \neg B) \quad \text{and} \quad (//) P(H|s) = P(H \wedge \neg B|s),$$

where s is sensory evidence. By $(/)$ and $(//)$, if s increases the probability of H , it must also increase the probability of $H \wedge \neg B$. So, if sensory evidence s supports H , it must also support $H \wedge \neg B$ to the same degree, although s does not support $\neg B$, as demonstrated before. But this result should give us pause. How can meager sensory evidence support $H \wedge \neg B$ which asserts, besides the presence of hands, the denial of a skeptical scenario? Just as it is odd to say that s supports $\neg B$, it is odd to say that s supports $H \wedge \neg B$.

Some might respond that while it is problematic that s would support $\neg B$, it is not so problematic that s would support $H \wedge \neg B$. To make this plausible, confirmation theory can be invoked. In confirmation theory, under suitable conditions, although the probability of a conjunction cannot be higher than the probability of its conjuncts, the degree of evidential support of a conjunction can be higher than the degree of evidential support of its conjuncts.²⁶ In fact, it is possible that a conjunct has no evidential support while the conjunction does. So, although the probability-based account of evidential support yields the conclusion that s supports $H \wedge \neg B$, insofar as this does not entail that s

²⁵ See, for example, the exchange between [Dretske(2005)] and [Hawthorne(2005)]. See, more recently, the exchange between [Yablo(2017)] and [Sharon and Spectre(2017)].

²⁶ For a proof, see [Crupi et al(2008)Crupi, Fitelson, and Tentori].

also supports $\neg B$, there would be no problem. All we want, some might argue, is to block the conclusion that s supports $\neg B$.

But such response, while persuasive within probability theory and confirmation theory, is unsatisfactory. Let us grant, for the sake of argument, that the conclusion that s supports $H \wedge \neg B$ is not problematic. Note that sensory evidence s supports $H \wedge \neg B$ to the same degree as it supports H because the extent of the probability increase is the same. If Amaril knows H , based on the sensory evidence, and if the evidential support is the same for H and $H \wedge \neg B$, we must conclude that Amaril knows $H \wedge \neg B$. Evidential support cannot identify any epistemic difference between H and $H \wedge \neg B$. If one is an item of knowledge, so is the other. But isn't this odd given our starting point?

Our starting point, if you recall, was that it is intuitively problematic to attribute knowledge of $\neg B$, at least whenever the only evidence available is mere sensory evidence showing the presence of hands. This intuitive judgment is what motivated the denial of full epistemic closure and the search for a theory of when epistemic closure fails. Now, if it is intuitively problematic to attribute knowledge of $\neg B$, it should be at least as problematic, if not even more problematic, to attribute knowledge of the logically stronger proposition $H \wedge \neg B$ while holding fixed the evidence available. And yet, this is what the theory of evidential support would force us to do. Given evidential support, although one would fail to know $\neg B$, on the basis of mere sensory evidence (as expected), one would nevertheless be in a position to know the logically stronger proposition $H \wedge \neg B$ (with no change in the evidence whatsoever). This seems odd.

Some might respond that any theory will clash with intuitions at some point, and intuitions in the end are not so clear-cut. This may be so. Still, let us examine the contours of the problem more generally and see what is at stake here. The crux of the matter is that probability-based evidential support validates a closure principle for a priori equivalent propositions. Let φ and ψ be a priori equivalent propositions. If so, whatever evidence supports φ must equally support ψ . As far as evidential support goes, then, there would be no epistemologically relevant difference between φ and ψ . Whenever ψ is an item of knowledge, so is φ , and vice versa. In other words, the strategy based on evidential support validates the following closure principle:

EQUIVALENCE CLOSURE: S knows φ if and only if S knows ψ , where φ and ψ are a priori equivalent propositions.

This principle is not problematic *per se*, but its implications in the context of the strategy based on evidential support are, as we shall soon see.

First note that EQUIVALENCE CLOSURE combined with another apparently innocuous closure principle, that is,

CONJUNCTION CLOSURE: If S knows $\varphi \wedge \psi$, then S knows φ and S knows ψ ,

must yield closure under known implication.²⁷ This means that if one is to deny the latter principle, one has to deny at least one of the two other closure principles above. Since the strategy based on evidential support validates EQUIVALENCE CLOSURE, it must necessarily deny CONJUNCTION CLOSURE, or else this strategy would stand no chance of vindicating the denial of closure under known implication. Less abstractly, given evidential support, if the epistemic agent knows H , she must also know the a priori equivalent proposition $H \wedge \neg B$, and thus by CONJUNCTION CLOSURE, the agent would know $\neg B$. This is precisely the counterintuitive result to be avoided. In order to avoid this result, the strategy based on evidential support must deny CONJUNCTION CLOSURE.

Given the above, the question that the strategy based on evidential support must now face is this. Is denying CONJUNCTION CLOSURE while also maintaining EQUIVALENCE CLOSURE a good way of denying closure under known implication? Some might respond in the affirmative by noting that since a priori equivalent propositions cannot be distinguished probabilistically, they should not be distinguished epistemically, and thus EQUIVALENCE CLOSURE must hold. But this reasoning assumes that probability captures everything there is to capture about epistemology, and this is precisely the point under dispute.²⁸ On the other hand, the controversy here cannot be resolved by appealing to the intuitive plausibility of each closure principle considered separately. After all, closure under known implication is plausible, and so are EQUIVALENCE CLOSURE and CONJUNCTION CLOSURE. They are all plausible principles. The question, then, is whether denying closure under known implication together with denying CONJUNCTION CLOSURE—which is the route the evidential support strategy must take—is more or less plausible than denying closure under known implication together with denying EQUIVALENCE CLOSURE—which is the alternative route for those who want to deny closure under known implication. The question is one of comparative plausibility or implausibility of a set of theoretical moves.

I claim that if one is to deny closure under known implication, the route that consists in denying CONJUNCTION CLOSURE is more implausible than the route that consists in denying EQUIVALENCE CLOSURE. Let me explain why.

Crucially, just as the propositions H and $\neg B$ can be used to put intuitive pressure against closure under known implication, they can also be used to put intuitive pressure against EQUIVALENCE CLOSURE. If one knows H on the basis of, say, sensory evidence as of hands, then one knows $\neg B$ by closure under known implication, and also one knows $H \wedge \neg B$ by EQUIVALENCE CLOSURE.

²⁷ Suppose S knows φ and knows (by the meaning of words, a priori) $\varphi \rightarrow \psi$. If φ is a priori equivalent to $\varphi \wedge \psi$, then S knows φ iff S knows $\varphi \wedge \psi$ by EQUIVALENCE CLOSURE. So S knows $\varphi \wedge \psi$ by *modus ponens*. By CONJUNCTION CLOSURE, S knows ψ . In short, if S knows φ and S knows $\varphi \rightarrow \psi$ (by the meaning of words, a priori), then S knows ψ . This is closure under known implication.

²⁸ [Sharon and Spectre(2017)] argue that ‘if a priori equivalent propositions are not supported or refuted jointly they can fall prey to Dutch-book and to other arguments’. But Dutch book arguments are not directed against those who, while assigning probabilities coherently, attribute knowledge claims in ways that do not entirely align with probability.

It seems no less odd to say that in such circumstances one knows $H \wedge \neg B$ than that one knows $\neg B$. If the oddity of the latter knowledge claim is used to put intuitive pressure against closure under known implication, the oddity of the former knowledge claim should be used to put intuitive pressure against EQUIVALENCE CLOSURE.

By contrast, the set of propositions that is often used to put pressure on closure under known implication, such as H and $\neg B$, can hardly form a counterexample against CONJUNCTION CLOSURE. If one knows $H \wedge \neg B$, there seems to be nothing counterintuitive in saying that one, thereby, knows one of the conjuncts.²⁹ In this sense, there is no counterexample on offer here. And if CONJUNCTION CLOSURE is less prone to counterexamples than EQUIVALENCE CLOSURE, denying the former would be more implausible than denying the latter. This conclusion is also corroborated by a brief survey of the existing, non-probabilistic approaches in the literature which deny epistemic closure under known implication (or variants thereof). They all aim to keep CONJUNCTION CLOSURE, and this means that they must let go of EQUIVALENCE CLOSURE.³⁰ This indicates they must consider the former principle more intuitively plausible than the latter.³¹

To summarize, the argument against the strategy based on evidential support highlighted two main difficulties. The first had to do with the sensitivity to partitioning of the logical space. This difficulty, however, need not cut too deep because the strategy could be suitably amended, for example, by specifying a range of acceptable partitions of the logical space.

The second, deeper difficulty had to do with conjunctive statements of the form $H \wedge \neg B$. The argument in this regard can be broken down into two components. The *first* component shows that if the strategy based on evidential support explains why closure under known implication fails in the problematic

²⁹ To be sure, the problematic part here is the if-statement that one knows $H \wedge \neg B$. How could one know that in the first place? But *if* this conjunction is known, the knowledge of the conjuncts seems unproblematic, in agreement with CONJUNCTION CLOSURE.

³⁰ Here are two representative approaches. The first takes knowledge to be *question sensitive*; see [Schaffer(2007)]. Knowledge can be understood as a ternary relation, between a subject S , a proposition p and a set of contrast questions Q . Proposition p would be known provided the alternatives to p that are selected by Q are ruled out. On this approach, knowledge need not be closed under known implication. After all, the contrast questions associated with $\neg B$ might well be different from the contrast questions associated with H . However, suppose one knows $H \wedge \neg B$. This means that all the contrast questions associated with $H \wedge \neg B$ have been addressed. The contrast questions about each conjunct will be included in the contrast questions about the whole conjunction. Hence, the approach based on question sensitivity does not deny CONJUNCTION CLOSURE. The second approach takes knowledge to be *subject matter sensitive*; see [Yablo(2014)]. On this approach, closure under known implication fails whenever the subject matter of the consequent extends beyond the subject matter of the antecedent, as in the case of $H \rightarrow \neg B$. But now suppose one has knowledge of $H \wedge \neg B$. Since the subject matter of H and $\neg B$, considered separately, does not extend beyond the subject matter of $H \wedge \neg B$, the subject matter account will not deny CONJUNCTION CLOSURE.

³¹ Some authors believe that failing to satisfy CONJUNCTION CLOSURE is a conclusive objection against a theory of knowledge; see [Hawke(2016)] and [Kripke(2011)].

case under consideration, the same strategy, combined with structural principles linking probability and knowledge, must also be committed to denying CONJUNCTION CLOSURE while at the same time endorsing EQUIVALENCE CLOSURE. The *second* component of the argument is that, as a route to denying closure under known implication, denying CONJUNCTION CLOSURE while keeping EQUIVALENCE CLOSURE is less plausible than denying EQUIVALENCE CLOSURE while keeping CONJUNCTION CLOSURE. Other, non-probabilistic approaches in the literature pursue the latter and not the former combination of theoretical moves. In short, the strategy based on evidential support fails to offer a satisfactory account of the failure of epistemic closure under known implication, at least it fails to offer an account that is at least as plausible, or no more implausible, than other non-probabilistic accounts in the literature.

This is not to say that the strategy based on evidential support does not make progress compared to the risk accumulation strategy. But it would be better if a probability-based account of the intuitive failures of closure and of its intuitive successes could be given which was less susceptible to the above shortcomings.

4 Assumptions

The previous two strategies were principally based on probability theory. To be sure, they both invoked plausible, and rather uncontroversial, bridge principles linking probability and knowledge, such as KNOWLEDGE-AS-THRESHOLD and KNOWLEDGE-REQUIRES-SUPPORTING-EVIDENCE. But the main explanatory work was done by probability theory, and precisely, by the idea of risk accumulation in the case of the first strategy, and by the changes in probability in the case of the second strategy. But as it turns out, these strategies are unsatisfactory. I shall now consider two further strategies. They combine insights from probability theory with insights that are not unique to probability theory. We shall see that they point towards a good account of the intuitiveness of epistemic closure as well as of its putative failure in the counterexample featuring Amaril. But we shall also see that such an account is likely to rest on insights that do not come from probability theory itself.

Turning now to the third strategy, let us postulate that the denial of any skeptical hypothesis, for example the denial of hypothesis *B* ‘Amaril is a brain-in-a-vat with simulated hands indistinguishable from real hands’, is always assumed as a background assumption in any probability assignment. This strategy is not very different from what some authors have proposed in the case of knowledge. They have argued that the reasons for one’s knowledge of a proposition include items of evidence as well as assumptions that are taken for granted in absence of directly supporting evidence.³² The proposal here is about probabilities, not knowledge, and the idea is to take the falsity of all skeptical hypotheses as a background assumption in any probability assignment.

³² See, among others, [Wright(2004)] and [Harman and Sherman(2004)].

The assumption-based strategy, in fact, retains all the ingredients of the earlier strategy, namely (1) a probability-based account of evidential support and (2) a principle that links evidential support and knowledge, such as KNOWLEDGE-REQUIRES-SUPPORTING-EVIDENCE. To ingredients (1) and (2), the assumption-based strategy adds that (3) any probability assignment must be conditional on the denial of all skeptical hypotheses.

This strategy is, at first blush, rather promising. The probability of $\neg B$, conditional on the sensory evidence s and the background anti-skeptical assumption $\neg B$, must be one. Although it might seem counterintuitive to say that the probability of $\neg B$ is one, the difficulty disappears after realizing that this is the probability of $\neg B$ conditional on $\neg B$. The important point is that by stipulating that $\neg B$ is always in the background in any probability assignment (together with the denial of many other anti-skeptical assumptions), the probability of $\neg B$ will never change. To illustrate, consider the following probability assignment:

$$\begin{aligned} \text{Without evidence } s: & P(H|\neg B) = r \text{ and } P(\neg B|\neg B) = 1; \\ \text{With evidence } s: & P(H|s, \neg B) = r' \text{ and } P(\neg B|s, \neg B) = 1, \text{ with } r' > r. \end{aligned}$$

The conditional probability of $\neg B$, with and without the sensory evidence s , remains the same. Since the probability of H increases given the sensory evidence s , while the probability of $\neg B$ does not change, it follows that s supports H but not $\neg B$. In terms of knowledge, this means that Amaril may well know H (on the basis of the sensory evidence), but not know $\neg B$ (on the basis of the sensory evidence), as expected.

Interestingly enough, the assumption-based strategy solves the first problem with the probability-based account of evidential support. If you recall, the problem arose because of the distinction between conceptual and sensory evidence. But this distinction need not be made here, or more precisely, every information that the conceptual evidence conveys is already contained in the denial of skeptical hypotheses. What the conceptual evidence ensures is that scenarios such as $H \wedge B$ are ruled out. In fact if $\neg B$ is kept as a background assumption, all scenarios in which B holds are automatically ruled out.

So far so good. Unfortunately, the assumption-based strategy cannot solve the problem with the conjunction $H \wedge \neg B$. Any increase in the probability of H must be accompanied by an equal increase in the probability of $H \wedge \neg B$. To see why, consider

$$\begin{aligned} \text{Without evidence } s: & P(H|\neg B) = r \text{ and } P(H \wedge \neg B|\neg B) = r; \\ \text{With evidence } s: & P(H|s, \neg B) = r' \text{ and } P(H \wedge \neg B|s, \neg B) = r', \end{aligned}$$

where $r' > r$. The probability of H increases from r to r' , and the same applies to $H \wedge \neg B$. So, the sensory evidence s supports H , as expected, but also supports $H \wedge \neg B$. This is the same problem we encountered earlier with the probability-based account of evidential support. This means we are still saddled with the conclusion that Amaril would know $H \wedge \neg B$, contrary to intuitions.

The assumption-based strategy, in this respect, fares no differently than the earlier strategy based on evidential support. Both strategies validate EQUIVALENCE CLOSURE and must therefore deny CONJUNCTION CLOSURE. The same argument that was given earlier for why these theoretical moves are implausible compared to denying EQUIVALENCE CLOSURE while keeping CONJUNCTION CLOSURE applies here. When it comes to the conjunction $H \wedge \neg B$, the assumption-based strategy seems to be in the same predicament as the strategy based on evidential support.

To be sure, the assumption-based strategy has room for further maneuvering. The peculiarity of $H \wedge \neg B$ is that it is a hybrid. It is the conjunction of the anti-skeptical proposition $\neg B$, taken as a background assumption, and the ordinary proposition H . The solution might be that knowledge claims be disallowed for such hybrid conjunctions. This will solve the problem with $H \wedge \neg B$ by declaring it not suitable for knowledge. But now the question arises why such hybrid conjunctions should be disqualified. A plausible answer rests on the need to avoid the circularity of supporting reasons. It is plain enough that one cannot know p simply by assuming p . So, $\neg B$ cannot count as an item of knowledge insofar as it is an assumption. Similarly, propositions which entail assumptions cannot be items of knowledge either. If p is among one's assumptions, $p \wedge q$ cannot be an item of knowledge because the path of reasons supporting $p \wedge q$ would include p itself. That would be circular. So, $H \wedge \neg B$ cannot count as an item of knowledge either.

Another approach is based on topic of inquiry. Every proposition that plays the role of an assumption would count as falling outside the scope of knowledge given a topic of inquiry. So $\neg B$ or B cannot be knowable because they fall outside the scope of knowledge for a given topic of inquiry. Similarly, every other proposition that is constructed from propositions that fall outside the scope of knowledge, cannot fall within the scope of knowledge. So $H \wedge \neg B$ is not knowable.³³

More needs to be said, but the important point is that a priori equivalent propositions can be distinguished in epistemologically meaningful ways. The problem is that if the main explanatory work is done by a scope restriction on knowledge or a ban on the circularity of reasons, probability theory will not be explanatorily central to that distinction. The assumption-based strategy, if at all successful, is likely to be probabilistic only in a marginal sense.

5 Resiliency

It is time for the fourth and final strategy. According to a probability-based account of knowledge, a proposition is known if it is resiliently highly probable on the evidence available.³⁴ On this account, high probability on the evidence, by itself, is not enough for knowledge. In addition, the high probability must be stable across a (reasonable) set of evidential updates that could be performed.

³³ For a more systematic exposition of this approach, see [Di Bello(2014)].

³⁴ See [Skyrms(1980)].

To illustrate, suppose the probability of p is high given some evidence e_0 and it remains high after considering items of evidence e_1, e_2, \dots, e_k , that is,

$$P(p|e_0) \text{ is high and } P(p|e_0, e_1, e_2, \dots, e_k) \text{ is still high.}$$

Then, p counts as an item of knowledge. If, instead, the probability of p is high given some evidence, but becomes much lower in light of other items of evidence, then p won't count as an item of knowledge. The set of evidential updates can be thought of as the set of items of evidence which could contradict, undermine or challenge one's claim to know p . Resiliency, in other words, is a guarantee that p will not be easily falsified or contradicted.

A qualification bears mentioning. The probability of p cannot be resiliently high relative to *any* possible item of evidence. To require that would be to require infallible knowledge. So, we should restrict the set of items of evidence to a *reasonable set*. But how are we to identify this set? Surely, this set must be restricted to the *relevant* items of evidence, where an item of evidence is—probabilistically—relevant for a proposition if it can affect the probability of the proposition, upwards or downwards. But this restriction is not enough. We cannot consider every item of relevant evidence because to require that, once again, would be to require infallible knowledge. The reasonable set, arguably, could be further restricted on the basis of pragmatic considerations such as practical interests, time, available resources. The crucial question here, however, is whether the selection of the reasonable set can be spelled out using the theoretical resources of probability theory itself. Let us leave this question open for the time being, but this is the question that decides whether the resiliency-based strategy is principally probabilistic or not.

Now, since this strategy focuses on knowledge without relying on an account of evidential support, it is not directly prone to the problems that the previous strategies faced. So, if knowledge is understood as resiliently high probability, the resiliency-based strategy is adequate provided the following three hold:

- (i) H has a resiliently high probability given Amaril's evidence;
- (ii) $\neg B$ does not have a resiliently high probability given Amaril's evidence;
- (iii) $H \wedge \neg B$ does not have a resiliently high probability given Amaril's evidence.³⁵

Consider first H and think of any other evidence that Amaril could realistically acquire. This evidence—mostly other sensory evidence—will not diminish the probability of H given the sensory evidence Amaril already acquired. After all, acquiring other evidence that disproves Amaril's having hands, after having acquired sensory evidence showing the presence of hands, is virtually impossible. Evidence that disproves having hands, however theoretically possible, does not belong to the reasonable set of evidential updates.

³⁵ These three desiderata require the denial of closure under known implication (its full version, at least) as well the denial of EQUIVALENCE CLOSURE while keeping CONJUNCTION CLOSURE. Why this package of theoretical moves is attractive has already been discussed in section 3.

The reasonable set was defined as “items of evidence Amaril could realistically acquire”. A probabilistic characterization can also be given, that is, “items of evidence whose probability of acquisition, given Amaril’s current evidence and her knowledge of the world, meets a certain, not extremely low, threshold”. The probability of acquiring evidence that would disprove having hands is not zero, but certainly, extremely low. All items of evidence in the reasonable set will thus confirm that H is highly probable. So, H counts as an item of knowledge given the proposed theory of knowledge as resiliently high probability.

Consider now the anti-skeptical proposition $\neg B$. In discussing evidential support, we have already seen that while sensory evidence increases the probability of H , it decreases the probability of $\neg B$. Sensory evidence, then, affects the probability of $\neg B$ in ways that are different from how it affects the probability of H . Since the changes in probability of H and $\neg B$ diverge, it follows that although the probability of H is resiliently high, that of $\neg B$ is not. If so, $\neg B$ does not count as an item of knowledge.

So far so good. Consider finally $H \wedge \neg B$. Here is where a difficulty arises. We know that the changes in probability of H must be paralleled by the changes in probability of $H \wedge \neg B$. So, it is not possible to distinguish the two propositions in terms of probability changes, at least insofar as we hold fixed the set of reasonable evidential updates. But suppose the reasonable set of evidential updates associated with H , call it U , is different from the reasonable set of evidential updates associated with $H \wedge \neg B$, call it U' . If so, Amaril could know H , insofar as H has a resiliently high probability relative to U , and not know $H \wedge \neg B$, insofar as $H \wedge \neg B$ does not have a resiliently high probability relative to U' . The question now is why there should be two reasonable sets of evidential updates.

Here is a possible answer. More items of evidence can challenge $H \wedge \neg B$ compared to H , because the former is a conjunction and the latter only one of the conjuncts. If a conjunction is more difficult to establish, it must be checked against a larger set of potentially contradicting items of evidence. Insofar as U' is larger than U , it is plausible to say that H is resiliently highly probable but $H \wedge \neg B$ is not. This is the result we want. Still, how are we to select different reasonable sets of evidential updates? The crucial question is whether probability theory can be of service here.

The bad news is that relevance, understood in a probabilistic way, is unhelpful. Any item of evidence that affects the probability of H will affect the probability of $H \wedge \neg B$ to the same extent. As far as probabilistic relevance is concerned, the two sets are the same. By contrast, considerations based on practical interests are helpful here. Arguably, more is at stake in establishing the truth of $H \wedge \neg B$ than in establishing the truth of H . The former denies a skeptical hypothesis, while the latter does not. If more is at stake with $H \wedge \neg B$ than with just H , the set of evidential updates is expected to be larger for $H \wedge \neg B$ than for H . But if different reasonable sets of evidential updates can be identified on the basis of what is at stake for different propositions, it is hard to see how probability theory can be of help in selecting them. So, the

resiliency-based strategy, if at all successful in solving the problem with the conjunction $H \wedge \neg B$, is likely to be a probabilistic strategy only in a marginal sense.

One might respond that the resiliency-based strategy can still retain its probabilistic character. Suppose we adopt a varying threshold for the probability of acquiring further evidence. This threshold determines how probable a piece of potentially contradicting evidence must be in order to be part of the set of reasonable evidential updates. The more is at stake, the lower such threshold. The lower the threshold, the larger the set of reasonable evidential updates. The larger the set of reasonable evidential updates, the harder to meet the knowledge requirement for resiliently high probability. With this machinery in place, the suggestion here is that the threshold for the reasonable set of evidential updates would be lower for $H \wedge \neg B$ than for H , and thus the reasonable set of evidential updates would be larger for the former than for the latter proposition. So H would count as knowledge but not $H \wedge \neg B$. While the criteria to decide on the threshold are non-probabilistic, the theory would still remain probabilistic.

But this response trivializes the resiliency-based strategy. If all we need is a varying threshold to determine when a piece of potentially contradicting evidence can count as part of the set of reasonable evidential updates, we might just as well get rid of the notion of resiliency. Having a varying probabilistic threshold for knowledge itself would give us all we need. Here is how. Even if the probabilities on the evidence of H and $H \wedge \neg B$ are the same, the threshold for knowledge would be higher for the latter than for the former, and thus H would count as an item of knowledge but not $H \wedge \neg B$. This would be, in essence, a contextualist account of why Amaril knows H , but not $H \wedge \neg B$.³⁶ It would be a contextualist account framed in a probabilistic language, but probability theory would do no significant explanatory work.

6 Ordinary scenarios

It is now time to see whether the difficulties for probability theory, which I have amassed in this paper, also apply to scenarios that do not involve far-fetched skeptical hypotheses. Here is an example inspired by standard bootstrapping counterexamples to reliabilism.³⁷ Suppose Elsa just looked at her well-functioning watch telling her that it is 3PM (abbreviated, *3PM*), formed the belief that it is 3PM, and this belief is true. So,

(E-1) Elsa knows *3PM*.

³⁶ Contextualism is, more specifically, a theory about the context-sensitivity of knowledge attributions rather than the context-sensitivity of knowledge itself. See, among others, [Cohen(1986)], [DeRose(2002)] and [Sosa(2000)].

³⁷ See [Vogel(2000)], [Cohen(2002)] and [Weisberg(2010)]. The example in the text is from [Sharon and Spectre(2013)].

If it is, in fact, 3PM, a watch telling you that it is 3PM is not wrong in that respect (abbreviated, $3PM \rightarrow \neg W$). The implication is true in virtue of the meaning of the words, and Elsa on reflection must know that. So,

(E-2) Elsa knows that $3PM \rightarrow \neg W$.

By the full principle of epistemic closure,

(E-3) Elsa knows $\neg W$.

While (E-1) and (E-2) are plausible, (E-3) is odd.³⁸ Elsa cannot know that her watch is not wrong just by looking at it and by thinking about the meaning of words.³⁹ Can probability theory offer an account of the oddity of (E-3)?

Consider first risk accumulation. The probability of $3PM \rightarrow \neg W$ must be one hundred percent, because the implication is true in virtue of the meaning of the words involved. If $P(3PM \rightarrow \neg W) = 1$, then $\neg W$ must have a probability that is at least as high as $3PM$. If the probability of $3PM$ is high enough to meet the threshold for knowledge, the probability of $\neg W$ must also be high enough to meet the threshold. Hence, risk accumulation does not offer the accouter we want.

Consider next the probability-based account of evidential support. Here a problem emerges with the conjunction $3PM \wedge \neg W$. If the introduction of evidence r increases the probability of $3PM$, as expected, the same evidence r must also increase, to the exact same degree, the probability of $3PM \wedge \neg W$. If r evidentially supports $3PM$, the account of evidential support tells us that the same evidence r must also support $3PM \wedge \neg W$. This is counterintuitive. How can reading one's watch support $3PM \wedge \neg W$ which asserts that one's watch is correct?⁴⁰

The assumption-based and resiliency-based strategies will face similar difficulties with the conjunction $3PM \wedge \neg W$. The assumption-based strategy rests on the same account of evidential support, so if r supports $3PM$, it will also support $3PM \wedge \neg W$. The resiliency-based strategy cannot distinguish between $3PM$ and $3PM \wedge \neg W$. If one is resiliently highly probable, so is the other. This holds provided the set of reasonable evidential updates is held fixed. If the set varies, the same difficulties discussed in the earlier section arise.

All in all, changing the putative counterexample, from one involving a far-fetched scenario to one involving a more ordinary one, leaves the difficulties unchanged.

³⁸ There is an ambiguity here. Proposition $\neg W$ should not be read as meaning that 'the watch, generally speaking, functions properly or is reliable', but rather, 'the watch, on this occasion, is not wrong about the time'. By merely reading a watch, you cannot know that the watch, on the occasion at hand, is right. Similarly, by merely listening to a witness testifying, you cannot know the witness, on the occasion at hand, is truthful.

³⁹ Watches, *typically*, tell the time correctly. So, Elsa's background knowledge warrants the belief that there is a high probability that her watch is correct. The oddity is to say that Elsa knows her watch is correct *by reading it*.

⁴⁰ In the section on evidential support, a way to circumvent this problem was discussed and set aside. The same considerations apply here.

7 Conclusion

Probability theory would seem to be a promising framework that can account for why epistemic closure is an intuitive principle and why it putatively fails in some cases. I have set out to show that this impression is mistaken. Probability theory falls victim of a number of difficulties. These difficulties are persistent. They cannot be eliminated by appealing to risk accumulation or a probability-based account of evidential support. They cannot be eliminated by postulating that certain propositions serve as background assumptions in all probability assignments or by appealing to a resiliency-based account of knowledge. To be sure, the assumption-based and resiliency-based strategy are promising. But their promise seems to stem from theoretical resources that lie outside probability theory. It is unlikely that probability theory can offer the account we need.

Acknowledgements For helpful comments and questions, many thanks to Dominik Klein, Olivier Roy, Gustavo Cevolani, Rahul Chaudhri, Justine Garson and an anonymous reviewer of this journal. I would also like to thank the participants in the CUNY Workshop for early-career philosophers, the Philosophy Breakfast at the University of Bayreuth and the 2016 conference of the Italian Society for Analytic Philosophy (SIFA). For financial support during my 2016-17 research leave, many thanks to Infosys Limited, the Institute for Advanced Study in Princeton and Lehman College of the City University of New York.

References

- [Achinstein(2001)] Achinstein P (2001) *The Book of Evidence*. Oxford University Press
- [Cohen(1986)] Cohen S (1986) Knowledge and context. *Journal of Philosophy* 83(10):574–583
- [Cohen(2002)] Cohen S (2002) Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research* 65(2):309–329
- [Cohen(2005)] Cohen S (2005) Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research* 70(2):417–430
- [Crupi et al(2008)] Crupi V, Fitelson B, Tentori K (2008) Probability, confirmation, and the conjunction fallacy. *Thinking and Reasoning* 14(2):189–199
- [Davies(1998)] Davies M (1998) Externalism, architecturalism, and epistemic warrant. In: Wright C, Smith B, Macdonald C (eds) *Knowing Our Own Minds*, Oxford University Press, pp 321–361
- [DeRose(2002)] DeRose K (2002) Assertion, knowledge and context. *The Philosophical Review* 111(2):167–203
- [Di Bello(2014)] Di Bello M (2014) Epistemic closure, assumptions and topics of inquiry. *Synthese* 191(16):3977–4002
- [Dretske(1970)] Dretske F (1970) Epistemic operators. *Journal of Philosophy* 67(24):1007–1023
- [Dretske(2005)] Dretske F (2005) Is knowledge closed under known entailment? The case against closure. In: Steup M, Sosa E (eds) *Contemporary Debates in Epistemology*, Blackwell
- [Harman and Sherman(2004)] Harman G, Sherman B (2004) Knowledge, assumptions, lotteries. *Philosophical Issues* 14(1):492–500
- [Hawke(2016)] Hawke P (2016) Questions, topics and restricted closure. *Philosophical Studies* 173(10):2759–2784
- [Hawthorne(2004)] Hawthorne J (2004) *Knowledge and Lotteries*. Oxford University Press

- [Hawthorne(2005)] Hawthorne J (2005) Is knowledge closed under known entailment? The case for closure. In: Steup M, Sosa E (eds) *Contemporary Debates in Epistemology*, Blackwell
- [Holliday(2015)] Holliday WH (2015) Fallibilism and multiple paths to knowledge. In: Gendler TS, Hawthorne JH, Holliday WH (eds) *Oxford Studies In Epistemology*, vol 5, Oxford University Press
- [Kripke(2011)] Kripke S (2011) Nozick on knowledge. In: S Kripke, *Philosophical troubles: Collected papers (Vol. 1)*, Oxford University Press
- [Lasonen-Aarnio(2008)] Lasonen-Aarnio M (2008) Single premise deduction and risk. *Philosophical Studies* 141(2):157–173
- [Leitgeb(2014)] Leitgeb H (2014) The stability theory of belief. *Philosophical Review* 123(2):131–171
- [Nozick(1981)] Nozick R (1981) *Philosophical Explanations*. Harvard University Press
- [Pryor(2012)] Pryor J (2012) When warrant trasmits. In: Coliva A (ed) *Wittgenstein, Epistemology, and Mind: Themes from the Philosohy of Crispin Wright*, Oxofrd University Press
- [Pynn(2013)] Pynn G (2013) The bayesian explanation of transmission failure. *Synthese* 203(190):1519–1531
- [Schaffer(2007)] Schaffer J (2007) Closure, contrast and answer. *Philosophical Studies* 133(2):233–255
- [Sharon and Spectre(2013)] Sharon A, Spectre L (2013) Epistemic closure under deductive inference: what is it and can we afford it? *Synthese* 190(14):2731–2748
- [Sharon and Spectre(2017)] Sharon A, Spectre L (2017) Replies to Comesaña and Yablo. *Philosophical Studies* 174(4):1073–1090
- [Sherman and Harman(2011)] Sherman B, Harman G (2011) Knowledge and assumptions. *Philosophical Studies* 156(1):131–140
- [Silins(2008)] Silins N (2008) Basic justification and the Moorean response to the skeptic. In: Gendler TS, Hawthorne J (eds) *Oxford Studies In Epistemology*, vol 2, Oxford University Press
- [Skyrms(1980)] Skyrms B (1980) *Causal Necessity*. Yale University Press
- [Sosa(2000)] Sosa E (2000) Skepticism and contextualism. *Philosophical Issues* 10:1–18
- [Vogel(1990)] Vogel J (1990) Are there counterexamples to the closure principle. In: Roth M, Ross G (eds) *Doubting: Contemporary Perspectives on Skepticism*, Kluwer, pp 13–27
- [Vogel(2000)] Vogel J (2000) Reliabilism leveled. *Journal of Philosophy* 97(1):602–623
- [Weisberg(2010)] Weisberg J (2010) Bootstrapping in general. *Philosophy and Phenomenological Research* 81:525–548
- [White(2006)] White R (2006) Problems for dogmatism. *Philosophical Studies* 131(3):525–557
- [Williamson(2000)] Williamson T (2000) *Knowledge and Its Limits*. Oxford University Press
- [Wright(2004)] Wright C (2004) Warrant for nothing (and foundation for free)? *Aristotelian Society Supplementary Volume* 78(1):167–212
- [Yablo(2014)] Yablo S (2014) *Aboutness*. Princeton University Press
- [Yablo(2017)] Yablo S (2017) Open knowledge and changing the subject. *Philosophical Studies* 174(4):1047–1071