Replies to Comesaña and Yablo

Assaf Sharon[1] · Levi Spectre[2]

Abstract There are few indulgences academics can crave more than to have their work considered and addressed by leading researchers in their field. We have been fortunate to have two outstanding philosophers from whose work we have learned a great deal give ours their thoughtful attention. Grappling with Stephen Yablo's, and Juan Comesaña's comments and criticisms has helped us gain a better understanding of our ideas as well as their shortcomings. We are extremely grateful to them for the attentiveness and seriousness with which they have considered our arguments and to philosophical studies for giving us this opportunity. Given the substantive difference between the two response papers, there is not much beyond sincere gratitude that we can covey to them jointly. We will therefore address them in turn.

1 Stephen Yablo

Stephen Yablo critically addresses our argument and sketches a different strategy for dealing with (apparent) cases of closure failure, building on the groundbreaking work in his recent *Aboutness* (2014). Naturally, our focus will be on points of disagreement, but this should not obscure our fundamental agreement and our great appreciation of his analysis. In fact, we uncomfortably find some of our views presented more lucidly in his article than they were in ours. Yet, in accordance with the standard "rules of engagement," besides a few brief comments, we restrict our reply to four points which we take to be the core of his criticism of our position.[1]

1.1 (ED)

Yablo introduces what he calls a "breakout scenario" which illustrates our argument against closure:

> Alma starts out not knowing p, and not knowing its consequence q either. She encounters evidence e that supports p but not q. The evidence is good enough that she winds up knowing p. Closure requires her also to know q. But e is irrelevant to q, or

Levi Spectre
levi.spectre@gmail.com

Assaf Sharon
assafsharon@post.tau.ac.il

[1] Department of Philosophy, Tel Aviv University, Tel Aviv, Israel

[2] Department of History, Philosophy, and Judaic Studies, Open University of Israel, Ra'anana, Israel
Associated Researcher at The Department of Philosophy, Stockholm University, Stockholm, Sweden

[1] We also have no objection and nothing substantial to add to his exceptional discussion of Hempel and related issues in the appendix to his paper.

> even negatively relative to it. Alma cannot come to know *q* on the basis of evidence pointing away from *q*! She gains knowledge that *p* but not that *q*, which is contrary to Closure. (Yablo, Sect. 3)

The argument here, Yablo notes, relies on the claim we labeled (ED). This is the claim that says:

> (ED) If A does not know at time t_0 that *q*, and the evidence she acquires between t_0 and t_1 counts against *q*, then A does not know *q* at t_1 either.

Yablo's first criticism of our argument targets this claim. Although appealing in Alma's case, Yablo claims that (ED) cannot be accepted as a general principle because "It makes no sense that e's claim to be regarded as good knowledge-conferring evidence depends on Alma not having earlier encountered highly seductive bad evidence." But this is what (ED) entails, as the following example shows:

> Let the hypothesis be *Jones own a Ford or Brown is in Barcelona*. Alma accepts it, initially, because of misplaced confidence in *Jones owning a Ford*. Brown writes, saying, "News of your mistake about Jones has reached me in Barcelona." Alma didn't know *h* to begin with, but she does now, since her confidence is well founded. Ah, but the new confidence may be slightly less confidence-inspiring than her ill-founded old confidence in Jones's owning a Ford. (ED) will in that case object, but it shouldn't. (Yablo, Sect. 3)

It is clear that a lot hangs on our (ED), and Yablo's objection to it give us a welcome opportunity to clarify some issues that should have been clearer in our original article. Yablo understands (ED) as an *epistemic principle* and target it by apparent counter-examples. However, as we were careful to specify "...(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" (Sect. 3.7). We went on to clarify in a footnote, that we do not regard (ED) as a general truth about knowledge: "Note that we do not regard (ED) as a principle. Nevertheless, it becomes evident that this claim 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." (footnote 44).

Nevertheless, we do take (ED) to be a true claim in standard cases that we relate to closure failure. It is true that we did not specify conditions for (ED)'s truth and we cannot specify them exhaustively here.[2] But we can identify one condition to which we gestured in the original article. (ED) is not true in cases where bad evidence is replaced by good evidence. That is to say, when high probability is based on misleading data taken as evidence for a truth, then lower probability on the basis of new, good evidence can facilitate knowledge. A good example of this is Gettier cases, of which Yablo's example is an instance (Alma's "misplaced confidence" in Jones' owning a Ford). The *typical cases* of (apparent or real) closure failure, which Yablo calls IONs, though, are not of this kind. At the initial stage, before the evidence is registered, one does not know that one is not a bodiless brain in a vat, not because this is false and not because one's belief in this statement is based on misleading evidence. In such cases, one's belief that evidence against *p*—which one happens to know—is not misleading, does not amount to knowledge not because it is based on bad evidence, but because it is, apparently, based on none. Thus, we have no quarrel with Yablo that one should be able to "live down early misadventures with the wrong disjunct," and when this happens—through replacing misleading evidence for a truth—it may be possible for knowledge to be gained despite a decrease in probability relative to some high misguided probability. Our point is that when this does not happen, as seems to be the case in typical IONs—when the incoming evidence does not dispense with older, misleading evidence—but merely lowers the probability of the proposition, knowledge isn't gained. Nevertheless, the new "good" evidence will have to count sufficiently in favor of the proposition in question. If not, then one who did not have the bad evidence would not have knowledge as well.

As we intended it, the example presented by Yablo, rather than undermine (ED) actually bolsters it. There is an identifiable piece of evidence that has come to light—Brown's telegram—between t_0 and t_1 which supports the proposition that Jones is in Barcelona. What (ED) says is that absent such new evidence, if all the evidence one gained supports the negation of some proposition, one doesn't (at least not typically) move from not knowing it to knowing it. To be sure, (ED) does not say that if one's rational credence with respect to something has decreased one could not have come to know it. Rather, the evidentialist thrust of (ED) is that typically (in non-Gettier, or other untypical cases of the kind we mentioned in the footnote quoted above) one cannot gain knowledge without gaining *some* evidence supporting the proposition one previously had not known. Our further claim—independent of (ED)—is that this is what happens in cases of closure failure, and can help explain them. Remembering where you parked your car gives you no evidence that it hasn't been stolen. Seeing a zebra-looking animal gives you no evidence that it is not a disguised mule, etc. In all of these cases, the *only* evidence gained is evidence that supports the negations of the propositions closure would have you know.

Why do we think (ED) holds typically? In typical ION cases, if you seek knowledge regarding q directly (where q is the proposition the knowledge of which is under dispute)

---

[2] We could slightly change (ED) so that Yablo's case wouldn't apply: (ED*) If A does not know at time t_0 that *q*, and the evidence she acquires between t_0 and t_1 counts against *q* only, then A does not know *q* at t_1 either. Since Alma gains some evidence against *q* but some evidence in favor of *q*, (ED*) wouldn't apply. The problem, however, with this type of claim is that it rules out many instances where some of the evidence counts in favor of *q* even though as a whole it counts against *q*. This is the case in many instances where it is clear that one gains no knowledge. See our comments below regarding Black winning the silver medal.

getting evidence *e* would not be any help. But closure is precisely the claim that if *p* entails *q* and *e* gets you knowledge that *p*, you do know *q*. So *e* is a way of getting to know that *q* after all (which, again, if we were looking at how *e* relates to *q* directly, we would agree that it isn't). This is what makes the cases we have in mind different from the Gettier type cases (or other possible cases where you get evidence that "defeats a defeater" and that lowers the probability that *q* relative to some former point in time). The fact that you get to know *p* by *e* plays no such Gettier or defeating a defeater type role.

A closer look at Yablo's own example of (ED)'s failure shows more directly why his criticism is mistaken. Think about Alma. Because her belief was based on "misplaced confidence in Jones owning a Ford," she does not know the disjunction Jones own a Ford or Brown is in Barcelona. But upon receiving Brown's telegram, Yablo says, "she does [know it] now." But does she? The evidence Alma gained by virtue of the telegram necessarily lowered the probability of one disjunct (*Jones owns a Ford*) and raised the probability of the other (*Brown is in Barcelona*). What she should believe now is that *Jones does not own a Ford and that Brown is in Barcelona*. Perhaps she even knows these truths. But it does not follow that she knows the original disjunction. Assuming that she must know it since she knows *Brown is in Barcelona*, is assuming the truth of epistemic closure, which in the present context would be begging the question.

In fact, the analysis that seems to be suggested by Yablo, would commit him to denying that Alma knows the disjunction. Yablo's *immanent closure*, "endorses Addition (If A knows that *p*, and competently infers $p \vee q$, A knows that $p \vee q$) only if *p* says in part that $p \vee q$" (Yablo, Sect. 9). But when "*q* introduce[s] matters that *p* says nothing about," Yablo says, knowledge of *p* does not entail knowledge of $p \vee q$. This is precisely Alma's situation. "Brown is in Barcelona" says nothing about Jones owning a car, and the telegram indicates that Jones does not own a Ford. So, while Alma can learn about Brown's location by reading his telegram, she gains evidence that *Jones does not own a Ford* and therefore does not know—not merely because it's false—*Jones owns a Ford* and, by Yablo's immanent closure, does not know the disjunction. Replacing bad evidence with good gives her knowledge of a disjunct she didn't previously know, but not necessarily knowledge of the disjunction, as epistemic closure would entail.[3,4]

Before proceeding to our second point, note that little hangs on our analysis above. Even if it's wrong and indeed (ED) is severely misguided, it isn't the case that it never correctly describes the epistemic state of agents (for instance in Yablo's Alma case that we repeat in the next section). Indeed if this were so we would lose the explanatory aspect of our argument, but all that is needed for our central argument against the *principle* of epistemic closure to work is that (ED) doesn't always fail.

1.2 Apparatus

Yablo's second main criticism is that our account has unacceptable consequences. This he tries to show by presenting two cases, one he says we'll like (and we do) and one that is consistent with our apparatus and that no one would like:

---

[3] In footnote 8 Yablo says: "Another way to see the problem with (ED). Evidence *e* that refutes *b* surely also refutes a hypothesis *c* that strictly entails *b*. (ED) cannot allow this when $c = e\&b$; *e* as a consequence of *b&e* cannot lower its probability. It seems like double-counting for *e*'s recurrence in a hypothesis to be what spares the hypothesis from refutation by *e*. The principle that suggests itself, letting *p–q* be what remains when *q* is extricated from *p*, is this.

(1) *e* is evidence for (against) *p* iff e makes *p–e* likelier (less likely).
Call that the remainder principle. Putting $e\&\neg h$ for *p*,
(2) *e* is evidence for (against) $e\&\neg h$ iff *e* raises (lowers) the probability of $(e\&\neg h)–e$.
Putting $\neg h$ for $(e\&\neg h)–e$,
(3) *e* is evidence for (against) $e\&\neg h$ iff e raises (lowers) the probability of $\neg h$."

It's hard to see why Yablo thinks that *e* surely refutes *c* if it refutes *b* that logically follows from *c*. If refutation of *p* is the confirmation of $\neg p$, this is a restatement of evidence closure (that he claims has been surrendered long ago). Even viewed as a principle, (ED) is about knowledge and its relation to evidence, (ED) is silent on the conception of the evidence-for relation. But even if we do engage with the suggestion that there is some double counting in his example, as we've shown, all that is required for refuting a probabilistic closed evidence principle is evidence that raises the probability of $\neg(e\&\neg p)$, we need not have *e* itself as evidence. Vogel's car being stolen from a certain place is more probable if memory of it being last parked there is his evidence. Yet it entails that the car is not where he last parked it.

The suggested principle is hard to evaluate. Suppose *e* is that the die landed on an even number and *p* is that it landed on 4. Now *e* is evidence for *p*, iff *e* supports *p–e*. But what is it to extract from 4 its being an equal number? We are not clear about the remainder principle, is the point.

[4] What we had in mind, and what we think makes (ED) hold in many cases, regards a comparison between what is sometimes called in Bayesian parlance Ur-Prior—the as-if confidence one has in a proposition without any evidence or background empirical knowledge—and the posterior probability. The first thing to note about Ur-Priors is that they are often understood as subjective prior guesses, or at least nothing like known propositions (even within a more objective framework). Second, the "as-ifness" allows us to disregard the "ups and downs" of confidence one has in a proposition due to biographical evidence-profiles that may include misleading evidence, undermining/undercutting evidence, etc. Some differences in biography can be abstracted away for the purpose of evaluating cases like those of Alma and her friends John and Brown. One way of seeing that these differences should not matter is checking to see if reversal of the time order of evidence Alma receives will make any difference. We will leave the excise of this comparison for another opportunity. We turn now to our second point.

> **Alma**: Alma watches Bolt win the 100 metre dash on TV, reads about it in the newspaper, etc. The evidence she gets from these sources does not address the issue of what will become of Bolt's refrigerated blood sample. Testing could reveal that Bolt had been using a banned substance. He would in that case be disqualified, with the result that he will not technically have competed or hence won the gold. Alma does know that Bolt is the Olympic champion, it seems. Does she know that blood tests won't be devised on the basis of which Bolt is disqualified? Surely not; she has no evidence on that issue.

> **Albert**: Albert learns from TV and newspapers that Bolt has won the gold and Blake the silver. The evidence he gets from these sources is all about Bolt. It does not address the issue of whether Blake won the silver. That is OK, though; knowledge that *p* does not require one to tick of all its necessary conditions *q*. Just as Alma needn't know about future blood tests, Albert needn't know about Blake. This is hard to make sense of. For Blake to win the silver is not one more necessary condition. It's part of what it is, we want to say, for Bolt to win the gold and Blake the silver. (Yablo, Sect. 4)

Yablo draws two conclusions from these cases: (A) "To know that *p*, it is not enough to know *some* of what it says; one must know all of it. This applies when *q* is a conjunct of *p*, but not in the [Alma] case, because *Bolt took the gold* only implies—it does not in part *say*—that Bolt won't be disqualified." (B) "Closure fails because of a factor"—Yablo continues—"that is present in Alma-type instances of the scenario, and absent in Albert-type instances. If a factor like that comes to light, we will want to investigate it further, to see how well its operations line up with the intuitive data on IONs generally." (Yablo, Sect. 4)

Yablo's footnote 10 brings A and B together to level a sharp argument against our view:

> Compare *Bolt won the gold* and *He won the gold and will never be disqualified*. These are a priori equivalent, so they ought to be equiknowable. The scenario does not bear this out. If to learn that Bolt won the gold and Blake the silver, Bina must know that Blake won the silver, then, it seems to me, to learn that Bolt won the gold and will not be disqualified, Alma must know that he will not be disqualified.

Why does it seem that Alma needs to know that Bolt won't be disqualified in order to know that he won the gold? There are three steps here: one is that *p* (*Bolt won the gold*) is a priori equivalent to *p&q* (*p&Bolt won't be disqualified*). Two (an application of A) is that to know *p&q*, one must know that *p* and one must know that *q*. Three, if we want to block the Albert case by claiming that he only knows that *p* but not that *r* (*Blake won the silver*) and therefore does not know *p&r*, we will have to also claim that Alma doesn't know *p*. Or more modestly (because later Yablo himself will question the equivalence principle), that if Albert doesn't know *p&r* Alma doesn't know *p&q* (which we claim she does).

Our view is indeed inconsistent with Yablo's A, because we think what matters with regard to KC (Yablo's knowledge closure principle) is the evidence-for relation. We do want to claim that Alma knows *p&q* because the evidence-for relation as we understand it (in line with the Bayesian confirmation framework as well as any view that accepts (EQ)) treats known a priori equivalent propositions on an evidential and epistemic par. So if we can't find a factor that is present in the Alma case that is absent in Albert's, we will have to agree with the absurd claim that Albert can know *p&r*.[5]

Note, first, that Yablo's proposal for distinguishing these two cases, faces serious theoretical challenges. If knowing p and knowing *q* is a condition of knowing *p&q*, then either Alma does not know *p*, or a priori equivalent propositions are not after all equi-knowable. Neither option is appealing in Alma's case. Denying Alma knowledge of *p* brings us dangerously close to skepticism. It seems that Yablo prefers the other option, namely the rejection of closure of evidence and of knowledge under equivalence. But if a priori equivalent propositions are not supported or refuted jointly they can fall prey to Dutch-book and to other arguments (like any view that diverge from the axioms of standard probability theory). We do not deny that rejecting equivalence sometimes fits intuitions more neatly, and that there are considerable advantages to Yablo's theory of *aboutness*. Our qualms are just with its theoretical consequences for epistemology, specifically with respect to evidence as understood not only within the Bayesian framework but also in non-probabilistic accounts (as we argued in our paper).[6] The evidence-for relation fares as bad as open-knowledge (or worse).

But whether or not Yablo's account has bad consequences does not relieve us of the argument he levels against us. So the second thing to note is that it isn't easy to see how the evidence in Albert's case could be all about Bolt and say nothing about Blake taking silver. Suppose there are three participants in a race: A, B, and C. The prior probabilities for winning are, e.g., as follows: $Pr(ABC) = 0.2$, $Pr(ACB) = 0.1$, $Pr(BAC) = 0.3$, $Pr(BCA) = 0.2$ $Pr(CAB) = 0.1$, $Pr(CBA) = 0.1$. With no additional background assumptions, getting evidence that, say, A is the winner, raises the probability of B winning the silver and of C winning the silver ($Pr(ABC|A)=Pr(ABC)/Pr(A)$ which is greater than $Pr(ABC)$, since $Pr(ABC \vee ACB) = Pr(A) = 0.3 < 1$). Now yes, this inevitability depends on the evidence being conclusive, while in Albert's example it is not. The point is, however, that evidence for Bolt taking the gold is often also evidence about Blake. But we shouldn't find this

[5] The problem Yablo notes about "for all Alma knows, *p*" (Yablo, Sect. 4) as opposed to "for all Alma knows ¬(*p&q*)" is a problem for knowledge failure of equivalence just as it is for closure generally. And this is true regardless of whether one has made inferences or drawn equivalences or is even in a position to do so (e.g., for all Alma knows the axiom of choice might be false has no worlds witnessing this possibility). A world witnessing Vogel's car being in the driveway is a world witnessing it's being there and not towed away. The point is, the possible world framework is ill fitted to capture any view regarding epistemic possibility that rejects "multi-premise closure" that in the present fallibilist framework is agreed to be invalid. It is hard to see how normal modal logic can be used even locally for any one proposition for a view that is fallible. Such a logic will entail inaccessible relations to possibilities that one has no evidence to rule out (and that are not ruled out a priori).

[6] Relatedly see John Hawthorne's argument along these lines in his (2004) against Herman and Sherman's denial of an epistemic equivalence principle (2004). The problem is worse for a deniers of evidence equivalence principles, e.g., because there isn't any foreseeable way of avoiding straightforward synchronic Dutch-Book situations (as there is by severing the tie between decision and knowledge).

surprising—evidential relations don't always line up with intuition (as we claimed above).[7] The problem for us, however, seems to have deepened. Intuitively Albert doesn't know that *p&r* even if he has *some* evidence counting in favor of *r* just by having (enough) evidence for *p*. But in the Alma case our view says that she does know *p&q* even though she doesn't know or have evidence for *q*. Yablo's challenge to us wasn't to show that Albert knows *p&r*, but to locate a factor that his case doesn't share with Alma's.

But if the equivalence principle for evidential support is valid (as we maintain it is), then there is a factor that is present in Alma's case and absent in Albert's. In fact it is the key factor of our account—evidential support. A conjunction's probability can't be greater than the probability of any of its conjuncts: $\Pr(q)<\Pr(p)\vdash\Pr(p\&q)\leq\Pr(q)$. In the Alma case, since $p\vdash q$, the probability of the conjunction will equal the probability of p ($\Pr(p\&q) = \Pr(p)$). The probability on the evidence $\Pr(p\&q|e)$ in Alma's case depends only on $\Pr(p|e)$. Thus, due to equivalence, there must be evidence for the conjunction when there is evidence for *p*. In the Albert case, on the other hand, the probability of the conjunction depends also on the probability of *q* (given the evidence). So it is not necessarily the case that the evidential support for the conjunction suffices for knowledge, even if it suffices for knowledge of *p*.

Having evidence in favor of a proposition, even if it justifies one's confidence in its truth enough for rational belief, does not necessarily amount to knowledge. George has a lottery ticket with one in a million chance of winning the draw. According to the more common view, he does not know he will lose, though his confidence that his ticket is not the winner is justified. Suppose an announcement is made that the lottery is sold out and the organizers decided to sell additional tickets. Presumably, this raises the probability that George's ticket will lose, but it does not necessarily provide George with knowledge that it will lose (assuming he will). Reading in the paper that his ticket lost, presumably does give him such knowledge, despite the fact that there is some probability that the paper is mistaken. This example, like lottery cases in general, shows that the probabilities alone do not determine the epistemic state. In both scenarios George receives evidence that presumably raises the probability that he will lose to the same degree, but only in the latter scenario and not in the former does he form knowledge. In the same vain, our evidential commitments entail that Alma has knowledge-conducive evidence for the conjunction, while Albert might not. So without denying that there are Albert-type cases, with the implicit assumption that he doesn't know r, we have a reason to attribute knowledge only to Alma.

1.3 Dogmatism

Perhaps Yablo's main criticism of our argument is the denial of our claim that evidence fails to transmit from premises to conclusion in the cases of apparent closure failure. "The problem," he writes, "is that evidential support can be understood so that Alma has it in these cases." Yablo illustrates this with respect to two types of cases—*dogmatism* and *easy knowledge*. Let us address them in turn.

The lesson that we draw from Kripke's dogmatism argument is that "having proper evidence that *p* is true can allow one to know *p*, but not ... that evidence against *p* is

[7] Intuitively, it seems that one cannot have evidence for *p&q* without having evidence for either *p* or for *q*. But arguments of the kind presented by Carnap and others show that this is unavoidable, as we elaborated in our paper.

misleading." But, Yablo claims, "the quick and easy way to find out whether anti-*p* evidence is evidence against a truth is to seek evidence directly on the question of whether *p* is true" (Yablo, Sect. 7). He illustrates this with the following example:

> Suppose we were wondering whether evidence against *p* was evidence against something funny. An analogue of the first approach would be to look, say, at the newspaper that gave us this evidence, for signs of bias in favor of funny falsehoods. That doesn't sound very promising. To determine if evidence against *p* is evidence against something funny, we should ask directly whether *p* is funny. Why should a similar strategy not sometimes work for truth? (Yablo, Sect. 7)

It seems to us that this argument doesn't address the issue we raised. Surely a good way of determining if evidence is misleading is to determine the truth or falsity of what it supports, but this can be done in different ways, only some of which will allow the inference that the evidence is misleading. Let as draw an analogy here in order to separate the proposition in question from the evidence for its funniness.[8] Suppose Alma tells you that Louis CK is not funny. To determine the veracity of her report, you can either attend one of his shows or you can ask Albert. In the first case, if you watch the show and see that CK is indeed funny, you have good evidence that Alma's report was misleading. But if you ask Albert and he says that Louis is very funny, you have counter-evidence to Alma's report, but, arguably, you are not in a position to know that her report was misleading (surely much will depend on the details and the particular reliability of Alma versus that of Albert, whether they give examples, etc. but we simplify for argument's sake).

The cases of closure failure that are relevant to our claim are akin to the latter scenario, not the former. Take Sorensen's (1988) parked car example. It is true, and we thank Yablo for pointing this out, that given equivalence knowing that your car is in the driveway also means that you know that any evidence that it is not there will be misleading. But closure failure of the sort we are defending also means that if all you have to go on is remembering that you parked your car in the driveway, you are not in position to know that Doug's report that the car was stolen is misleading. Your knowledge that the car is in the driveway is based on evidence that does not count against this possibility of theft and therefore does not allow you to dismiss counter-evidence of the kind issued by Doug. If, however, you know that your car is in the driveway not by virtue of remembering where you parked it but because you are looking at it, you know Doug's report is misleading. In other words, the epistemic implications of knowledge depend not on what the known proposition entails, but on what the evidence on which it is based supports. The evidence supports the truth of the generalized proposition but not all of its instances.[9]

The absurdity entailed by closure in the dogmatism cases can be connected to the problem of *easy knowledge*. Note that the dogmatist conclusion can be reached even without Albert's, or anyone else's, report. Assuming that Alma's original report facilitated

[8] Thanks here to Yablo for correcting a mistake we made.

[9] We thank Yablo for helping us clarify that our commitment to closure under equivalence entails a commitment to knowledge of general propositions of the type *evidence against p is evidence against a truth*, while rejecting knowledge of some particular propositions, such as *Doug's report is misleading*. This is another instance of closure-failure of the kind we argue for.

knowledge, you are in a position to know that her report is not misleading. We take it that this result of closure should be by anyone's lights, as Yablo stresses, at least an anomaly to be reckoned with. His way, we expect, would be to say that knowledge cannot be gained in this way because in and of itself Alma's report, while informative about Louis CK, tells you nothing about the credibility of Alma's report. True. But in the same vain, Alma's report tells you nothing about Albert's. Knowing directly that Louis CK isn't funny, on the other hand, gives you evidence about both. Without this direct access to Louis CK's funniness, you are not in a position to dismiss reports to the contrary. In both cases, evidence that supports a proposition does not support something entailed by it. Knowledge, we claim, fails to transmit in both cases, and for the same reason—lack of evidential support for the conclusion. Let us now turn to the issue of *easy knowledge*.

1.4 Easy knowledge

Perhaps the most telling difference between our view and Yablo's comes out in his argument against our easy knowledge analysis.

> A watch reading 3am is evidence for It is now 3am, but not, we are told, for *If the watch says it is 3am, then it is*. How better to establish that *Either my watch doesn't say so, or it is 3am* ($\neg e \vee p$), than to look at the watch? Seeing that it reads 3am tells me that the first disjunct ($\neg e$) is false, which means that q stands or falls with its other disjunct p: It is now 3am. By hypothesis, The watch reads 3am counts in favor of this disjunct. How can evidence for a hypothesis p that is known to agree in truth-value with q fail to be evidence for q?
> I am not questioning here that e makes q ($= e \rightarrow p$) unlikelier, or that e to confer knowledge should boost something's probability. What I am questioning is whether the something has to be q itself, as opposed to the disjunct p on which q's truth-value turns out to depend? (Yablo, Sect. 7)

Appealing though it no doubt sounds, this idea has some unwelcome consequences. If a disjunction can be known whenever the evidence raises the probability of "the disjunct on which its truth-value turns out to depend" then disjunctions or the form $\neg e \vee p$ can be known no matter what the evidence turns out to be. In other words, if seeing that the watch reads 3 a.m., is evidence for the disjunction $\neg e \vee p$, because it raises the probability of the disjunct on which the truth value of the disjunction "turns out to depend", then one has evidence for $\neg e \vee p$ no matter what the watch shows. If the watch shows some other time, then $\neg e$ is the case, and again the probability of the disjunct on which the truth value of the disjunction depends is raised (to 1). We know in advance, then, that we will have evidence for $\neg e \vee p$ no matter what the watch reads. The same is true for any other possible reading of the watch (i.e. for any other value of e and p). But this just means that before looking at the watch we already have knowledge-conducive evidence that the watch is telling the correct time.

This shows that if looking at the watch provides evidence for the truth of the disjunction, this cannot be simply because it raises the probability of p. In accordance with standard confirmation theory, we assume that to count as evidence of a proposition, the evidence must raise the probability of the proposition (in this case the disjunction). In fact, as we understand his view, Yablo shares this conviction—evidence must be about all of the proposition, not just about the part on which its truth value turns out to depend.

There is a further shortcoming to Yablo's proposal as compared with our's. The problem is most evident when presented with respect to conjunction introduction. If the constraint on knowledge by inference is that one know all the parts of what is inferred, then knowing p and knowing *q* one can know *p-and-q* simply by inferring it (after all there isn't any subject matter the conjunction operator introduces that wan't already known). Yablo's immanent closure principle, that is, seems committed to the closure of knowledge under multi-premise inferences such as conjunction introduction. But as familiar lottery and preface arguments show, this leads to untenable consequences. The evidentialist apparatus we propose is well equipped to handle such cases. Even if each and every claim in a book is well supported by the evidence, the conjunction of all of them may not be. Yablo would, at the very least have to draw on other resources in order to avoid these unhappy implications. In fact, we think he will have to draw on the very same evidence-knowledge relations that underly the brand of open knowledge we advocate. That is, he will need to take into account the probability of the propositions that he doubted were relevant for tracking evidential support relations in the context of *easy knowledge*.

Before moving on to address Comesaña's worries let us stress that for all we have said in support of our view as opposed to Yablo's, the differences are more on detail than on the general view of epistemic closure. We take ourselves to be on his side with regard to the central epistemological questions.

2 Juan Comesaña

Juan Comesaña's response includes several novel (at times revisionary) ideas about the evidence-for relation and how this relation should be understood with regard to knowledge. His response also includes a reconstruction and some developments of our main and secondary arguments. Let us follow his lead, then, and address several of the issues he raises in his insightful response.

Comesaña presents the general line of argument with two principles: evidence closure (EC) and knowledge closure (CP) (Sect. 1). The argument he ascribes to us runs as follows:

1. ''EC is false'' (:EC): The *evidence for* relation is not closed under deductively valid logical operations.

2. ''If EC is false, then CP is false'' (¬EC ⊃ ¬CP): If the evidence for relation is not closed under deductively valid logical operations, then the principle of knowledge closure is not valid.

"Therefore,

3. CP is false'' (¬CP): The principle of knowledge closure is not valid.

Although this conveys the gist of our main argument, it incorrectly suggests that we endorse premise 2. Comesaña rightly hesitates to attribute this premise to us since it is obviously false that ¬(EC)⊢¬(CP). What is, then, our main argument? Our central argument indeed contains the openness of the evidence-for relation. The following example displays its essentials and makes clear why we do not rely on ¬(EC)⊢¬(CP): Suppose Fred wants to find a mule that looks (or is made to look) like a zebra. He is at the very beginning of his search and so he first wants to find out where the zebra-looking animals are or are expected to be. He decides to visit a nearby zoo, but at this initial stage he isn't even informed about whether or not there are zebra-looking animals in this local zoo (including zebras). Presumably, Fred knows neither whether there is a zebra at the zoo, nor whether there is a zebra-looking mule there. Upon entering the zoo he sees an animal that looks like a zebra. It seems that his confidence that his search has been successful should increase at least somewhat. Seeing a zebra-looking animal provides Fred evidence that there is a zebra at the zoo, but—since he can't tell zebras from zebra-looking mules—it also gives him evidence that there is a zebra-looking mule there too. The natural thing for Fred to do is look more closely, ask experts, etc. It would be unreasonable for Fred to forego all these examinations and take himself to know that the animal in the pen is not a disguised mule simply by inferring it from the proposition *there is a zebra in the pen*. The evidence he got (that supports the possibility that there's a zebra-looking mule at the zoo) is precisely what Fred was looking for. It would be absurd for him to call his search off on that basis alone.

Our argument is that the total evidence that allows for knowledge that p need not support propositions that follow from p. Because making an inference doesn't add evidence, knowledge, we argue, is either infallible—contrary to the framework within which we are working—or the evidentially unsupported propositions are known a priori, i.e. without evidence. For the latter to be a full response to our argument (as we claimed above with respect to Yablo's watch case and as we argue below regarding the rationalist account of justification) is for it to be a variant of the former infallibilist conception of knowledge.

Let us further clarify that we do not claim that Fred's belief that the animal in the pen is not a mule disguised to look like a zebra is not justified, or that his failure to know this is due to the absence of justification for this belief. Our claim is that he gained no evidence for this claim via the evidence he gained supporting the proposition that there is a zebra at the zoo. Fred's belief, even if justified, that there is no disguised mule does not become knowledge on the basis of the evidence he received by looking at the animal in the pen. Beliefs may be justified in the absence of evidence, but they do not become knowledge (at least not typically). This is the claim, which we labeled (ED), we take to hold at least for most Fred-like cases.[10]

Again, we are not arguing from lack of rational confidence, nor are we inferring directly from the openness of evidence to the openness of knowledge, not even via the assumption that evidence is necessary for knowledge or with the added assumption that skepticism is false. The openness of evidence does not establish, on our account, that knowledge is not

---

[10] What we mean by Fred-like cases is cases in which one's evidence is not conclusive (i.e. does not entail the proposition it supports) and when one does not have prior knowledge of the relevant entailments of what one comes to know on the basis of the evidence. We are not wedded to any of the examples we use to demonstrate these features, only to the existence of such possibilities on any fallibilist account of knowledge, which can be proved as we show in our article (footnote 43). Our argument, then, is structural, not one that is based on cases. That the cases in the literature also fit the structural features of knowledge we highlight—in particular Yablo's ION's—gives added support to our conclusion.

closed, but that evidence is lacking in the cases we investigate. We provided a detailed example—the watch case (Sect. 2.1)—that shows how one gains evidence for a proposition while gaining no evidence for propositions that follow from it. This together with the assumption that evidence is—in these cases—necessary to go from ignorance to knowledge of a proposition (because if we wanted to know these propositions directly we would look for different kinds of evidence) creates a formidable *challenge* for knowledge closure. To be clear, this alone does not strictly entail that knowledge is not closed (as Comesaña thinks we suppose), but it puts serious restrictions on possible defences of closure (which we address in the article and will touch on here in the last section).

We do not think, then, that the entailment of premise 2 holds. Rather, we *argue* from ¬(EC) to ¬(CP)—as the Fred example was meant to show in a different way—by pointing to a dependence on evidence that we identify in cases of closure failure (ED). In addition, our argument also assumes that (some) knowledge is fallible. We return to Comesaña's more specific argument against premise 2 later, but first we will focus on his more detailed response to premise 1.

2.1 Premise 1: evidence openness

Comesaña presents a very concise reconstruction of our argument against (EC) using principles that he claims are weaker than ours. In general we have no problem with his way of arguing against (EC), but let us address a few critical claims that he makes along the way and present one shortcoming of his version of the argument.

Our goal in presenting the argument from basic principles against the idea that the evidence-for relation is deductively closed (i.e. against EC) was:

i. To present a non-probabilistic argument that would make clear that theoretically unappealing measures (besides rejecting the standard probabilistic understanding of the evidential support relation) are required to defend evidence closure.[11]
ii. To mirror John Hawthorne's argument in favor of knowledge closure and show why and where it won't hold for evidence.

Comesaña's first comment is that there are weaker principles than the ones we use that entail the same conclusion, namely the failure of evidence closure (EC). Even if Comesaña is right, since he uses weaker principles his argument does not mirror Hawthorne's argument and therefore will not satisfy our objective ii. What about purpose i? Here things are a bit complicated. On the face of it, between two reductio arguments that have the same conclusion, the one with the weaker principles is preferable. Nevertheless, the weaker premises might not fit all the cases we are interested in. Compare Comesaña's argument to ours. The first step is not problematic:

> Comesaña employs (Underdetermination): It is possible for *e* to be evidence for *p* even if *e* doesn't entail *p*;

[11] Jonathan Vogel defends (in an unpublished paper) an evidence-for relation based view that is precisely the kind of non-probabilistic view aimed at preserving evidence closure that our argument was meant to address.

whereas we employ (UD): It is possible for *e* to support two (or more) inconsistent hypotheses.

Clearly Underdetermination follows from our (UD) principle. Indeed it would be harder to endorse the claim (although surprisingly Comesaña himself seems to) that *e* cannot support a proposition *p* unless *e* entails *p*, than it is to reject (UD). That is to reject the idea that evidence supports, or at least can support, two (or more) inconsistent hypotheses. So far so good, then.

The second premise Comesaña uses is:

(Entailment): If *p* entails *e*, then *e* is not evidence for ¬*p*.

ours is this:

(CS) if e supports *p*, *e* does not support ¬*p*.

Suppose a theorist Hakim claims that there's a case x where *p* entails *e* and is evidence for ¬*p*. Case x is a counter instance to Entailment but Hakim can still endorse (CS) if he doesn't accept the following principle:

(*) if *p* entails *e*, *e* is evidence for *p*.

It is true that (*) seems like (and we believe it to be) a valid principle. But Hakim (who claims (Entailment) has counter instances) would reasonably claim that x is a counter instance. More specifically, though (*) may be considered one of the basic principles of hypothetical deductivism, Hakim is claiming that x shows it to be invalid. He need not go as far as Popper in claiming that there is no support relation at all (or only conclusive support for refutation), he can consistently maintain that (CS) is valid while (Entailment) is not (because (*) isn't). In other words, (CS) does not entail (Entailment) unless supplemented by something like (*).

Our argument from (CS) is a reductio designed to capture any conception of evidence that can be plausibly thought of as workable. Entailment and (*) are silent with regard to cases where *p* does not entail *e*. *I remember parking my car in the driveway* does not follow from *the car is in the driveway*. Nor does the claim that *the wall is blue* entail, at least without further premises, that *the wall looks blue*. Our purpose was to show that any plausible non-probabilistic (perhaps qualitative) evidential account must either give up on the idea that evidence is deductively closed or give up on (CS). Rejecting (CS) is a heavy theoretical price for any conception of evidence, a price that we think no plausible conception of evidence can afford to pay. (Entailment) is not as difficult to give up.[12]

Up to this point, our differences with Comesaña regard details: (1) we think our argument applies more widely, i.e. not only when the hypotheses entail the evidence; (2) our argument reflects Hawthorn's argument with regard to knowledge; and (3) even for those who are skeptical about Comesaña's (Entailment) principles pay a heavy price in giving up our (CS). We now turn to his suggestion to reject (ED), or even more radically, reject his Underdetermination principle (that e can evidentially support p while $e \nvdash p$).

The idea that one can know that a proposition is true only if one has evidence that entails that proposition, is as Comesaña says, not an idea we spend much time on in our paper. This idea is perhaps most famously defended by Timothy Williamson.[13] The reason we don't say much about this idea is that our paper is meant to focus on closure within a fallibilist framework for knowledge (this is why we focus on single-premise closure).

Comesaña goes further than infallibilism about knowledge and claims that it can be defended with regard to evidence. This idea is quite radical, if adopted it would have serious ramifications with regard to rationality of belief and scientific practice. Take rational belief. If evidence is factive and Underdetermination is false, there are no rational false beliefs. Indeed, Williamson has recently suggested just such a view,[14] but even he does not think that evidential support is a relation of entailment. Quite the opposite, he claims that though false beliefs are never fully rational, one has an excuse for believing them, which partially depends on the existence of compelling, though misleading, evidence. Comesaña's rejection

---

[12] Comesaña also thinks that our argument that we claimed does not appeal to (CS) tacitly appeals to his (Entailment) principle. Here is his presentation of our argument:

> Sharon and Spectre suggest at one point that they can do without anything like Entailment. I reconstruct their argument as follows. Suppose that our evidence is given by an atomic proposition a and that it supports two alternative theories: $(a \wedge b)$ and $(a \wedge \neg b)$ (where b is another atomic proposition). Then, by Closure, a also supports the following proposition: $[(a \wedge b) \vee \neg(a \wedge \neg b)]$. But that proposition is logically equivalent to its second disjunct: $\neg(a \wedge \neg b)$, which in turn is logically equivalent to $\neg a \vee b$. Therefore, assuming that logically equivalent propositions are supported by the same evidence, a supports $\neg a \vee b$. But that last claim, Sharon and Spectre say, is "absurd." Now, why would it be absurd to say that a supports $\neg a \vee b$? It is, of course, incompatible with Entailment, but if we are not assuming Entailment or anything like it, I do not see how Sharon and Spectre's argument goes through. (Comesaña, Sect. 2)

Comesaña is correct that without further principles we cannot derive a contradiction. The reductio argument, however, was meant to show an absurd conclusion from (EC). The absurdity does not stem from the claim that the evidence is entailed by the hypothesis $a \wedge b$ and $a \wedge \neg b$ and hence should not support their negations $\neg(a \wedge \neg b)$ and $\neg(a \wedge b)$. The absurdity stems from two things: (1) that a and b are atomic propositions, and (2) that a is silent with respect to b. The absurdity is that, if evidence is deductively closed, a supports $\neg a \vee b$ although it is silent on b.

[13] See Williamson (2000). For our arguments on this issue see our (2013).

[14] Interestingly Comesaña, along with Stewart Cohen, have advanced powerful arguments against this idea. See their (2013a, 2013b).

of Underdetermination is, therefore, highly revisionary in that it entails that false beliefs are never even supported by evidence.

If Comesaña's suggestion about evidence is correct, then it turns out that many scientists are confused about how they should view their stance with regard to the theories and hypotheses they explore and advocate. If the evidence for evolutionary theory in biology does not entail evolutionary theory, we are all mistaken in thinking we have good evidence for it. The same holds for Quantum mechanics, its interpretations, (special) relativity, etc. In fact, it's hard to think of a theory that is logically entailed by the available evidence. So if Underdetermination is false, scientists are either mistaken about the quality of their evidence or wrong to think they have evidence for most if not all of their theories.[15] Moreover induction in science becomes particularly problematical. At no point until the evidence is conclusive does one have any evidence for the inductive conclusion no matter how much inductive evidence one gets. Having no evidence at every step it is hard to justify looking for further evidence for a hypothesis rather than its negation.

Perhaps what Comesaña is suggesting is a sufficiency condition: unless one's evidence entails a proposition, one does not have sufficient evidence to rationally believe it. This idea —similar to the one Williamson defends—does not carry the objectionable consequences we mentioned. But it also does not contradict our argument (even in Comesaña's version of it), which pertains to incremental evidence, not to a sufficiency notion of evidence.

2.2 Premise 2: from evidence openness to knowledge openness

As we explained above, premise 2 is not something we appeal to as a logical or a priori truth. Rather, we argued—as we illustrated with the example of Fred—that knowledge is open via our claim that in order for one to know a proposition one typically needs evidence for it. Since closure is a principle, its truth depends on there being not a single exception. Yet, for every item of knowledge that is not entailed by one's evidence there are many propositions that follow from the known proposition for which one has no evidence (our footnote 43). So for closure to fail it suffices that there is one case in which evidence is required for knowledge of the entailed proposition and yet it is lacking. We argue that this is the case in the standard examples in which closure seems to fail (Yablo's IONs) and provide a detailed example that shows the evidential relations that are involved. The fact that our intuitions with respect to closure in these cases fit the theoretical analysis of their evidential features, lends further support to our account. Comesaña in fact agrees that in such cases knowledge (if it obtains) is not evidence based. His question is: "why must my rational confidence in the conditional be evidence-based in order for me to know it? Isn't it enough that it is *rational*?" (Comesaña, Sect. 3). This amounts to questioning the claim that (ED) is typically true (at least for ordinary empirical knowledge).

But mere rationality is certainly not sufficient for knowledge. Comesaña's comments suggest that it is sufficient in the special case of what following Dretske he calls "metaphysically heavyweight propositions." But, as John Hawthorne convincingly argued,

[15] The only way out seems to be ascribing to scientists conclusive evidence that maybe theory T is correct, that it is highly probable that T is correct, or something of this sort. However, this would require showing that they have evidence that entails these kinds of propositions. But this is highly doubtful for the same reasons that conclusive evidence is doubtful in the first place.

there is nothing metaphysically special about these propositions and there is no plausible way to define heavyweight propositions in a way that would capture all and only the propositions targeted.[16] Metaphysically, my car has not been stolen, the table is not white under red lighting, etc. seem perfectly ordinary. What makes them unique in Comesaña's view is their epistemic role as "gap-fillers between the evidence and what the evidence justifies the subject in believing." But if all the gaps are filled, then evidence is essentially conclusive and we find ourselves in an infallibilist framework (see a more careful argument for this claim in footnote 54 of our paper).[17]

Comesaña proposes the neo-rationalist framework recently advanced by Cohen (2010) and Wedgwood (2013). Cohen and Wedgwood argue, along with Roger White (2006) and others, that some empirical beliefs must have a priori justification. We have no qualms with this view. In fact, since such neo-rationalism pertains to justified belief, it aligns quite well with our position, which holds that rational belief is closed under single-premise entailment. Since such beliefs do not have evidential justification, the idea that they are justified a priori sits comfortably within our framework.[18] The question is whether they amount to knowledge. Even if in the special cases neo-rationalists are concerned with (the reliability of our perceptual faculties) they do, there is no special reason to think this holds for every proposition that is entailed by ordinary empirical propositions but not supported by their evidence (e.g., that memory is a priori reliable can hardly be reworked into knowledge that cars haven't been stolen). If it does, then we seem to be forced out of the fallibilist framework, and then also saddled with multi-premise closure and its implausible implications (regarding not merely knowledge but also rational beliefs). If not all entailed propositions which are not supported by the evidence are known a priori, then the anti-closure argument will hold with respect to them.[19] The choice is between rejecting closure and giving up either non-skepticism or fallibilism. This indeed is what our original argument was meant to show.

Perhaps some hybrid theory will work: some of the evidence gap is filled by a priori knowledge, some propositions are known only relative to certain contexts, interests, etc. While this kind of idea will bring us closer to closure, it's hard to see a non-question begging reason to assume that it would close off all possibilities that leave knowledge open.

---

[16] See Hawthorne (2005).

[17] If all the "gap-fillers" are known, then the conditional probability of each entailed proposition will be 1 (see Sect. 1.4 above). For an additional argument against the a priori account see footnote 54 in our paper.

[18] Another possibility, highlighted by Cohen, is to deny that we need justification for believing that our perceptual faculties are reliable in order to be justified in believing their deliverances.

[19] This is what holds when we think about the standard role of a priori knowledge in conditional probability. As far as we know, there isn't any worked out theory regarding such matters. The challenge in developing such an account and remain within fallibilists is to say, first, why the entire "evidential gap" is to be filled by a priori knowledge, and, if so, how is it different from other a priori knowledge (regarding its probability, its role in conditionalization etc.) that allows the probability of known propositions to be <1.

Acknowledgments Levi Spectre's research was supported by the Israel Science Foundation (grant no. 463/12).

References

Cohen, S. (2010). Bootstrapping, defeasible reasoning, and a priori justification. *Philosophical Perspectives, 24*(1), 141–159.

Cohen, S., & Comesaña, J. (2013a). Williamson on Gettier cases and epistemic logic. *Inquiry, 56*(1), 15–29.

Cohen, S., & Comesaña, J. (2013b). Williamson on Gettier cases in epistemic logic and the knowledge norm for rational belief: A reply to a reply to a reply. *Inquiry, 56*(4), 400–415.

Harman, G., & Sherman, B. (2004). Knowledge, assumptions, lotteries. *Philosophical Issues, 14*(1), 492–500.

Hawthorne, J. (2004). Replies to Harman, Sherman, Vogel, Cohen. *Noûs, 38*, 510–523.

Hawthorne, J. (2005). The case against closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology*. Malden, Ma: Blackwell.

Sharon, A., & Spectre, L. (2013). Epistemic closure under deductive inference: What is it and can we afford it? *Synthese, 190*(14), 2731–2748.

Sorensen, R. A. (1988). Dogmatism, junk knowledge, and conditionals. *Philosophical Quarterly, 38*(153), 433–454.

Wedgwood, R. (2013). A priori bootstrapping. In A. Casullo & J. Thurow (Eds.), *The a priori in philosophy* (pp. 226–246). Oxford: Oxford University Press.

White, R. (2006). Problems for dogmatism. *Philosophical Studies, 131*(3), 525–557.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.