Inferential Knowledge, Counter Closure, and Cognition

Abstract

In previous work we argued that alleged counterexamples to counter closure for knowledge fail and should be construed as cases of knowledge despite (rather than from) falsehood. On our view, what really grounds our belief in the mentioned examples are tacit or subconscious inferences from tacitly or subconsciously known propositions. In this paper we develop this view further by embedding it in the psychological literature on implicit reasoning. As will emerge from the discussion, our view is firmly supported by recent results in cognitive psychology. In the middle sections of the paper, we address recent criticisms of our view, and discuss an interesting example by Roy Sorensen. We finally conclude by formulating a wider outlook for the current debate on counter closure and knowledge despite falsehood, including a discussion of the Socratic method of Elenchus.

Introduction

It has been argued that it is possible, in certain circumstances, for an agent to come to know something by inference from premises that are not themselves known—perhaps because one or more of them is false. For instance, according to Warfield (2005), one might come to know that one has sufficiently many handouts to distribute to the audience members at one's talk if one miscounts the latter slightly, provided the excess number of the former is sufficiently great—in such a case, he claims, we are able to obtain knowledge from falsehood (KFF). However, we have elsewhere upheld the principle of counter closure for knowledge (Ball and Blome-Tillmann, 2014), according to which, if a subject knows a proposition on the basis of competent deduction from one or more premises, the subject must also know those premises. In cases like that of Warfield's handouts, we do not dispute that the subject comes to know the conclusion—but we contend that this piece of knowledge is secured on the basis of inference from other tacitly, or subconsciously known premises, rather than from those that are explicitly and consciously considered, including the crucial falsehood.

In this paper, we further develop our view. We begin by relating it to the recent literature from cognitive psychology, showing that our appeal to tacit belief and implicit reasoning is quite uncontroversial. We then show how this empirically grounded perspective can be applied in certain problem cases, thereby responding to objections raised by Luzzi (2019). Next, we discuss an interesting example due to Sorensen (this volume). We conclude with some broader reflections on the debate.

Inference and Implicit Reasoning

What is inference? Recent theoretical characterizations vary widely. For instance, according to Ludwig and Munroe (2019: 20), inference involves "a transition from one set of propositional attitudes (e.g., beliefs, intentions, suppositions) to another"—though not every such transition is an inference; crucially, the transition must be—or at least be capable of being brought—under the agent's conscious control. Thus, on their view, inference is a personal-level act (involving the propositional attitudes of the subject), that is potentially conscious. By contrast, Quilty-Dunn and Mandelbaum (2018, 2019) allow that inference may be unconscious and even sub-personal (e.g., taking place within the visual system); but at the same time, they insist it is a formal, rule-governed transition. This further requirement strikes us as potentially disastrous:

the central lesson of Goodman's (1955) new riddle is that no formal logic of induction is possible; yet we do not want to preclude that some inference is probabilistic in character!

We prefer a characterization of inference on which it is a transition between content-bearing states that is sensitive to rational relations between those contents, however that sensitivity is achieved. Thus, our view allows inference to be tacit/subconscious without insisting it is formal, and/or that it involves rule-following. Ultimately, though, it does not matter to us whether such transitions are properly called 'inferences'—we can set such terminological issues to one side. What matters is that transitions of the kind in question (a) are widely appealed to in the cognitive sciences, and (b) are what underpin knowledge despite falsehood (KDF) in our view. In the next section, we illustrate claim (b) in connection with some cases in the literature on knowledge from/despite falsehood; but we begin, in the current section, by providing evidence of claim (a).

As Rescorla (2019) notes, the notion of unconscious inference has been appealed to in perceptual psychology since it was introduced by Helmholtz in the late 19th century. More recently, specifically Bayesian models of such inference have been widely deployed, not only in the science of perception, but also in relation to phenomena that are clearly cognitive in character (cf. Rescorla, 2019: 41). In short, Bayesian cognitive psychologists hypothesize tacit or subconscious *reasoning*.[1]

Another tradition of theorizing within recent cognitive psychology has been the two systems approach, popularized as *Thinking, Fast and Slow* by Kahneman (2011). System 1 thinking is characterized as fast, automatic, and unconscious, while System 2 thinking is described as slow, voluntary, and conscious. Like Bayesianism, two-systems theory has been widely deployed: for instance, in debates on mindreading (Apperly and Butterfill, 2009), delusion (Bongiorno and Bortolotti, 2019), numerical cognition (Graziano, 2018), and elsewhere—with the obvious consequence that unconscious inference has been postulated in this wide range of domains.[2]

Of course, we are not committed to any particular theory of implicit reasoning—thus, we need not endorse either Bayesianism or two-systems theory specifically. We only rely on the more general claim that people engage in specifically cognitive processing that involves transitions between attitudinal states manifesting rational sensitivity to their informational content and that is nevertheless unconscious. And there is a broad consensus in the literature on the existence of implicit reasoning as relied upon by us. For instance, prominent critics of two-systems theory, such as Kruglanski and Gigerenzer (2011), share the view that there is intuitive or implicit

[1] One particular version of Bayesianism that has been attracting attention recently is that of predictive processing (Clark, 2013; Hohwy, 2013). This particular theoretical approach, on which Bayesian reasoning is used specifically to minimize prediction error, has proven controversial—see, e.g., Cao (2020) for a recent critical discussion. But the more general Bayesian view, on which subjects are in states that can be characterized by probability distributions on variables, and updating from one time to another is undertaken in accordance with Bayes' rule, has been *much* more widely adopted.

[2] Our preferred theoretical approach (cp. Evans and Stanovich 2013) is one in which rapid autonomous processes (Type 1) are assumed to yield default responses unless intervened on by distinctive higher order reasoning processes (Type 2). Evans and Stanovich go on to give three kinds of evidence in favour of two types of reasoning (autonomous vs decoupling/working memory): experimental manipulations; neuroscientific; and individual differences. (They also contrast their preferred default-interventionist approach—shared with Kahneman—with parallel-competitive accounts.)

reasoning—they just account for it differently than two-systems theory. According to them, there is a single, rule-governed system that underlies both intuitive and deliberate judgment. Our reliance on implicit or intuitive reasoning is thus far from controversial in the current literature in cognitive psychology.

This discussion of implicit reasoning allows us to substantiate the view we took previously, by specifying more clearly what we had in mind by saying that subjects in cases of KDF implicitly rely on known premises: their knowledge in these cases is supported, in our view, by implicit reasoning from tacitly known premises. In the next section, we respond to some recent objections to our view (due to Luzzi (2019)), thereby showing the advantage of our reliance on the existence of actual causal processes involving implicit beliefs (rather than dispositions towards possible beliefs and inferences).

Response to Luzzi

In his (2019) book, *Knowledge from Non-Knowledge*, Federico Luzzi discusses views, like ours, according to which there is knowledge despite falsehood (KDF) in cases such as that of Warfield's handouts, because the known conclusion is doxastically justified (or "epistemized", as Luzzi (2019: 10) puts it) by knowledge of a true "proxy premise" that is "in the neighbourhood"[3] of the consciously and explicitly believed falsehood. The idea is that such views can navigate between the Scylla of Gettier cases (where we do not want to say that the conclusion ostensibly believed on the basis of a false premise is known), and the Charybdis of admitting that there is knowledge from falsehood, and that the principle of counter closure for knowledge is false. Luzzi, however, is sceptical: "the ultimate effectiveness of this strategy," he says, "remains dubious" (2019: 13).

Unfortunately, although he notes that we are advocates of a view of this kind (2019: 13, fn. 3), he does not consider our published views explicitly in any detail, preferring to focus on Montminy's (2015) development of the position. This matters because, not only was our view developed independently of Montminy's—despite Luzzi's (2019: 22, fn. 8) claim that we "endorse" Montminy's strategy in at least one crucial respect, in fact we do not cite, and had not read Montminy's article at the time of writing our (2014) article—it differs from it in at least one crucial respect. Whereas Montminy appeals to *dispositional* belief, we speak of *tacit* belief. Crucially, dispositional belief is (at least roughly) a disposition to consciously judge, whereas tacit belief is not a disposition of any kind:[4] that is, it is not a potentiality but an actuality; thus, it can be subconscious but causally active.[5] As we will see, this allows us to respond to the worries Luzzi raises for the general approach.

To begin with, consider Warfield's *Handout* case of alleged KFF.

> *Handout*
> Counting with some care the number of people present at my talk, I reason: 'There are 53 people at my talk; therefore my 100 handout copies are sufficient'. My premise is false. There are 52 people in attendance—I double counted one person who

[3] Luzzi (2019: 10).
[4] Of course, it may *ground* various dispositions, including the disposition to judge. But this is not to say that it *is* a disposition.
[5] For instance, it may figure in implicit reasoning.

> changed seats during the count. And yet I know my conclusion. (Warfield 2005, pp. 407-408)

What is Montminy's account of this case? According to Montminy, Warfield's conclusion that he has enough handouts is based on a true, dispositionally known proposition—namely, the proposition that there are *approximately* 53 people at the talk. According to Luzzi (2019: 17-8), however, that proxy proposition is essentially evidentially based on the false premise that there are exactly 53 people at the talk. Luzzi's response to Montminiy's line of reasoning is thus straightforward: he points out that, if Warfield's belief that he has enough handouts is indirectly still based on the falsehood that there are exactly 53 people in attendance, then his belief still qualifies as knowledge from falsehood. The only difference being that, on Montminy's account, it isn't *directly* based on the relevant falsehood, but it remains based on it, nevertheless.

As we showed in our original paper, this difficulty can be avoided fairly easily. Here is the view we endorsed in 2014:

> *Knowledge Despite Falsehood (KDF)*
> In apparent cases of KFF, there are two true propositions t_1 and t_2 such that:
> (1) t_1 evidentially supports both p and t_2 for S;
> (2) t_2 is entailed by p;
> (3) S knows both t_1 and t_2;
> (4) S*'s* belief that q is properly based on her knowledge that t_2.

To see how this account handles the above case, note that the variables in *Handout* take the following values:

> t_1: The result of my (Warfield's) count was '53'.
> t_2: There are 53 people at my talk *give or take a few*.
> p: There are 53 people in the room.
> q: I have enough handouts.

In Warfield's example, t_1 is clearly known by the subject: it also evidentially supports or justifies both t_2 and p to at least some degree; and t_2 is entailed by p. Finally, note that the truth t_2 is a rough-and-ready approximation of the falsehood p, which is, in addition, known in the examples. Moreover, since Warfield's belief that q is properly based in his knowledge that t_2—namely, by competent deduction—it follows that the examples are not cases of knowledge from falsehood but rather cases of *knowledge despite falsehood*.

What is crucial about this account is that, contrary to Luzzi's assumption, we do not claim that Warfield's belief that t_2 is derived or epistemically based upon his false belief that p. Instead, we take, as Sorensen (this volume, n. 8) notes, the far more plausible view that Warfield's tacit belief that t_2 is based on his knowledge that t_1.[6] Thus, Luzzi's objection that Warfield's belief is *indirectly* based on the falsehood p misfires as an objection to our view. In short, Montminy

[6] We also think that Warfield's belief that q is, more or less directly, based on his knowledge that t_1.

as construed by Luzzi must appeal to an 'approximately' belief that is supported by a false belief, whereas we do not.

Luzzi also presents another objection to the KDF view based on the following example (due to Fitelson, 2010):

> *Fancy Watch**: I have extreme confidence in the accuracy of my fancy watch. Having lost track of the time I look carefully at my watch. I reason: 'It is exactly 2:58 p.m.'. Having learned from my logic professor earlier that day that precision entails approximation, I conclude, without any loss of confidence in my premise: 'Therefore, it is approximately 2:58 p.m.'. I know my conclusion but as it happens it is exactly 2:56 p.m., not 2:58 p.m. (Luzzi (2019: 24))

First, here are the values for our propositions:[7]

- t_1: My watch reads '2:58'.
- t_2: It is approximately 2:58pm.
- p: It is exactly 2:58pm.
- q: It is approximately 2:58pm.

Luzzi thinks that the case is problematic for the defender of KDF, because his belief that q is intuitively knowledge, but is *explicitly derived* from and thus based on the false belief that p. Now, as we pointed out in our 2014 paper, philosophers' stipulations are sometimes incoherent or at least psychologically implausible. We think that this is such a case. According to Luzzi's stipulation, his belief that it is approximately 2:58pm is caused and based *solely* by and on the explicit derivation from the belief that it is exactly 2:58pm.

However, in standard cases of the kind suggested here Luzzi believes and knows (at least tacitly) that it is approximately 2:58pm and he bases this belief (at least partly and possibly tacitly) on his knowledge that t_1—that is, on his knowledge that his watch reads '2:58'. By contrast, if there are possible cases in which a subject has the combination of attitudes that Luzzi stipulates, that subject simply will not count as knowing his or her conclusion. This can be illustrated further by the counterfactual consideration that if, in Fancy Watch*, Luzzi were told that it is not *exactly* 2:58pm, he would retain his belief that it is approximately 2:58pm, despite the fact that what he thought would ground his belief has now been shown to be false.

Luzzi (2019: 25) suggests further that the defender of KDF must "attribute[…] to cognisers an unpalatable insensitivity to the character of their reasoning." We are not sure the insensitivity in question is unpalatable: Kornblith (2002: 111), for instance, emphasizes empirical findings to the effect that human subjects are unaware of significant causal factors influencing their beliefs; and we would be unsurprised to discover similar findings in relation to beliefs arising specifically through inference. But in any case, we need not deny that Luzzi in Fancy Watch* comes to believe t_2 by deductive inference. What we deny is that he knows t_2 in virtue of any such competent deduction. On our view, there is a tacitly known truth—namely, t_1—that grounds Luzzi's knowledge that t_2. That ground for Luzzi's belief that t_2 is not just available, but it is an actual ground. This is illustrated by the counterfactual (and causal) sensitivity of the

[7] Note that in this case $t_2 = q$.

belief that t_2 to Luzzi's belief that t_1, and it's counterfactual insensitivity to his belief that *p*. In short, tacit beliefs do more psychological—and, therefore, epistemological—work than our opponents acknowledge.

In summary, defenders of KFF misidentify the causally active grounds for belief because they under-appreciate the role of subconscious psychological mechanisms and information processing. Perhaps they do so because, like some defenders of KDF, they speak of dispositional belief rather than tacit/subconscious belief.

The final objection of Luzzi's to be discussed here is based on a novel and somewhat complicated case. We quote it here in full (Luzzi 2019: 20):

> *One Short*: Let *n* be the smallest number such that observing *n* red balls drawn with replacement from a bag containing *m* balls allows one to know inductively that the next ball drawn will be red.
>
> In Gladys's office there is a large pile of visually indistinguishable bags, each known by Gladys to contain exactly *m* balls and each holding varying ratios of red balls to black balls. There are balls of no other color in the bags. Every morning at 10am, Gladys selects a bag randomly, extracts one ball from that bag and places it back in that bag. If the ball was red, she then takes the whole bag and its contents and puts it on the table in the common room. If the ball was black, she keeps the bag in her office. Each evening, if she left a bag on the table, she takes it back to her office and places it randomly among the pile of bags.
>
> Everyone in Gladys's workplace, including Sam, is aware of the above. Around noon one day, Sam walks into the common room, sees the bag and extracts exactly *n-1* balls with replacement, all of which are red. Knowing there to be *m* balls in the bag, but momentarily forgetting that Gladys has already extracted one, she reasons: (p) *I have drawn a red ball from the bag n-1 times*; so (q) *the next drawn ball will be red*.

As Luzzi points out, in this case Sam doesn't know her conclusion, since, by stipulation, she has drawn one ball too few to obtain inductive knowledge of *q*.[8] If Sam had, on Luzzi's assumptions, made one more draw, her premise *p* would be strong enough to ground her knowledge that *q*. Luzzi now objects to the defenders of KFF that Sam in fact has *dispositional* knowledge of the potential proxy premise *p*': a red ball has been drawn from the bag n times (n-1 times by me [Sam] and once by Gladys). Since Sam has this dispositional knowledge, defenders of KDF are committed to the implausible view, Luzzi claims, that Sam knows *q*.

Our response should be obvious by now: Sam may have dispositional knowledge that *p*'—that is, she may be disposed to know p'; but she doesn't have tacit knowledge that *p*'. As Luzzi (2019: 21) himself puts it elegantly, we need to distinguish between "a basis that one *actually exploits* in inferential reasoning and a basis that is *merely available* but which remains idle." In this case the available basis isn't exploited, but remains idle. The extra premise invoked may be dispositionally known, but it is not tacitly known—nor therefore is it reasoned from. The

[8] Epistemicists about vagueness will have no qualms about granting such a stipulation in general; though even they might suspect that what the number in question is varies with context, or circumstance.

subject in the case (Sam) knows the proposition that someone else drew one red ball, but is not currently recalling it, *or basing her conclusion on it*. Sure, she *could* recall it and make the inference, thereby securing knowledge of the conclusion (in the case where m=n); but she doesn't. So, in fact her belief in the conclusion is not justified. In the counterfactual case where she recalls the additional premise and then infers, her belief *is* justified. As non-dispositionalists, we are under *no* pressure to admit she's justified in the actual case.

There is one further complication with Luzzi's One Short example. We take it that Luzzi assumes that, in his case, $n \neq m$. But it is worthwhile noting that, if $n \neq m$, Luzzi is committed to the view that we can know lottery propositions (cp. Hawthorne 2004, Williamson 2000). We take this implication to be implausible—in fact, sufficiently so to dismiss the case outright; but an alternative version of the example might be designed that doesn't have this problem.

It is time to sum up. In this section, we have considered objections to KDF views, but found them wanting in relation to our own variant of the position, which relies on tacit, rather than dispositional, beliefs that serve as actual, and not merely potential, grounds for belief in the target propositions in the cases under consideration. This is as we should expect. In general, for a belief to be knowledge there must be some good, "normal" explanation of its being held (cf. Ball, 2013; Peet and Pitcovski, 2018). That normal explanation will not appeal to the falsehood that is explicitly believed, but to the relevant truths that are tacitly believed—and, in our view, to implicit reasoning from them to the targets.

Response to Sorensen

Sorensen (this volume) considers a nice example in which, he suggests, an agent generates knowledge from non-knowledge through deduction. Mini thinks, but does not know, that not all deductive arguments reason from general to particular (in brief, ¬D), from which she infers that not all deductive arguments reason from general to particular (that is, once again, ¬D). The inference has a deductively valid form (P, therefore, P); and Sorensen suggests that Mini comes to know the conclusion on the basis of this inference. The example is especially nice, however, because the argument under consideration itself verifies the conclusion: that is, it is an instance of an argument that does not proceed from general to particular; and so it makes clear that not all arguments proceed in this way.

Unfortunately for Sorensen, however, this very feature of the argument undermines his case. To see this, consider Medi, who guesses that water is H_2O, then argues on this basis to the conclusion that water is H_2O. Pursuing this line of argument of course does nothing to justify her belief in the conclusion; her inference does not result in knowledge. Or, perhaps better, suppose Maxi reasons as follows: some argument proceeds from particular to general; therefore, some argument proceeds from particular to general. Whether the sole claim involved here is analysed as particular or general, the argument does not proceed from particular to general, and Maxi cannot come to know the conclusion in the way Mini is said to by Sorensen.

These examples suggest that what justifies Mini's belief that not all deductive arguments reason from general to particular is not her inference from her belief in this very premise, but rather the fact that she appreciates that the argument is itself deductive and does not reason from general to particular. Thus, she is able to come to see the truth of the conclusion by virtue of her awareness of a fact that verifies it. Whether this process itself involves inference or insight is not entirely obvious. Does Mini simply come to see the truth of the conclusion as a

result of going through her argument (rather than on the basis of its premise)? Or does she tacitly reason (as follows)? This argument is deductive but does not reason from general to particular; therefore, not all deductive arguments reason from general to particular. Either way, Mini does not come to know something merely by way of inference from an unknown premise.[9]

Sorensen says, "[a]dmittedly, Mini must be aware of the nature of her deduction. But this information is completely endogenous to her competent deduction." (manuscript: 5) This is interesting. We agree that "[i]t is not as if [Mini] reasoned, '*Some arguments are expressed in italics, therefore, some arguments expressed in italics*'." (manuscript: 5) Reasoning is not necessarily expressed at all, let alone expressed in italics—and so one might engage in the very same reasoning with or without italics being in play. The awareness of the features of the argumentation that Mini displays is not like that—it is integral to her reasoning. And yet we think that the information in question in Mini's case is exogenous to the explicitly mentioned deduction itself: it is possible to engage in that deduction—to advance deductively from the premise of Mini's argument to its conclusion—*without* noticing the features of the argument that support Mini's knowledge.[10] Accordingly, there are two cases to consider: in one of them, Mini does not base her belief in the conclusion on this information, and she fails to come to know her conclusion; [11] in the other, she does know the conclusion, but on the basis of the further information, not just the premise.

Indeed, we can see that, in the case where knowledge is achieved, the conclusion is not believed on the basis of (belief in) the premise, since merely entertaining the premise, or assuming it for the sake of argument, would suffice. This is perhaps particularly clear with the following, earlier example: it is possible for a conclusion to entail its premise; therefore, it is possible for a conclusion to entail its premise (Sorensen, 1991: 249). Here, considering the possibility that a conclusion entails its premise, or assuming it to obtain, together with an appreciation of the form of the argument, puts the subject in a position to know that the conclusion of the argument is (actually) true. But this knowledge cannot be justified on the basis of a belief in the premise—for there is no such belief in the case in question!

"One vague suspicion," says Sorensen (in his earlier work), "is that an argument [of the kind(s) with which we are concerned] *causes* belief in the conclusion without providing a *reason* for the conclusion." (1991: 252) This is at least related to our worry about alleged cases of

---

[9] We might compare Descartes' cogito: when he forms the (justified) belief (and indeed knowledge) that he exists, does he infer this from the premise that he thinks? Some commentators deny that he does: perhaps it is a simple insight—after all, Descartes explicitly denies that the cogito is a syllogism. On the other hand, the inference from the claim that one thinks to the claim that one exists is modally, if not formally, valid: it is impossible for the premise to be true without the conclusion also being true; and since Descartes can see the truth of the premise, perhaps he likewise intuits (or singularly grasps) the validity (i.e., necessary truth-preservation) of the transition.

[10] Sorensen appears, by contrast, to be suggesting that her awareness (of the facts that her argument form is valid and that it manifests the truth of the conclusion) is, in effect, built into the requirement that her deduction is competent. We do not think that can be right. In more usual cases of competent deduction only the first of these conditions can be required, since the second of them is not met!

[11] Sorensen says, "Thanks to the narrowness of Mini's reasoning, there is nothing that sustains her belief in her conclusion beyond her competent deduction of the conclusion from the premise." (manuscript: 5). We take this to be true in the case just considered—with negative consequences for the assessment of her as knowing the conclusion.

knowledge from non-knowledge, which is that belief in the premises cannot justify belief in the conclusion. Our worry is perhaps especially clear in cases where a key premise is false: belief in a false premise, even if causally implicated in generating belief in the conclusion, cannot serve to justify that latter belief, since evidence is (and reasons are) factive. But equally, in our view, true belief that falls short of knowledge cannot serve as evidence that justifies belief in the conclusion, since E=K.

Sorensen denies the truth of the vaguely suspected claim on the grounds that the arguments in question are *rationally* persuasive. More specifically, he distinguishes ontic and propositional reasons: roughly, the former are things (the meanings of noun phrases) that are reasons, while the latter are propositions (the meanings of sentences); thus, a broken leg might be an ontic reason for a hospitalization, while that the leg is broken is a propositional reason for the same. Armed with this distinction, Sorensen claims that some arguments are ontic (rather than propositional) reasons for believing their conclusions—and that this is what makes them rationally persuasive.

We have no reason to deny any of this. It is clear that in the (good) case of Mini (in which she ends up with knowledge), her belief in the conclusion is justified because—and in virtue—of her consideration of the argument. We think that, in attending to the argument, she comes to know certain truths about it—e.g., that it is formally valid and that it does not proceed from general to particular—and that this knowledge explains her knowledge of the conclusion. So we think the argument provides Mini with propositional reasons to believe her conclusion. This is consistent with—though does not entail—the argument being an ontic reason for her to believe the conclusion. (Our view does not imply that there *are* any ontic reasons; though neither does it imply that there aren't.) Crucially, it does not commit us to saying that the premise is Mini's reason for believing her conclusion, or that it justifies it. And since Sorensen does not claim that it is—he says the *argument* is a reason for believing the conclusion in Mini's case, not that the premise is—we do not find ourselves in conflict with his views on this point.

Zooming out a little, it seems to us that Sorensen's case of Mini (this volume), and some of his earlier (1991) examples, encourage us to consider the relationship between doxastic and propositional justification. In these terms, the central lesson he wishes us to draw, it seems, is that the former cannot be understood solely in terms of the latter. Whether this is because there are ontic reasons for belief (and these include arguments), or for more prosaic reasons,[12] we agree.

[12] While the doxastic variety of justification entails the propositional (in our view), it does not follow that something (such as causation) can be added to this necessary condition to yield sufficient conditions (after all, there can be deviant causal chains).

Concluding Remarks

In our view, there is no knowledge from falsehood (KFF); though there *are* cases of knowledge *despite* falsehood (KDF). We explain how belief in the conclusions of arguments can be justified in such cases by appeal to implicit reasoning from, and tacit knowledge of, alternative 'proxy' premises that are true (Ball and Blome-Tillman, 2014). But as we showed in the first section of the current paper, it is quite uncontroversial in cognitive psychology to think that such unconscious mental states and reasoning occur. It would, of course, be problematic to appeal to merely dispositional belief in the proxies (cf. Luzzi, 2019), i.e. a disposition to consciously judge these propositions true, in explaining the presence of the knowledge in these cases (as some have admittedly attempted—cf. Montminy, 2014); but as we made clear in the second section, we do no such thing, as our tacit mental states are causally involved in the implicit reasoning we posit, and serve as actual, not merely available (hence potential, or possible), grounds for belief in the conclusion. Implicit cognition does real epistemological work, and appeal to it can help us to explain how we know what we do.

In the third section, we considered the ingenious case of Mini, who acquires the knowledge that not all deductive arguments proceed from general to particular, by engaging in a process of reasoning that begins from her unjustified belief in this very proposition (Sorensen, this volume). We argued that Mini also knows some additional information which serves to justify her final belief. But we also suggested that the case may help us to see that doxastic justification cannot be reduced to propositional justification—though possibly not for the reasons that others have indicated (Sorensen, 1991); consideration of the cognitive processes involved is likely to be required. We conclude by briefly expanding on this point.

Sorensen (this volume) suggests that Mini is in something of the position of the slave boy in the Meno, generating knowledge out of her own ignorance, through a process of reasoning—and he intimates that a commitment to the principle of counter closure, which we endorse, may explain (though not justify) the Platonic doctrine of recollection, according to which we have innate knowledge (e.g. of the axioms from which the boy can deduce the theorem he comes to recognize as true). But Sorensen takes the lesson of reflection on the Socratic method to be a different one: "Competent deduction can add justification beyond what is present in the premise." (manuscript: 9) We see no reason to abandon the principle of counter closure in this—though we welcome the encouragement to investigate deduction itself… and other cognitive processes.

Socrates, of course, asks *questions* of the slave boy. The boy, in turn, considers these questions, advances answers to them, and by recognizing dead ends as such, comes to arrive at some (mathematical) knowledge. It is through the process of pursuing this investigation—which begins with the mental state of wondering (what the answers to Socrates' questions are)—that he generates knowledge. But of course, acknowledging this fact does not yield pressure to regard the knowledge in question as based on the wondering—which has, as its object a question, not a proposition. Similarly, we see no reason to think that the cognitive processes involved in cases of KDF, significant as they may be in explaining the (doxastic) justification of the target belief, provide us with reason to regard that belief as 'epistemized' by (belief in) the premise.

References

Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological review*, *116*(4), 953.

Ball, B. (2013). Knowledge is normal belief. *Analysis*, *73*(1), 69-76.

Ball, B., & Blome-Tillmann, M. (2014). Counter closure and knowledge despite falsehood. *The Philosophical Quarterly*, *64*(257), 552-568.

Bongiorno, F., & Bortolotti, L. (2019). The Role of Unconscious Inference in Models of Delusion Formation. In *Inference and Consciousness* (pp. 74-96). Routledge.

Cao, R. (2020). New labels for old ideas: predictive processing and the interpretation of neural signals. *Review of Philosophy and Psychology*, *11*(3), 517-546.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181-204.

Evans, J. St. B. T. and Stanovich, K. E. (2013). Dual process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*, 223-241.

Graziano, M. (2018). *Dual-Process Theories of Numerical Cognition*. Springer.

Hohwy, J. (2013). *The predictive mind*. Oxford University Press.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kornblith, H. (2002). *Knowledge and its Place in Nature*. Oxford University Press.

Kruglanski, A. W., & Gigerenzer, G. (2011). Intuitive and deliberate judgements are based on common principles. *Psychological Review*, *118*(1), 97-109.

Ludwig, K., & Munroe, W. (2019). Unconscious Inference Theories of Cognitive Achievement. In *Inference and Consciousness* (pp. 15-39). Routledge.

Luzzi, F. (2019). *Knowledge from Non-Knowledge*. Cambridge University Press.

Montminy, M. (2014). Knowledge despite falsehood. *Canadian Journal of Philosophy*, *44*(3-4), 463-475.

Peet, A., & Pitcovski, E. (2018). Normal knowledge: Toward an explanation-based theory of knowledge. *The Journal of Philosophy*, *115*(3), 141-157.

Quilty-Dunn, J., & Mandelbaum, E. (2018). Inferential transitions. *Australasian Journal of Philosophy*, *96*(3), 532-547.

Quilty-Dunn, J., & Mandelbaum, E. (2019). Non-inferential transitions: Imagery and association. In *Inference and Consciousness* (pp. 151-171). Routledge.

Rescorla, M. (2019). A realist perspective on Bayesian cognitive science. In *Inference and Consciousness* (pp. 40-73). Routledge.

Sorensen, R. A. (1991). P, therefore, P' without circularity. *The Journal of Philosophy*, *88*(5), 245-266.

Sorensen, R. (this volume). Mini-Meno: How to Get More Out of Your Transmission.

Warfield, T. A. (2005). Knowledge from falsehood. *Philosophical Perspectives*, *19*, 405-416.